

TMA4267 Linear Statistical Models V2017 (L15)

Part 3: Hypothesis testing and analysis of variance
One- and two-way ANOVA [H:8.1.1]

Mette Langaas

Department of Mathematical Sciences, NTNU

To be lectured: March 10, 2017

Today: Analysis of variance (ANOVA) and analysis of covariance (ANCOVA)

- ▶ Good news: really nothing new, just linear regression where we have one or more categorical covariates.
- ▶ Bad news: a bit technical with respect to coding the covariates in the design matrix.
- ▶ Bad or good news: also tell the story of ANOVA without linear regression since that is the classical way to do things - so you will be able to recognize that this is a problem that you can solve.
- ▶ Good news: we are taking one step toward the last topic Part 4: Design of experiments.

Rothamsted Experimental Station

- founded in 1843 by John Bennet Lawes on his inherited 16th century estate, Rothamsted Manor,
 - wanted to investigate the impact of inorganic and organic fertilizers on crop yield
 - had founded a fertilizer manufacturing company in 1842
- Lawes appointed the chemist Joseph Henry Gilbert to the directorship of the chemical laboratory
- the two began a series of field experiments to examine the effects of inorganic fertilizers and organic manures on the nutrition and yield of a number of important crops



<http://www.stats.uwo.ca/faculty/bellhouse/stat499lecture13.pdf>

The Broadbalk Field Trial at Rothamsted

- this was the first field trial started by Lawes and Gilbert
- began in 1843
- purpose was to investigate the relative importance of different plant nutrients (N, P, K, Na, Mg) on grain yield of winter wheat
- weeds were controlled by hand hoeing and fallowing
 - now some herbicides are used
- The experiment continues to this day



<http://www.stats.uwo.ca/faculty/bellhouse/stat499lecture13.pdf>

Concrete aggregates example



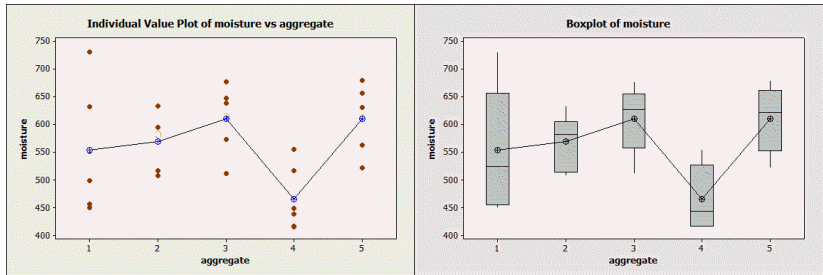
- ▶ Aggregates are inert granular materials such as sand, gravel, or crushed stone that, along with water and portland cement, are an essential ingredient in concrete.
- ▶ For a good concrete mix, aggregates need to be clean, hard, strong particles free of absorbed chemicals or coatings of clay and other fine materials that could cause the deterioration of concrete.
- ▶ We could like to examine 5 different aggregates, and measure the absorption of moisture after 48hrs exposure (to moisture).
- ▶ A total of 6 samples are tested for each aggregate.
- ▶ Research question: Is there a difference between the aggregates with respect to absorption of moisture?

Concrete aggregates data

Aggregate:	1	2	3	4	5	
	551	595	639	417	563	
	457	580	615	449	631	
	450	508	511	517	522	
	731	583	573	438	613	
	499	633	648	415	656	
	632	517	677	555	679	
Total	3320	3416	3663	2791	3664	16,854
Mean	553.33	569.33	610.50	465.17	610.67	561.80

Table 13.1 of Walepole, Myers, Myers, Ye: Statistics for Engineers and Scientists – our textbook from the introductory TMA4240/TMA4245 Statistics course.

Concrete aggregates example



One-way Analysis of Variance (ANOVA)

Model

$$Y_{ij} = \mu_i + \varepsilon_{ij} \text{ for } i = 1, 2, \dots, p \text{ and } j = 1, 2, \dots, n_i$$

alternative parameterization

$$Y_{ij} = \mu + \alpha_i + \varepsilon_{ij}$$

The sample sizes for each group, n_i may vary. $\varepsilon_{ij} \sim N(0, \sigma^2)$. Let $n = \sum_{i=1}^p n_i$ be the total number of observations.

Aim: look at parameter estimates and test if there is any difference between the groups.

How can that be done using our linear regression model?

Concrete aggregates data

```
# means for each recipe
> means=
    aggregate(ds,by=list(ds$aggregate),FUN=mean)$moisture
> grandmean=mean(ds$moisture)
> grandmean
[1] 561.8
> alphas=means-grandmean
> alphas
[1] -8.466667    7.533333   48.700000  -96.633333   48.866667
```

Concrete aggregates data

```
# the same with regression
> options(contrasts=c("contr.sum","contr.sum"))
> obj <-lm(moisture~as.factor(aggregate),data=ds)
> summary(obj)
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	561.800	12.859	43.688	< 2e-16	***
as.factor(aggregate)1	-8.467	25.719	-0.329	0.744743	
as.factor(aggregate)2	7.533	25.719	0.293	0.772005	
as.factor(aggregate)3	48.700	25.719	1.894	0.069910	.
as.factor(aggregate)4	-96.633	25.719	-3.757	0.000921	***

Concrete aggregates data

```
#comparing means and regression estimates
>cbind(c(grandmean,alphas),
      c(obj$coefficients,-sum(obj$coefficients[2:5])))
```

	[,1]	[,2]
(Intercept)	561.800000	561.800000
as.factor(aggregate)1	-8.466667	-8.466667
as.factor(aggregate)2	7.533333	7.533333
as.factor(aggregate)3	48.700000	48.700000
as.factor(aggregate)4	-96.633333	-96.633333
	48.866667	48.866667

Run R code from course lectures tab for model matrix.

Concrete aggregates data (1)

```
# checking manually with linear hypotheses
r=4
C=cbind(rep(0,r),diag(r))
d=matrix(rep(0,r),ncol=1)
betahat=matrix(obj$coefficients,ncol=1)
sigma2hat=summary(obj)$sigma^2
Fobs=(t(C%*%betahat-d)%*%
solve(C%*%solve(t(X)%*%X)%*%t(C))%*%
(C%*%betahat-d))/(r*sigma2hat)
> Fobs
      [,1]
[1,] 4.301536
> 1-pf(Fobs,r,n-r-1)
      [,1]
[1,] 0.008751641
```

Concrete aggregates data (2)

```
> fitA=obj  
> fitB=lm(moisture~1,data=aggregates)  
> anova(fitA,fitB)
```

Analysis of Variance Table

Model 1: moisture ~ as.factor(aggregate)

Model 2: moisture ~ 1

	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)
1	25	124020				
2	29	209377	-4	-85356	4.3015	0.008752 **

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Concrete aggregates data (3)

```
# performing ANOVA using method anova -
```

```
> anova(obj)
```

```
Analysis of Variance Table
```

```
Response: moisture
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
as.factor(aggregate)	4	85356	21339.1	4.3015	0.008752 **
Residuals	25	124020	4960.8		

One factor: unequal sample sizes

Classical formulation with ANOVA decomposition

$$Y_{ij} - Y_{..} = (Y_{ij} - Y_{i.}) + (Y_{i.} - Y_{..})$$

$$\sum_{i=1}^p \sum_{j=1}^{n_i} (Y_{ij} - Y_{..})^2 = \sum_{i=1}^p \sum_{j=1}^{n_i} (Y_{ij} - Y_{i.})^2 + \sum_{i=1}^p \sum_{j=1}^{n_i} (Y_{i.} - Y_{..})^2$$

$$\sum_{i=1}^p \sum_{j=1}^{n_i} (Y_{ij} - Y_{..})^2 = \sum_{i=1}^p \sum_{j=1}^{n_i} (Y_{ij} - Y_{i.})^2 + \sum_{i=1}^p n_i (Y_{i.} - Y_{..})^2$$

$$\text{SST} = \text{SSE} + \text{SSA}$$

One factor: unequal sample sizes

ANOVA decomposition: what happened to the cross-term?

$$2 \sum_{i=1}^p \sum_{j=1}^{n_i} (Y_{ij} - Y_{i.})(Y_{i.} - Y_{..}) = 2 \sum_{i=1}^p (Y_{i.} - Y_{..}) \sum_{j=1}^{n_i} (Y_{ij} - Y_{i.}) = 0$$
$$\sum_{j=1}^{n_i} (Y_{ij} - Y_{i.}) = \sum_{j=1}^{n_i} Y_{ij} - \sum_{j=1}^{n_i} Y_{i.} = n_i Y_{i.} - n_i Y_{i.} = 0$$

One factor: unequal sample sizes

$H_0 : \mu_1 = \mu_2 = \cdots = \mu_p = 0$ vs. H_1 : At least one pair of μ_i different
is then tested based on

$$F = \frac{\frac{SSA}{p-1}}{\frac{SSE}{n-p}}$$

Where H_0 is rejected if $f_{\text{obs}} > f_{\alpha, (p-1), (n-p)}$.


Machine example

- ▶ Response: time (s) spent to assemble a product.
- ▶ Factor: this is done by four different machines;
 M_1, M_2, M_3, M_4 .
- ▶ Question: Do the machines perform at the same mean rate of speed?

TABLE 13.12 Time, in Seconds, to Assemble Product

Machine	Operator:	1	2	3	4	5	6	Total
1		42.5	39.3	39.6	39.9	42.9	43.6	247.8
2		39.8	40.1	40.5	42.3	42.5	43.1	248.3
3		40.2	40.5	41.3	43.4	44.9	45.1	255.4
4		41.3	42.2	43.5	44.2	45.9	42.3	259.4
Total		163.8	162.1	164.9	169.8	176.2	174.1	1010.9

Data from Walepole, Myers, Myers, Ye: "Statistics for Engineers and Scientists", Example 13.6= our TMA4245/40 textbook.



One factor ANOVA

```
> options(contrasts=c("contr.sum","contr.sum"))
> fit <- lm(time~as.factor(machine),data=dsmat)
> summary(fit)
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	42.1208	0.3706	113.647	<2e-16	***
as.factor(machine)1	-0.8208	0.6419	-1.279	0.216	
as.factor(machine)2	-0.7375	0.6419	-1.149	0.264	
as.factor(machine)3	0.4458	0.6419	0.695	0.495	

Residual standard error: 1.816 on 20 degrees of freedom

Multiple R-squared: 0.1945, Adjusted R-squared: 0.07372



F-statistic: 1.61 on 3 and 20 DF, p-value: 0.2186

```
> anova(fit)
```

Response: time


	Df	Sum Sq	Mean Sq	F value	Pr(>F)
as.factor(machine)	3	15.925	5.3082	1.6101	0.2186
Residuals	20	65.935	3.2968		

Residuals



Machine example: operators

- ▶ The 6 repeated measurements for each machine was in fact made by 6 different operators.
- ▶ The operation of the machines requires physical dexterity and differences among the operators in the speed with which they operate the machines is anticipated.
- ▶ All of the 6 operators have operated all the 4 machines, and the machines were assigned in random order to the operators=*randomized complete block design*.
- ▶ By including a blocking factor called Operator, we will reduce the variation in the experiment that is due to random error. Thus, we reduce variation due to *anticipated factors*.
- ▶ By randomizing the order the machines were assigned to the operators we aim to reduce the variation due to *unanticipated factors*.



Model and Sums of squares

Model

$$Y_{ij} = \mu + \alpha_i + \gamma_j + \varepsilon_{ij} \text{ for } i = 1, 2, \dots, r \text{ and } j = 1, 2, \dots, s$$

Sums of Squares Identity

$$Y_{ij} = Y_{..} + (Y_{i.} - Y_{..}) + (Y_{.j} - Y_{..}) + (Y_{ij} - Y_{i.} - Y_{.j} + Y_{..})$$

$$\begin{aligned} \sum_{i=1}^r \sum_{j=1}^s (Y_{ij} - Y_{..})^2 &= s \sum_{i=1}^r (Y_{i.} - Y_{..})^2 + r \sum_{j=1}^s (Y_{.j} - Y_{..})^2 \\ &\quad + \sum_{i=1}^r \sum_{j=1}^s (Y_{ij} - Y_{i.} - Y_{.j} + Y_{..})^2 \end{aligned}$$

$$\text{SST} = \text{SSA} + \text{SSB} + \text{SSE}$$

$$r \cdot s - 1 = (r - 1) + (s - 1) + (r - 1)(s - 1)$$

Effect of factor A:

$H_0 : \alpha_1 = \alpha_2 = \cdots = \alpha_r = 0$ vs. $H_1 : \text{At least one } \alpha_i \text{ different from } 0$

is then tested based on

$$F_1 = \frac{\frac{SSA}{r-1}}{\frac{SSE}{(r-1)(s-1)}}$$

Where H_0 is rejected if $f_1 > f_{\alpha, (r-1), (r-1)(s-1)}$.

Block effect present?

$H_0 : \gamma_1 = \gamma_2 = \cdots = \gamma_s = 0$ vs. $H_1 : \text{At least one } \gamma_j \text{ different from } 0$

is then tested based on

$$F_2 = \frac{\frac{SSB}{s-1}}{\frac{SSE}{(r-1)(s-1)}}$$

Where H_0 is rejected if $f_2 > f_{\alpha, (s-1), (r-1)(s-1)}$.

RCBD ANOVA

```
> fit2 <- lm(time~as.factor(machine)+as.factor(operator),  
data=dsmat)
```



```
> anova(fit2)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
as.factor(machine)	3	15.925	5.3082	3.3388	0.047904	*
as.factor(operator)	5	42.087	8.4174	5.2944	0.005328	**
Residuals	15	23.848	1.5899			

Effect of operator with linear hypotheses

```
fit2 <- lm(time~as.factor(machine)+as.factor(operator),
data=dsmat)
r=5
C=cbind(rep(0,5),rep(0,5),rep(0,5),rep(0,5),diag(5))
d=matrix(rep(0,r),ncol=1)
betahat=matrix(fit2$coefficients,ncol=1)
X=model.matrix(fit2)
sigma2hat=summary(fit2)$sigma^2
Fobs=(t(C%*%betahat-d)%*%solve(C%*%solve(t(X)%*%X)%*%t(C))
%*%(C%*%betahat-d))/(r*sigma2hat)
> Fobs
      [,1]
[1,] 5.294435
> 1-pf(Fobs,r,n-dim(C)[2])
      [,1]
[1,] 0.005327541
```

Residuals



A second look at the RCBD: additive effects

Previously, randomized complete block design (RCBD) with the machine example:

$$Y_{ij} = \mu + \alpha_i + \gamma_j + \varepsilon_{ij}$$

where $\sum_{i=1}^r \alpha_i = 0$ and $\sum_{j=0}^s \gamma_j = 0$.

This is called *additive effects of treatment and blocks*.

- ▶ This means that if we compare two operators there is a constant difference in time to assemble the product,
- ▶ or, if we compare machines, these are ranked in the same order of (wrt time) for each operator.

Estimates

$$\hat{\mu} = 42.1208$$

$$\hat{\alpha}_1 = -0.8208$$

$$\hat{\alpha}_2 = -0.7375$$

$$\hat{\alpha}_3 = 0.4458$$

$$\hat{\alpha}_4 = 1.1125$$

$$\hat{\gamma}_1 = -1.1708$$


$$\hat{\gamma}_2 = -1.5958$$

$$\hat{\gamma}_3 = -0.8958$$

$$\hat{\gamma}_4 = 0.3292$$

$$\hat{\gamma}_5 = 1.9292$$

$$\hat{\gamma}_6 = 1.404167$$



Estimates

$$\hat{\mu} = 42.1208$$

$$\hat{\alpha}_1 = -0.8208$$

$$\hat{\alpha}_2 = -0.7375$$

$$\hat{\alpha}_3 = 0.4458$$

$$\hat{\alpha}_4 = 1.1125$$

$$\hat{\gamma}_1 = -1.1708$$


$$\hat{\gamma}_2 = -1.5958$$


$$\hat{\gamma}_3 = -0.8958$$

$$\hat{\gamma}_4 = 0.3292$$

$$\hat{\gamma}_5 = 1.9292$$

$$\hat{\gamma}_6 = 1.404167$$





Interaction effect?

But, it could be interactions present. What if one of the operators really could not manage one of the machines?

Model with interaction between treatment and block:

$$Y_{ij} = \mu + \alpha_i + \gamma_j + (\alpha\gamma)_{ij} + \varepsilon_{ij}$$

where $\sum_{i=1}^r (\alpha\gamma)_{ij} = \sum_{j=1}^s (\alpha\gamma)_{ij} = 0$ (for all i and j) in addition to $\sum_{i=1}^r \alpha_i = 0$ and $\sum_{j=1}^s \gamma_j = 0$.

But, since we only have one observation for each combination of i and j , we can not separate $(\alpha\gamma)_{ij}$ and ε_{ij} .

Interaction effect?

$$SSE = \sum_{i=1}^r \sum_{j=1}^s (Y_{ij} - Y_{.i} - Y_{j.} + Y_{..})^2$$

$$E\left(\frac{SSE}{(r-1)(s-1)}\right) = \sigma^2 + \frac{\sum_{i=1}^r \sum_{j=1}^s (\alpha\gamma)_{ij}^2}{(s-1)(r-1)}$$

A large value of SSE will either mean that we have an interaction term present, or that σ^2 is large. We can not assess interaction in a RCBD. We need more than one observation for each observation to distinguish between $(\alpha\gamma)_{ij}$ and ε_{ij} .

Age and memory

- ▶ Why do older people often seem not to remember things as well as younger people? Do they not pay attention? Do they just not process the material as thoroughly?
- ▶ One theory regarding memory is that verbal material is remembered as a function of the degree to which it was processed when it was initially presented.
- ▶ Eysenck (1974) randomly assigned 50 younger subjects and 50 older (between 55 and 65 years old) to one of five learning groups.
- ▶ After the subjects had gone through a list of 27 items three times they were asked to write down all the words they could remember.

Eysenck study of recall of older and younger subjects under conditions of differential processing, Eysenck (1974) and presented in Howell (1999).


The Age and Memory data set

- ▶ Number of words recalled: After the subjects had gone through the list of 27 items three times they were asked to write down all the words they could remember.
- ▶ Age: Younger (18-30) and Older (55-65).

The Age and Memory data set: Process

- ▶ The Counting group was asked to read through a list of words and count the number of letters in each word. This involved the lowest level of processing.
- ▶ The Rhyming group was asked to read each word and think of a word that rhymed with it.
- ▶ The Adjective group was asked to give an adjective that could reasonably be used to modify each word in the list.
- ▶ The Imagery group was instructed to form vivid images of each word, and this was assumed to require the deepest level of processing.
None of these four groups was told they would later be asked to recall the items.
- ▶ Finally, the Intentional group was asked to memorize the words for later recall.

Data taken from: <http://www.statsci.org/data/general/eysenck.html>



Y=younger (blue), O=older (red), A=adjective, C=counting, Im=Imagery, In=intentional, R=rythming.

Eysenck ANOVA

```
> res <- lm(Words~Age*Process)
> summary(res)
Call:
lm(formula = Words ~ Age * Process)
```

Residuals:

Min	1Q	Median	3Q	Max
-7.0	-1.6	-0.5	2.0	9.6

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	11.6100	0.2833	40.982	< 2e-16 ***
Age1	-1.5500	0.2833	-5.471	3.98e-07 ***
Process1	1.2900	0.5666	2.277	0.025170 *
Process2	-4.8600	0.5666	-8.578	2.60e-13 ***
Process3	3.8900	0.5666	6.866	8.24e-10 ***
Process4	4.0400	0.5666	7.130	2.43e-10 ***
Age1:Process1	-0.3500	0.5666	-0.618	0.538312
Age1:Process2	1.8000	0.5666	3.177	0.002040 **
Age1:Process3	-0.5500	0.5666	-0.971	0.334288
Age1:Process4	-2.1000	0.5666	-3.706	0.000363 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 2.833 on 90 degrees of freedom
Multiple R-squared: 0.7293, Adjusted R-squared: 0.7022
F-statistic: 26.93 on 9 and 90 DF, p-value: < 2.2e-16

Eysenck model matrix

```
> X=model.matrix(res)
> X[c(1,11,21,31,41,51,61,71,81,91),]
      (Intercept) Age1 Process1 Process2 Process3 Process4 Age1:Process1
1              1    -1         0         1         0         0         0
11             1    -1        -1        -1        -1        -1        1
21             1    -1         1         0         0         0       -1
31             1    -1         0         0         1         0         0
41             1    -1         0         0         0         1         0
51             1     1         0         1         0         0         0
61             1     1        -1        -1        -1        -1       -1
71             1     1         1         0         0         0         1
81             1     1         0         0         1         0         0
91             1     1         0         0         0         1         0
      Age1:Process2 Age1:Process3 Age1:Process4
1              -1         0         0
11             1         1         1
21             0         0         0
31             0        -1         0
41             0         0        -1
51             1         0         0
61            -1        -1        -1
71             0         0         0
81             0         1         0
91             0         0         1
```


Model and Sums of Squares

Model:

$$Y_{ijk} = \mu + \alpha_i + \gamma_j + (\alpha\gamma)_{ij} + \varepsilon_{ijk}$$

for $i = 1, 2, \dots, r$ and $j = 1, 2, \dots, s$ and $k = 1, \dots, m$

$$\varepsilon_{ijk} \sim N(0, \sigma^2)$$

Two-way ANOVA questions

There are three main questions that we might ask in two-way ANOVA:

- ▶ Does the response variable depend on Factor A?
- ▶ Does the response variable depend on Factor B?
- ▶ Does the response variable depend on Factor A differently for different values of Factor B, and vice versa?

All of these questions can be answered using hypothesis tests, first we test the interaction.

Effect of interaction AB

$H_0^A : (\alpha\gamma)_{11} = (\alpha\gamma)_{12} = \cdots = (\alpha\gamma)_{rs} = 0$ vs.

H_1 : At least one $(\alpha\gamma)_{ij}$ different from 0

is then tested based on

$$F_3 = \frac{\frac{SS(AB)}{(r-1)(s-1)}}{\frac{SSE}{rs(m-1)}}$$

Where H_0 is rejected if $f_3 > f_{\alpha, (r-1)(s-1), rs(m-1)}$.

What do we do after testing for interaction?

- ▶ If the interaction is significant (we reject H_0^{AB}).
 - ▶ Then it is not recommended to test for main effects (that is, the marginal contributions of the two factors A and B separately). This is since the interpretation of the marginal “main effect” is unclear in the presence of interaction. How can we “separate out” the effect of A from the interaction?
 - ▶ Instead, it is usually preferable to examine contrasts in the treatment combinations.
- ▶ If the interaction is not found to be significant (do not reject H_0^{AB}).
 - ▶ We are then interested in the main effects. These can now be tested within the complete model.

Effect of factor A:

$H_0^A : \alpha_1 = \alpha_2 = \cdots = \alpha_r = 0$ vs. $H_1 : \text{At least one } \alpha_i \text{ different from } 0$

is then tested based on

$$F_1 = \frac{\frac{SSA}{r-1}}{\frac{SSE}{rs(m-1)}}$$

Where H_0^A is rejected if $f_1 > f_{\alpha, (r-1), rs(m-1)}$.

Effect of factor B:

$H_0^B : \gamma_1 = \gamma_2 = \cdots = \gamma_s = 0$ vs. $H_1 : \text{At least one } \gamma_i \text{ different from } 0$

is then tested based on

$$F_2 = \frac{\frac{SSB}{s-1}}{\frac{SSE}{rs(m-1)}}$$

Where H_0^B is rejected if $f_2 > f_{\alpha, (s-1), sn(m-1)}$.

Eysenck ANOVA

```
> res <- lm(Words~Age*Process)
> summary(res)
Call:
lm(formula = Words ~ Age * Process)
```

Residuals:

Min	1Q	Median	3Q	Max
-7.0	-1.6	-0.5	2.0	9.6

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	11.6100	0.2833	40.982	< 2e-16 ***
Age1	-1.5500	0.2833	-5.471	3.98e-07 ***
Process1	1.2900	0.5666	2.277	0.025170 *
Process2	-4.8600	0.5666	-8.578	2.60e-13 ***
Process3	3.8900	0.5666	6.866	8.24e-10 ***
Process4	4.0400	0.5666	7.130	2.43e-10 ***
Age1:Process1	-0.3500	0.5666	-0.618	0.538312
Age1:Process2	1.8000	0.5666	3.177	0.002040 **
Age1:Process3	-0.5500	0.5666	-0.971	0.334288
Age1:Process4	-2.1000	0.5666	-3.706	0.000363 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 2.833 on 90 degrees of freedom
Multiple R-squared: 0.7293, Adjusted R-squared: 0.7022
F-statistic: 26.93 on 9 and 90 DF, p-value: < 2.2e-16

Eysenck ANOVA

```
> res <- lm(Words~Age*Process)
```

```
> anova(res)
```

Analysis of Variance Table

Response: Words

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
Age	1	240.25	240.25	29.9356	3.981e-07	***
Process	4	1514.94	378.74	47.1911	< 2.2e-16	***
Age:Process	4	190.30	47.58	5.9279	0.0002793	***
Residuals	90	722.30	8.03			

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Next: maybe want to compare different combinations of age and process? Then, easiest to just combine the two factors into a new joint factor and skip the intercept.

Summing up

Topic today: the one-way and two-way ANOVA models.

- ▶ Classical formulation has focus on comparing sums of squares.
- ▶ We don't have to prove the classical results because we instead fit the ANOVA model using linear regression with effect coding of covariates.
- ▶ It is important to plot results and to understand when an interaction term is needed.
- ▶ To test ANOVA hypotheses we use linear hypotheses in the regression – where we automatically have theoretical results for F-distributions.
- ▶ We will meet linear regression models with k factors with two levels each in Part 4: Design of Experiments (DOE).