Faculty of Life Sciences

Comparison of groups

One-way ANOVA

Ib Skovgaard & Claus Ekstrøm
E-mail: ims@life.ku.dk

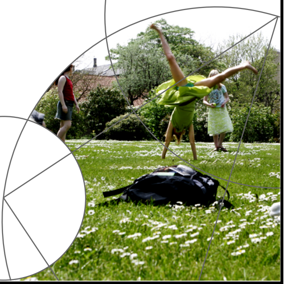

Program

- Comparison of two groups: overview
- Comparison of more than two groups: one-way ANOVA
  - Data: antibiotics and decomposition of organic material
  - Statistical model
  - Estimation and confidence intervals
  - Comparison of the groups (test)
  - Pairwise comparisons

Comparison of two samples: overview

	x, y indep.?	Same sd.?	R
Example?	Yes	Yes	`t.test(x,y, var.equal=T)`
Example?	Yes	No	`t.test(x,y)`
Example?	No		`t.test(x,y, paired=T)`

For comparison of two groups some form of t-test may be used.

What about comparison of three or more groups?
One-way ANOVA!

Antibiotics and the decomposition of organic material

Data

- Five types antibiotics and a control
- 36 heifers allocated to 6 treatment groups. Feed added antibiotics
- Dung suspended in bags in the ground
- Amount of organic material measured after 8 weeks.
- For spiramycin: only four measurements available

Problem:

- Does the antibiotics affect the decomposition of organic material?

Group means and standard deviations

Type	n_j	$\bar{y}_j$	s_j
Control	6	2.603	0.119
α-cyperm.	6	2.895	0.117
Enrofloxacin	6	2.710	0.162
Fenbendaz.	6	2.833	0.124
Ivermectin	6	3.002	0.109
Spiramycin	4	2.855	0.054

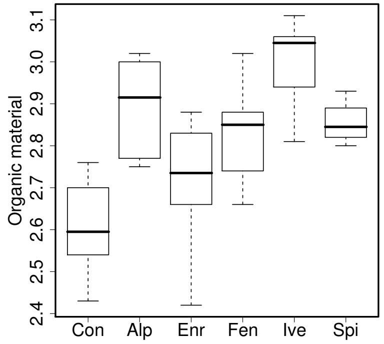

Pooled estimate of the standard deviation:

$$s = \sqrt{\frac{1}{28}(5 \cdot s_1^2 + \cdots + 3 \cdot s_6^2)} = \sqrt{\frac{1}{34-6}\sum_{i=1}^{n}(y_i - \bar{y}_{g(i)})^2} = 0.1217$$

Statistical model

Recall that $g(i)$ denotes the group for observation i. For example

$$g(1) = \cdots = g(6) = \text{control}, \quad g(31) = \cdots = g(34) = \text{Spiramycin}$$
$$g(1) = \cdots = g(6) = 1, \quad g(31) = \cdots = g(34) = 6.$$

Statistical model: $y_1, \ldots, y_{34}$ are independent and

$$y_i \sim N(\alpha_{g(i)}, \sigma^2)$$

Parameters: $\alpha_1, \ldots, \alpha_6$ and σ.

Equivalently:

$$y_i = \alpha_{g(i)} + e_i, \quad e_1, \ldots, e_{34} \sim N(0, \sigma^2) \text{ independent}$$

Statistical model

$$y_i = \alpha_{g(i)} + e_i, \quad e_1, \ldots, e_n \sim N(0, \sigma^2) \text{ independent}$$

Assumptions:

- Normal distribution?
- Mean?
- Standard deviation?
- Independence?

Estimation and confidence intervals

Statistical model:

$$y_i = \alpha_{g(i)} + e_i, \quad e_1, \ldots, e_n \sim N(0, \sigma^2) \text{ independent}$$

Parameters: $\alpha_1, \ldots, \alpha_k$ and σ. In particular, we are interested in differences, $\alpha_j - \alpha_l$!

Estimates and standard errors:

$$\hat{\alpha}_j = \bar{y}_j; \quad \mathrm{SE}(\hat{\alpha}_j) = s\sqrt{1/n_j} = s/\sqrt{n_j}$$
$$\hat{\alpha}_j - \hat{\alpha}_l = \bar{y}_j - \bar{y}_l; \quad \mathrm{SE}(\hat{\alpha}_j - \hat{\alpha}_l) = s\sqrt{1/n_j + 1/n_l}$$
$$\hat{\sigma} = s$$

Confidence intervals from the usual recipe:

$$\text{estimate} \pm t_{0.975, n-k} \cdot \mathrm{SE}(\text{estimate})$$

NB. The pooled s is used, also when comparing two groups!

One-way ANOVA in R

Fitting the model:

```
> model1 <- lm(org~factor(type))
> summary(model1)
```

R chooses a reference group — the first in alphabetic order — and estimates changes compared to that group.

We prefer the control group as the reference group:

```
> type <- relevel(type, ref="Control")
> model1 <- lm(org~factor(type))
> summary(model1)
```


One-way ANOVA in R

Output from summary(model1):

```
Coefficients:
                      Estimate Std. Error t value Pr(>|t|)
(Intercept)            2.60333    0.04970  52.379  < 2e-16 ***
factor(type)Alfacyp    0.29167    0.07029   4.150 0.000281 ***
factor(type)Enroflox   0.10667    0.07029   1.518 0.140338
factor(type)Fenbenda   0.23000    0.07029   3.272 0.002834 **
factor(type)Ivermect   0.39833    0.07029   5.667  4.5e-06 ***
factor(type)Spiramyc   0.25167    0.07858   3.202 0.003384 **

Residual standard error: 0.1217 on 28 degrees of freedom
```

Interpretations:

- Estimate and CI for α_{cont}, $\alpha_{\text{Fenb}} - \alpha_{\text{cont}}$ and α_{Fenb}? Estimat for σ?
- Why are the SE's not the same?

One-way ANOVA in R

If we prefer to see the group means:

```
> model2 <- lm(org~factor(type)-1)
> summary(model2)
Coefficients:
                    Estimate Std. Error t value Pr(>|t|)
factor(type)Control   2.60333    0.04970   52.38   <2e-16 ***
factor(type)Alfacyp   2.89500    0.04970   58.25   <2e-16 ***
factor(type)Enroflox  2.71000    0.04970   54.53   <2e-16 ***
factor(type)Fenbenda  2.83333    0.04970   57.01   <2e-16 ***
factor(type)Ivermect  3.00167    0.04970   60.39   <2e-16 ***
factor(type)Spiramyc  2.85500    0.06087   46.90   <2e-16 ***

Residual standard error: 0.1217 on 28 degrees of freedom
```


Hypothesis. Variation within and between groups

Hypothesis, $H_0 : \alpha_1 = \cdots = \alpha_k$.

Alternative, H_A : at least two α's are different.

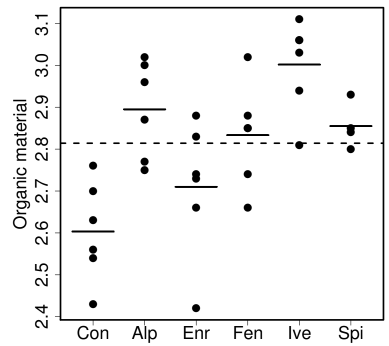

- Variation within groups — points around the lines
  $$SS_e = \sum_{i=1}^{n} \left(y_i - \bar{y}_{g(i)}\right)^2$$
- Variation between groups — Lines around the dashed line
  $$SS_{\text{grp}} = \sum_{j=1}^{k} n_j (\bar{y}_j - \bar{y})^2$$
- Test statistic
  $$F = \frac{MS_{\text{grp}}}{MS_e} = \frac{SS_{\text{grp}}/(k-1)}{SS_e/(n-k)}$$

Comparison af alle groupsne

Do not use `model2` for this — only `model1`

```
> anova(model1)
             Df  Sum Sq Mean Sq F value    Pr(>F)
factor(type)  5 0.59082 0.11816  7.9726 8.953e-05 ***
Residuals    28 0.41500 0.01482
```

Test statistic

$$F = \frac{\mathrm{MS}_{\mathrm{grp}}}{\mathrm{MS}_e} = \frac{\mathrm{SS}_{\mathrm{grp}}/(k-1)}{\mathrm{SS}_e/(n-k)}$$

Large values of F are in disagreement with the hypothesis. Hence, the p-value is

$$p = P(F \geq F_{\mathrm{obs}}) = P(F \geq 7.97) = 0.00009$$

There is overwhelming evidence that the hypothesis is not true.

How did we get the p-value?

The F-distribution

If the hypothesis is true, then the F-test statistic is F-distributed with $(k-1, n-k)$ degrees of freedom.

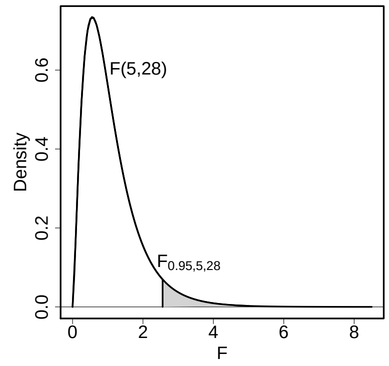

$p = P(F \geq 7.97) = 0.00009$

F-probabilities and quantiles in R:

```
> pf(7.97, df1=5, df2=28)
[1] 0.9999102
> qf(0.95, df1=5, df2=28)
[1] 2.558128
```


Pairwise comparisons

Suppose we want to compare the control group (1) with the Fenbendazole group (4): $\alpha_4 - \alpha_1$.

Estimate and its standard error:

$$\hat{\alpha}_4 - \hat{\alpha}_1 = 2.833; \quad \mathrm{SE}(\hat{\alpha}_4 - \hat{\alpha}_1) = 0.07029$$

- Confidence interval for $\alpha_4 - \alpha_1$?
- Test for the hypothesis $H_0 : \alpha_1 = \alpha_4$?
- Do all the groups differ significantly from the control group?

LSD-value: least significant difference

A confidence interval for the difference $\alpha_j - \alpha_l$ is

$$\hat{\alpha}_j - \hat{\alpha}_l \pm \mathrm{LSD}$$

where

$$\mathrm{LSD}_{j,l} = t_{0.975,n-k} \cdot \mathrm{SE}(\hat{\alpha}_j - \hat{\alpha}_l) = t_{0.975,n-k} \cdot s \cdot \sqrt{1/n_j + 1/n_l}.$$

A t-test for the hypothesis that the difference is zero uses the test statistic

$$T = \frac{|\hat{\alpha}_j - \hat{\alpha}_l|}{\mathrm{SE}(\hat{\alpha}_j - \hat{\alpha}_l)}$$

which is t-distributed with $n-k$ degrees of freedom.
LSD for control and fenbend.: $2.048 \cdot 0.1217 \cdot \sqrt{1/6 + 1/6} = 0.144$

If all group sizes are the same, then so are the LSD-values:

$$\mathrm{LSD} = t_{0.975,n-k} \cdot \mathrm{SE}(\hat{\alpha}_j - \hat{\alpha}_l) = t_{0.975,n-k} \cdot s \cdot \sqrt{2/n'}$$

Conclusion

Different effects of the different types has been shown with high degree of certainty ($p < 0.0001$)

For all types except Enrofloxacin the amount of organic material is significantly higher than for the control group.

These statements should be supplemented by estimates and confidence intervals for α's and/or for differences to the control group.

Multiple comparisons

Any time we make a test a type I error may occur. The risk depends on the level of significance — often 5%.

One test: risk of type I error: 5%

By m independent tests:

$$1 - 0.95^m$$

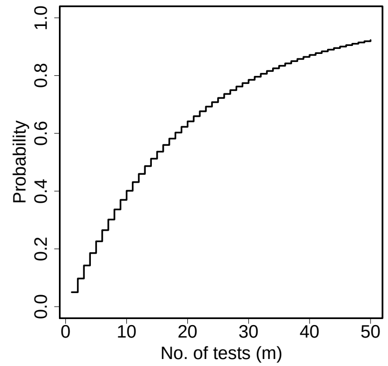

Summary: one-way ANOVA

- Statistical model: normal distribution with same SD in the groups; independence
- Estimation: group means and pooled SD
- Confidence interval: $\text{estimat} \pm t_{0.975,n-k} \cdot \text{SE(estimate)}$
- Hypothesis of equal group means tested by $F = \text{MS}_{\text{grp}}/\text{MS}_e$.
- Pairwise comparisons conducted "within" the model, using all the observations to estimate the SD.

With only two groups, t-tests suffice. Different versions:

- Paired or unpaired?
- If unpaired: same SD or not?

Lecture summary: main points

- One-way ANOVA
- Assumptions for one-way ANOVA
- Hypotheses for one-way ANOVA
- Test statistic and the F-fordelingen