STAT22200 Spring 2014 Chapter 4

Yibi Huang

April 24, 2014

Chapter 4 Contrasts

Analysis of Factor Level Means

- When F-test is significant; there exist differences among the means. Now what?
- Want to determine which means are different and identify treatments statistically of the same effect

Visual Assessment

- Can often get an idea by looking at plots
  - Side-by-side Box Plots
  - Scatter plots

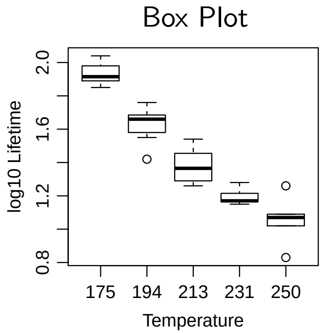

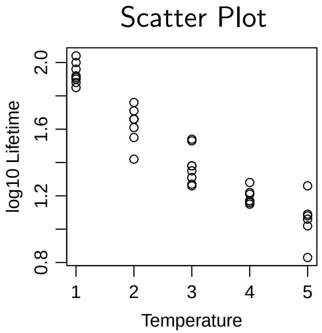

- These plots do not give any information about the precision of the estimates. Need to consider standard errors.

Confidence Interval for Treatment Means

Recall the models for a CRD experiment:

$$\begin{aligned} y_{ij} &= \mu_i + \varepsilon_{ij} && \text{(Means Model)} \\ &= \mu + \alpha_i + \varepsilon_{ij} && \text{(Effects Model)} \end{aligned}$$

where ε_{ij}'s are i.i.d. $N(0, \sigma^2)$

In the means model, μ_i is estimated by the **sample mean** of treatment group i

$$\widehat{\mu}_i = \overline{y}_{i\bullet} = \frac{1}{n_i} \sum_{j=1}^{n_i} y_{ij} = \frac{1}{n_i}(y_{i1} + \ldots + y_{in_i}).$$

$$\begin{aligned} \mathbb{E}(\overline{y}_{i\bullet}) = \mathbb{E}\left[\frac{1}{n_i}(y_{i1} + \ldots + y_{in_i})\right] &= \frac{1}{n_i}[\mathbb{E}(y_{i1}) + \ldots + \mathbb{E}(y_{in_i})] \\ &= \frac{1}{n_i}(\underbrace{\mu_i + \ldots + \mu_i}_{n_i \text{ copies}}) = \mu_i. \end{aligned}$$

So $\overline{y}_{i\bullet}$ is an unbiased estimate of μ_i.

Confidence Interval for Treatment Means

$$\begin{aligned}
\mathrm{Var}(\overline{y}_{i\bullet}) &= \mathrm{Var}\left(\frac{1}{n_i}\sum\nolimits_{j=1}^{n_i} y_{ij}\right) = \frac{1}{n_i^2}\mathrm{Var}\left(\sum\nolimits_{j=1}^{n_i} y_{ij}\right) \\
&= \frac{1}{n_i^2}\sum\nolimits_{j=1}^{n_i} \mathrm{Var}(y_{ij}) \quad \text{(by the indep. of } y_{i1}, \ldots, y_{in_i}\text{)} \\
&= \frac{1}{n_i^2}\sum\nolimits_{j=1}^{n_i} \sigma^2 = \frac{n_i\sigma^2}{n_i^2} = \frac{\sigma^2}{n_i}
\end{aligned}$$

So the standard deviation of $\overline{y}_{i\bullet}$ is $\mathrm{SD}(\overline{y}_{i\bullet}) = \sqrt{\mathrm{Var}(\overline{y}_{i\bullet})} = \dfrac{\sigma}{\sqrt{n_i}}$

The unknown σ is estimated by $\sqrt{MSE}$.
Recall the **standard error** is an *estimate* of SD, in which the unknown σ is replaced by some estimate $\widehat{\sigma}$. So

$$\mathrm{SE}(\overline{y}_{i\bullet}) = \widehat{\mathrm{SD}(\overline{y}_{i\bullet})} = \frac{\sqrt{MSE}}{\sqrt{n_i}}.$$

Confidence Interval for Treatment Means

Most of the confidence intervals have the general from:

$$\textit{Confidence Interval} = \textit{estimate} \pm (\textit{multiplier} \times \textit{standard error})$$

The $100(1-\alpha)\%$ Confidence Interval (C.I.) for μ_i is

$$\overline{y}_{i\bullet} \pm t_{\alpha/2,N-g} \frac{\sqrt{MSE}}{\sqrt{n_i}}$$

For resin glue failure time example, we have $N = 37$, $g = 5$. For $\alpha = 0.05$, $1 - \alpha/2 = 0.975$, we know $t_{37-5,0.975} \approx 2.037$.

```
> qt(1-0.05/2, df = 37-5)
[1] 2.036933
```

We have also found $MSE = 0.294$ and hence

Temp	n_i	$\overline{y}_{i\bullet}$	95% C.I. for μ_i
175	8	1.933	$1.933 \pm 2.037 \times \sqrt{0.294/8} = 1.933 \pm 0.390$
194	8	1.629	$1.629 \pm 2.037 \times \sqrt{0.294/8} = 1.629 \pm 0.390$
213	8	1.378	$1.378 \pm 2.037 \times \sqrt{0.294/8} = 1.378 \pm 0.390$
231	7	1.194	$1.194 \pm 2.037 \times \sqrt{0.294/7} = 1.194 \pm 0.417$
250	6	1.057	$1.057 \pm 2.037 \times \sqrt{0.294/6} = 1.057 \pm 0.451$

Pairwise Comparison of Treatments

Consider the general pairwise comparison (difference between two means): $\mu_i - \mu_k$

- the estimator is $\overline{y}_{i\bullet} - \overline{y}_{k\bullet}$
- Since $\overline{y}_{i\bullet}$ and $\overline{y}_{k\bullet}$ are independent,

$$\mathrm{Var}(\overline{y}_{i\bullet} - \overline{y}_{k\bullet}) = \mathrm{Var}(\overline{y}_{i\bullet}) + \mathrm{Var}(\overline{y}_{k\bullet}) = \frac{\sigma^2}{n_i} + \frac{\sigma^2}{n_k}$$

- $\mathrm{SD}(\overline{y}_{i\bullet} - \overline{y}_{k\bullet}) = \sqrt{\mathrm{Var}(\widehat{\mu}_i - \widehat{\mu}_k)} = \sqrt{\sigma^2\left(\frac{1}{n_i} + \frac{1}{n_k}\right)},$
- $\mathrm{SE}(\overline{y}_{i\bullet} - \overline{y}_{k\bullet}) = \widehat{\mathrm{SD}}(\overline{y}_{i\bullet} - \overline{y}_{k\bullet}) = \sqrt{\mathrm{MSE}\left(\frac{1}{n_i} + \frac{1}{n_k}\right)}.$

Confidence Intervals for Differences

The $100(1-\alpha)\%$ confidence interval (C.I.) for $\mu_i - \mu_k$ is

$$\overline{y}_{i\bullet} - \overline{y}_{k\bullet} \pm t_{\alpha/2,N-g}\sqrt{\text{MSE}\left(\frac{1}{n_i}+\frac{1}{n_k}\right)}$$

Note this is NOT equivalent to a two-sample comparison between treatment i and treatment k, which is

$$\overline{y}_{i\bullet} - \overline{y}_{k\bullet} \pm t_{\alpha/2,n_i+n_k-2}\sqrt{s_p^2\left(\frac{1}{n_i}+\frac{1}{n_k}\right)}$$

- $\text{MSE} = \frac{\sum_{i=1}^{g}\sum_{j=1}^{n_i}(y_{ij}-\overline{y}_{i\bullet})^2}{N-g}$ is a more accurate estimator of σ^2 than the pooled sample variance

$$s_p^2 = \frac{\sum_{j=1}^{n_i}(y_{ij}-\overline{y}_{i\bullet})^2 + \sum_{j=1}^{n_2}(y_{kj}-\overline{y}_{k\bullet})^2}{n_i+n_k-2}$$

- The multiplier for the two-sample C.I. is larger

$$t_{\alpha/2,n_i+n_k-2} > t_{\alpha/2,N-g}$$

Hypothesis Testing for Difference

For the hypotheses

$$\begin{cases} H_0: & \mu_i - \mu_k = 0 \\ H_a: & \mu_i - \mu_k \neq 0 \end{cases}$$

the test statistic is

$$t_0 = \frac{\overline{y}_{i\bullet} - \overline{y}_{k\bullet}}{\mathrm{SE}(\overline{y}_{i\bullet} - \overline{y}_{k\bullet})} = \frac{\overline{y}_{i\bullet} - \overline{y}_{k\bullet}}{\sqrt{\mathrm{MSE}}\sqrt{\frac{1}{n_i} + \frac{1}{n_k}}} \sim t_{N-g}$$

which is of $t-$distribution with d.f. $= N - g$. The P-value is

$$\mathrm{P}(|t_{N-g}| > |t_0|) = 2\mathrm{P}(t_{N-g} > |t_0|).$$

For one-sided alternatives, the P-value is

$$\begin{aligned} &\mathrm{P}(t_{N-g} > t_0), \qquad \text{if } H_a: \mu_i - \mu_k > 0, \\ &\mathrm{P}(t_{N-g} < t_0), \qquad \text{if } H_a: \mu_i - \mu_k < 0. \end{aligned}$$

Example — Beet Lice

- Goal: efficacy of 4 chemical treatments for beet lice
- Experimental units: 100 beet plants in individual pots in total, 25 plants per treatment, randomly assigned
- Response: # of lice on each plant at the end of the 2nd week
- The plots are spatially separated (in this random order)

```
> beet = read.table("beetlice.txt", header=TRUE)
> attach(beet)
> plot(ttt,licecount,xlab="Treatment",ylab="Number of Lice")
```

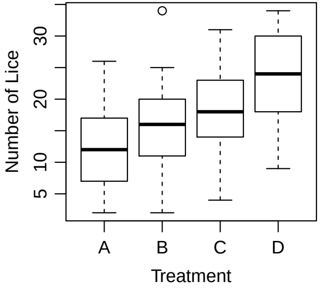

Example — Beet Lice

The group means of the 4 treatments are

```
> tapply(licecount,ttt,mean)
    A     B     C     D
12.00 14.96 18.36 24.00
```

From the ANOVA table below, we get MSE $= 47.8$.

```
> aov1 = aov(licecount ~ ttt, data = beet)
> summary(aov1)
            Df Sum Sq Mean Sq F value   Pr(>F)
ttt          3   1989   663.1   13.86 1.39e-07 ***
Residuals   96   4593    47.8
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Example — Beet Lice

Chemical	A	B	C	D
$\overline{y}_{i\bullet}$	12.00	14.96	18.36	24.00

, MSE = 47.8

Say we want to compare chemical A and B, $\mu_B - \mu_A$.
The estimate is $\overline{y}_{B\bullet} - \overline{y}_{A\bullet} = 14.96 - 12.00 = 2.96$ with SE

$$\text{SE}(\overline{y}_{B\bullet} - \overline{y}_{A\bullet}) = \sqrt{\text{MSE}}\sqrt{\frac{1}{n_A} + \frac{1}{n_B}} = \sqrt{47.8}\sqrt{\frac{1}{25} + \frac{1}{25}} = 1.956$$

A 95% confidence interval for $\mu_B - \mu_A$ is

$$\begin{aligned} &\overline{y}_{B\bullet} - \overline{y}_{A\bullet} \pm t_{96,0.025} \times \text{SE}(\overline{y}_{B\bullet} - \overline{y}_{A\bullet}) \\ =&\, 2.96 \pm 1.984984 \times 1.956 = (-0.923, 6.843) \end{aligned}$$

in which $t_{96,0.025}$ is found using the R command

```
> qt(0.975,df=96)
[1] 1.984984
```

The two chemicals are not significantly different at 5% level because 0 is in the 95% confidence interval.

Example — Beet Lice

If one wants to test whether chemical A and chemical B have the same effect

$$\begin{cases} H_0: & \mu_B - \mu_A = 0 \\ H_a: & \mu_B - \mu_A \neq 0 \end{cases}$$

the test statistic is

$$t_0 = \frac{\overline{y}_{B\bullet} - \overline{y}_{A\bullet}}{\mathrm{SE}(\overline{y}_{B\bullet} - \overline{y}_{A\bullet})} = \frac{2.96}{1.956} = 1.513$$

with $100 - 4 = 96$ degrees of freedom. The P-value is

```
> 2*pt(-1.513, df = 96)
[1] 0.1335649
```

As the P-value > 0.05, we again confirm that the two chemicals are similar in efficacy.

Example — Beet Lice

```
> lmA = lm(licecount ~ as.factor(ttt), data=beet)
> summary(lmA)
...(some output omitted)

Coefficients:
                 Estimate Std. Error t value Pr(>|t|)
(Intercept)        12.000      1.383   8.675 1.04e-13 ***
as.factor(ttt)B     2.960      1.956   1.513  0.13356
as.factor(ttt)C     6.360      1.956   3.251  0.00159 **
as.factor(ttt)D    12.000      1.956   6.134 1.91e-08 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 6.917 on 96 degrees of freedom
Multiple R-squared: 0.3022,     Adjusted R-squared: 0.2804
F-statistic: 13.86 on 3 and 96 DF,  p-value: 1.394e-07
```

Example — Beet Lice

```
> beet$ttt = relevel(as.factor(beet$ttt), ref = 2)
> lmB = lm(licecount ~ as.factor(ttt), data=beet)
> summary(lmB)
...(some output omitted)

Coefficients:
                Estimate Std. Error t value Pr(>|t|)
(Intercept)       14.960      1.383  10.814  < 2e-16 ***
as.factor(ttt)A   -2.960      1.956  -1.513   0.1336
as.factor(ttt)C    3.400      1.956   1.738   0.0854 .
as.factor(ttt)D    9.040      1.956   4.621 1.19e-05 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 6.917 on 96 degrees of freedom
Multiple R-squared: 0.3022,     Adjusted R-squared: 0.2804
F-statistic: 13.86 on 3 and 96 DF,  p-value: 1.394e-07
```

Example — Beet Lice

```
> beet$ttt = relevel(as.factor(beet$ttt), ref = 3)
> lmC = lm(licecount ~ as.factor(ttt), data=beet)
> summary(lmC)
...(some output omitted)

Coefficients:
                Estimate Std. Error t value Pr(>|t|)
(Intercept)       18.360      1.383  13.272  < 2e-16 ***
as.factor(ttt)B   -3.400      1.956  -1.738  0.08543 .
as.factor(ttt)A   -6.360      1.956  -3.251  0.00159 **
as.factor(ttt)D    5.640      1.956   2.883  0.00486 **
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 6.917 on 96 degrees of freedom
Multiple R-squared: 0.3022,     Adjusted R-squared: 0.2804
F-statistic: 13.86 on 3 and 96 DF,  p-value: 1.394e-07
```

Pairwise Comparison Summary

Chemical Comparison	Estimate	Standard Error	t-value	p-value	
$\mu_B - \mu_A$	2.960	1.956	1.513	0.13356	
$\mu_C - \mu_A$	6.360	1.956	3.251	0.00159	**
$\mu_D - \mu_A$	12.000	1.956	6.134	1.91×10^{-8}	***
$\mu_C - \mu_B$	3.400	1.956	1.738	0.0854	.
$\mu_D - \mu_B$	9.040	1.956	4.621	1.19×10^{-5}	***
$\mu_D - \mu_C$	5.640	1.956	2.883	0.00486	**

Display Results — **Underline Diagram** (p.88, Section 5.4.1)

1. Write out treatment labels horizontally in increasing order sorted by treatment means
2. Draw a line segment under a group of treatments if no pair of treatments in that group is significantly different.

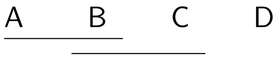

4.1-4.2 Contrasts

Contrasts is a more general form of treatment comparisons.

A **contrast** is a linear combination of group means μ_i's

$$C = \sum_{i=1}^{g} \omega_i \mu_i \quad \text{such that} \quad \sum_{i=1}^{g} \omega_i = 0$$

- **Every pairwise comparison is a contrast!** $(C = \mu_i - \mu_k)$
  $\omega_i = 1$, $\omega_k = -1$, all other ω_i's are 0, and
  $\sum_{i=1}^{g} \omega_i = 1 + (-1) + 0 + \cdots + 0 = 0$
- **A single treatment mean $C = \mu_i$ is NOT a contrast**
- Is $C = \dfrac{\mu_1 + \mu_2}{2} - \dfrac{\mu_3 + \mu_4 + \mu_5}{3}$ a contrast? Yes.
  $\omega_1 = \omega_2 = \frac{1}{2}$, $\omega_3 = \omega_4 = \omega_5 = -\frac{1}{3}$, which adds up to 0.
- Is $C = \dfrac{\mu_1 + \mu_2}{2} - \dfrac{\mu_3 + \mu_4}{2} + \mu_5$ a contrast? No.
  $\omega_1 = \omega_2 = \frac{1}{2}$, $\omega_3 = \omega_4 = -\frac{1}{2}$, $\omega_5 = 1$, which add up to 1, not 0.

Recall

$$\text{the means model: } y_{ij} = \mu_i + \varepsilon_{ij}$$
$$\text{and the effects model: } y_{ij} = \mu + \alpha_i + \varepsilon_{ij}.$$

In the effects model, the "treatment effect" $\alpha_i = \mu_i - \mu$ is arbitrary, because it depends on how the grand mean μ was defined, e.g.,

$$\mu = \mu_1, \quad \text{or} \quad \mu = \frac{1}{N}\sum_{i=1}^{g} n_i \mu_i, \quad \text{or} \quad \mu = \frac{1}{g}\sum_{i=1}^{g} \mu_i.$$

If possible, avoid talking about the size of a treatment effect without clearly stating the definition of μ and α_i's.

Contrasts make comparisons and do not have the problem of effects depending on an arbitrary definition of μ.

$$C = \sum_{i=1}^{g} \omega_i \mu_i = \sum_{i=1}^{g} \omega_i (\mu + \alpha_i) = \mu \underbrace{\sum_{i=1}^{g} \omega_i}_{=0} + \sum_{i=1}^{g} \omega_i \alpha_i = \sum_{i=1}^{g} \omega_i \alpha_i$$

Estimator and Confidence Interval for a Contrast

A natural estimator for a contrast $C = \sum_{i=1}^{g} \omega_i \mu_i$ is

$$\widehat{C} = \sum_{i=1}^{g} \omega_i \overline{y}_{i\bullet}$$

By the independence of $\overline{y}_{i\bullet}$, $i = 1, \cdots, g$, we know

$$\mathrm{Var}\Big(\sum_{i=1}^{g} \omega_i \overline{y}_{i\bullet}\Big) = \sum_{i=1}^{g} \mathrm{Var}(\omega_i \overline{y}_{i\bullet}) = \sum_{i=1}^{g} \omega_i^2 \mathrm{Var}(\overline{y}_{i\bullet}) = \sum_{i=1}^{g} \omega_i^2 \frac{\sigma^2}{n_i}.$$

The estimator $\widehat{C}$ has standard deviation and standard error

$$\mathrm{SD}(\widehat{C}) = \sqrt{\sigma^2 \sum_{i=1}^{g} \frac{\omega_i^2}{n_i}}, \qquad \mathrm{SE}(\widehat{C}) = \sqrt{\mathrm{MSE} \times \sum_{i=1}^{g} \frac{\omega_i^2}{n_i}}$$

A $(1-\alpha)100\%$ confidence interval for the contrast C is

$$\widehat{C} \pm t_{\alpha/2, N-g} \times \mathrm{SE}(\widehat{C})$$

Hypothesis Testing for a Contrast

For the hypotheses

$$\begin{cases} H_0: & C = 0 \\ H_a: & C \neq 0 \end{cases}$$

the test statistic is

$$t_0 = \frac{\widehat{C}}{\mathrm{SE}(\widehat{C})} = \frac{\sum_{i=1}^{g} \omega_i \bar{y}_{i\bullet}}{\sqrt{\mathrm{MSE} \times \sum_{i=1}^{g} \frac{\omega_i^2}{n_i}}} \sim t_{N-g}$$

which is of $t-$distribution with d.f. $= N - g$. For a two-sided test, the P-value is

$$2\mathrm{P}(t_{N-g} > |t_0|)$$

For one-sided alternative, the one-sided P-value is

$$\begin{cases} \mathrm{P}(t_{N-g} > t_0), & \text{if } H_a : C > 0, \\ \mathrm{P}(t_{N-g} < t_0), & \text{if } H_a : C < 0. \end{cases}$$

Example — Beet Lice

Recall

Chemical	A	B	C	D
$\overline{y}_{i\bullet}$	12.00	14.96	18.36	24.00

, MSE = 47.8

We want a comparison between treatment A, B, C (all liquid) and treatment D (powder). We thus consider the contrast

$$C = \frac{\mu_1 + \mu_2 + \mu_3}{3} - \mu_4$$

in which $(\omega_1, \omega_2, \omega_3, \omega_4) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, -1)$. The contrast is estimated by

$$\begin{aligned} \widehat{C} &= \frac{\overline{y}_{A\bullet} + \overline{y}_{B\bullet} + \overline{y}_{C\bullet}}{3} - \overline{y}_{D\bullet} \\ &= \frac{12.00 + 14.96 + 18.36}{3} - 24.00 = -8.893 \end{aligned}$$

with standard error

$$\begin{aligned} \text{SE}(\widehat{C}) = \sqrt{\text{MSE} \times \sum_{i=1}^{g} \frac{\omega_i^2}{n_i}} &= \sqrt{47.8(\frac{(1/3)^2}{25} + \frac{(1/3)^2}{25} + \frac{(1/3)^2}{25} + \frac{(-1)^2}{25})} \\ &= \sqrt{47.8 \times \frac{4}{75}} = 1.597. \end{aligned}$$

Example — Beet Lice

The 95% confidence interval for C is

$$\begin{aligned}&\widehat{C} \pm t_{0.025,96} \times \text{SE}(\widehat{C})\\ = &-8.893 \pm 1.984984 \times 1.597 = (-12.063, -5.723)\end{aligned}$$

in which $t_{0.025,96}$ is found in R via the command

```
> qt(0.975,df=96)
[1] 1.984984
```

The t-statistic is

$$t_0 = \frac{\widehat{C}}{\text{SE}(\widehat{C})} = \frac{-8.893}{1.597} = -5.568,$$

with $100 - 4 = 96$ degrees of freedom. The two-sided p-value is

```
> 2*pt(-5.568,df=96)
[1] 2.339348e-07
```

Because 0 is not in the confidence interval, or because of the small P-value, we can conclude that the 3 liquid chemicals are *less effective* than chemical D.

4.3 Orthogonal Contrasts

Two contrasts $C = \sum_{i=1}^{g} \omega_i \mu_i$ and $D = \sum_{i=1}^{g} \omega_i^* \mu_i$ are said to be **orthogonal** if

$$\sum_{i=1}^{g} \frac{\omega_i \omega_i^*}{n_i} = 0.$$

Note if the design is balanced $n_1 = n_2 = \cdots = n_g$, the condition is equivalent to $\sum_{i=1}^{g} \omega_i \omega_i^* = 0$.

When two contrasts are orthogonal, their estimators are independent.