

The Statistical Sleuth in R:

Chapter 6

Linda Loi Ruobing Zhang Kate Aloisio Nicholas J. Horton*

January 21, 2013

Contents

1	Introduction	1
2	Discrimination Against the Handicapped	2
2.1	Summary statistics and graphical display	2
2.2	One-way ANOVA	5
2.3	Contrasts and linear combination	6
3	Pre-existing Preference of Fish	11
3.1	Summary statistics and graphical display	11
3.2	One-way ANOVA	12
3.3	Contrasts and linear combination	14

1 Introduction

This document is intended to help describe how to undertake analyses introduced as examples in the Third Edition of the *Statistical Sleuth* (2013) by Fred Ramsey and Dan Schafer. More information about the book can be found at <http://www.proaxis.com/~panorama/home.htm>. This file as well as the associated **knitr** reproducible analysis source file can be found at <http://www.math.smith.edu/~nhorton/sleuth3>.

This work leverages initiatives undertaken by Project MOSAIC (<http://www.mosaic-web.org>), an NSF-funded effort to improve the teaching of statistics, calculus, science and computing in the undergraduate curriculum. In particular, we utilize the mosaic package, which was written to simplify the use of R for introductory statistics courses. A short summary of the R needed to teach introductory statistics can be found in the mosaic package vignette (<http://cran.r-project.org/web/packages/mosaic/vignettes/MinimalR.pdf>).

To use a package within R, it must be installed (one time), and loaded (each session). The package can be installed using the following command:

*Department of Mathematics and Statistics, Smith College, nhorton@smith.edu

```
> install.packages("mosaic") # note the quotation marks
```

Once this is installed, it can be loaded by running the command:

```
> require(mosaic)
```

This needs to be done once per session.

In addition the data files for the *Sleuth* case studies can be accessed by installing the **Sleuth3** package.

```
> install.packages("Sleuth3") # note the quotation marks
```

```
> require(Sleuth3)
```

We also set some options to improve legibility of graphs and output.

```
> trellis.par.set(theme = col.mosaic()) # get a better color scheme for lattice
> options(digits = 3)
```

The specific goal of this document is to demonstrate how to calculate the quantities described in Chapter 6: Linear Combinations and Multiple Comparisons of Means using R.

2 Discrimination Against the Handicapped

Do equivalent candidates with the same qualifications but different disabilities get treated differently? This is the question addressed in case study 6.1 in the *Sleuth*.

2.1 Summary statistics and graphical display

We begin by reading the data and summarizing the variables.

```
> case0601$Handicap = relevel(case0601$Handicap, ref = "Amputee")
> summary(case0601)
```

Score		Handicap	
Min.	:1.40	Amputee	:14
1st Qu.	:3.70	Crutches	:14
Median	:5.05	Hearing	:14
Mean	:4.93	None	:14
3rd Qu.	:6.10	Wheelchair	:14
Max.	:8.50		

```
> favstats(Score ~ Handicap, data = case0601)
```

	min	Q1	median	Q3	max	mean	sd	n	missing
Amputee	1.9	3.30	4.30	5.72	7.2	4.43	1.59	14	0
Crutches	3.7	4.50	6.10	7.15	8.5	5.92	1.48	14	0
Hearing	1.4	3.02	4.05	5.30	6.5	4.05	1.53	14	0
None	1.9	3.73	5.00	6.05	7.8	4.90	1.79	14	0
Wheelchair	1.7	4.73	5.70	6.35	7.6	5.34	1.75	14	0

A total of 70 undergraduate students from a U.S. university were randomly assigned to view the tapes, 14 to each tape. The five kinds of tapes are: *None*, *Amputee*, *Crutches*, *Hearing* and *Wheelchair*. After reviewing the tape, each subject rated the qualifications of the application on 0-10 scale. Among the five handicap conditions, the *Crutches* group gave the highest mean score, while the *Hearing* group gave the lowest mean score. This is summarized on page 150 and in Display 6.1 of the *Sleuth*.

```
> with(subset(case0601, Handicap == "None"), stem(Score, scale = 2))
```

```
The decimal point is at the |
```

```
1 | 9
2 | 5
3 | 06
4 | 129
5 | 149
6 | 17
7 | 48
```

```
> with(subset(case0601, Handicap == "Amputee"), stem(Score, scale = 2))
```

```
The decimal point is at the |
```

```
1 | 9
2 | 56
3 | 268
4 | 06
5 | 3589
6 | 1
7 | 2
```

```
> with(subset(case0601, Handicap == "Crutches"), stem(Score, scale = 1))
```

```
The decimal point is at the |
```

```
3 | 7
```

```
4 | 033
5 | 18
6 | 0234
7 | 445
8 | 5

> with(subset(case0601, Handicap == "Hearing"), stem(Score, scale = 2))

The decimal point is at the |

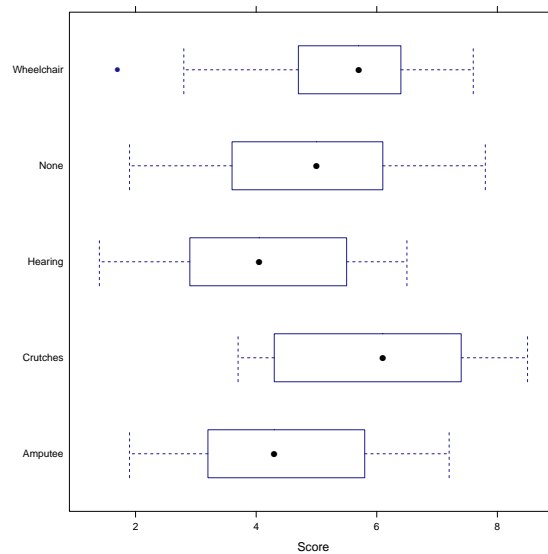
1 | 4
2 | 149
3 | 479
4 | 237
5 | 589
6 | 5

> with(subset(case0601, Handicap == "Wheelchair"), stem(Score, scale = 2))

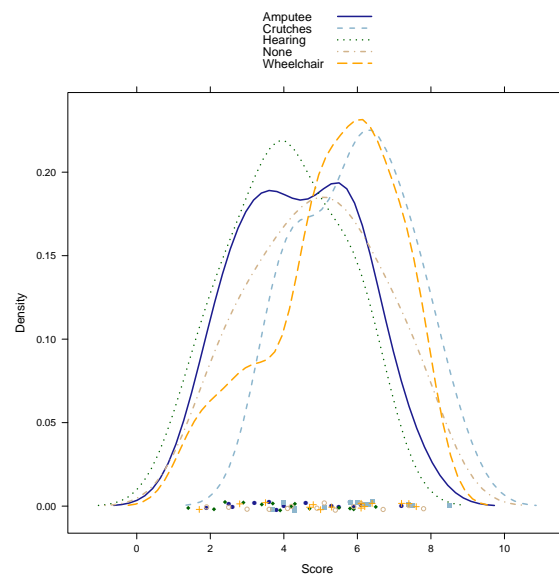
The decimal point is at the |

1 | 7
2 | 8
3 | 5
4 | 78
5 | 03
6 | 1124
7 | 246

> bwplot(Handicap ~ Score, data = case0601)
```



```
> densityplot(~Score, groups = Handicap, auto.key = TRUE, data = case0601)
```



The stem plots show the applicant qualification scores given by objectives. The boxplots and the density plots show that all the distributions are approximately normally distributed.

2.2 One-way ANOVA

First we fit the one way analysis of variance (ANOVA) model, using all of the groups. This corresponds to the interpretations on page 151.

```
> anova(lm(Score ~ Handicap, data = case0601))
```

Analysis of Variance Table

Response: Score

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Handicap	4	30.5	7.63	2.86	0.03 *
Residuals	65	173.3	2.67		

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

The p-value provides some evidence that subjects rate qualifications differently according to handicap status.

By default, the use of the linear model (regression) function displays the pairwise differences between the first group and each of the other groups. Note that the overall test of the model is the same.

```
> summary(lm(Score ~ Handicap, data = case0601))
```

Call:

lm(formula = Score ~ Handicap, data = case0601)

Residuals:

Min	1Q	Median	3Q	Max
-3.643	-1.209	0.114	1.329	2.900

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	4.429	0.436	10.15	5e-15 ***
HandicapCrutches	1.493	0.617	2.42	0.018 *
HandicapHearing	-0.379	0.617	-0.61	0.542
HandicapNone	0.471	0.617	0.76	0.448
HandicapWheelchair	0.914	0.617	1.48	0.143

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

Residual standard error: 1.63 on 65 degrees of freedom

Multiple R-squared: 0.15, Adjusted R-squared: 0.0974

F-statistic: 2.86 on 4 and 65 DF, p-value: 0.0301

The reference group here is *Amputee*, followed by *Crutches*, *Hearing*, *None* and *Wheelchair*.

Another way of viewing these results is through a model table, which displays the differences between the grand mean and the group means.

```
> model.tables(aov(Score ~ Handicap, data = case0601))
```

Tables of effects

```

Handicap
Handicap
  Amputee  Crutches  Hearing    None Wheelchair
-0.5000    0.9929  -0.8786  -0.0286    0.4143

```

Or by:

```

> mean(Score ~ Handicap, data = case0601) - mean(~Score, data = case0601)

  Amputee  Crutches  Hearing    None Wheelchair
-0.5000    0.9929  -0.8786  -0.0286    0.4143

```

2.3 Contrasts and linear combination

The Tukey-Kramer test is a reasonable method for these data. We can use this to verify the calculation on page 151.

```

> TukeyHSD(aov(lm(Score ~ Handicap, data = case0601)), "Handicap", ordered = TRUE,
+          c(0, 1, -1, 0, 0), conf.level = 0.95)

```

```

Tukey multiple comparisons of means
 95% family-wise confidence level
factor levels have been ordered

```

```

Fit: aov(formula = lm(Score ~ Handicap, data = case0601))

```

```

$Handicap
      diff      lwr      upr p adj
Amputee-Hearing  0.379 -1.353  2.11 0.972
None-Hearing     0.850 -0.882  2.58 0.644
Wheelchair-Hearing 1.293 -0.439  3.02 0.235
Crutches-Hearing  1.871  0.140  3.60 0.028
None-Amputee     0.471 -1.260  2.20 0.940
Wheelchair-Amputee 0.914 -0.817  2.65 0.578
Crutches-Amputee  1.493 -0.239  3.22 0.123
Wheelchair-None   0.443 -1.289  2.17 0.952
Crutches-None     1.021 -0.710  2.75 0.469
Crutches-Wheelchair 0.579 -1.153  2.31 0.881

```

Based on the Tukey-Kramer procedure, the difference is estimated to be higher for the *Crutches* tapes.

Next, we calculate the comparison of *Amputee/Hearing* to *Crutches/Wheelchair*.

```
> require(gmodels)

Loading required package: gmodels

> fit.contrast(lm(Score ~ Handicap, data = case0601), "Handicap", c(-1, 1, -1,
+ 0, 1), conf.int = 0.95)
```

	Estimate	Std. Error	t value	Pr(> t)	lower CI
Handicap c=(-1 1 -1 0 1)	2.79	0.873	3.19	0.00218	1.04
					upper CI
Handicap c=(-1 1 -1 0 1)	4.53				

The results indicate a statistically significant difference between the average scores given to the *Wheelchair* and *Crutches* handicaps and the average scores given to the *Amputee* and *Hearing* handicaps.

To verify the calculations on page 155 we used the following contrast:

```
> fit.contrast(lm(Score ~ Handicap, data = case0601), "Handicap", c(-0.5, 0.5,
+ -0.5, 0, 0.5), conf.int = 0.95)
```

	Estimate	Std. Error	t value	Pr(> t)
Handicap c=(-0.5 0.5 -0.5 0 0.5)	1.39	0.436	3.19	0.00218
				lower CI upper CI
Handicap c=(-0.5 0.5 -0.5 0 0.5)	0.521	2.26		

Other multiple comparison procedures could also be implemented. The following shows the calculation on page 164.

```
> require(agricolae)

Loading required package: agricolae

> LSD.test(aov(lm(Score ~ Handicap, data = case0601)), "Handicap") # LSD
```

Study:

LSD t Test for Score

Mean Square Error: 2.67

Handicap, means and individual (95 %) CI

	Score	std.err	r	LCL	UCL	Min.	Max.
Amputee	4.43	0.424	14	3.58	5.27	1.9	7.2
Crutches	5.92	0.396	14	5.13	6.71	3.7	8.5

```

Hearing      4.05    0.410 14 3.23 4.87  1.4  6.5
None         4.90    0.479 14 3.94 5.86  1.9  7.8
Wheelchair   5.34    0.467 14 4.41 6.28  1.7  7.6

alpha: 0.05 ; Df Error: 65
Critical Value of t: 2

Least Significant Difference 1.23
Means with the same letter are not significantly different.

Groups, Treatments and means
a   Crutches      5.92
ab  Wheelchair    5.34
abc None          4.9
bc  Amputee       4.43
c   Hearing        4.05

> HSD.test(aov(lm(Score ~ Handicap, data = case0601)), "Handicap") # Tukey-Kramer

Study:

HSD Test for Score

Mean Square Error:  2.67

Handicap, means

      Score std.err  r Min. Max.
Amputee   4.43   0.424 14  1.9  7.2
Crutches  5.92   0.396 14  3.7  8.5
Hearing   4.05   0.410 14  1.4  6.5
None      4.90   0.479 14  1.9  7.8
Wheelchair 5.34   0.467 14  1.7  7.6

alpha: 0.05 ; Df Error: 65
Critical Value of Studentized Range: 3.97

Honestly Significant Difference: 1.73

Means with the same letter are not significantly different.

Groups, Treatments and means
a   Crutches      5.92
ab  Wheelchair    5.34
ab  None          4.9

```

```

ab    Amputee      4.43
b     Hearing      4.05

> LSD.test(aov(lm(Score ~ Handicap, data = case0601)), "Handicap", p.adj = c("bonferroni")) #

Study:

LSD t Test for Score
P value adjustment method: bonferroni

Mean Square Error:  2.67

Handicap, means and individual ( 95 %) CI

      Score std.err  r  LCL  UCL Min. Max.
Amputee   4.43   0.424 14 3.58 5.27  1.9  7.2
Crutches  5.92   0.396 14 5.13 6.71  3.7  8.5
Hearing   4.05   0.410 14 3.23 4.87  1.4  6.5
None      4.90   0.479 14 3.94 5.86  1.9  7.8
Wheelchair 5.34   0.467 14 4.41 6.28  1.7  7.6

alpha: 0.05 ; Df Error: 65
Critical Value of t: 2.91

Least Significant Difference 1.79
Means with the same letter are not significantly different.

Groups, Treatments and means
a    Crutches      5.92
ab   Wheelchair    5.34
ab   None          4.9
ab   Amputee       4.43
b    Hearing        4.05

> scheffe.test(aov(lm(Score ~ Handicap, data = case0601)), "Handicap") # Scheffe

Study:

Scheffe Test for Score

Mean Square Error : 2.67

Handicap, means

```

```

      Score std.err  r Min. Max.
Amputee    4.43   0.424 14  1.9  7.2
Crutches    5.92   0.396 14  3.7  8.5
Hearing     4.05   0.410 14  1.4  6.5
None        4.90   0.479 14  1.9  7.8
Wheelchair  5.34   0.467 14  1.7  7.6

alpha: 0.05 ; Df Error: 65
Critical Value of F: 2.51

Minimum Significant Difference: 1.96

Means with the same letter are not significantly different.

Groups, Treatments and means
a  Crutches      5.92
a  Wheelchair    5.34
a  None          4.9
a  Amputee       4.43
a  Hearing        4.05

```

The “Significant Difference” in each test result is the “95% interval half-width” described in the book.

3 Pre-existing Preference of Fish

Was Charles Darwin right that sexual selection is driven by females? This is the question addressed in case study 6.2 in the *Sleuth*.

3.1 Summary statistics and graphical display

We begin by reading the data and summarizing the variables.

```

> case0602$NewProportion = case0602$Proportion * 100
> case0602$Proportion = case0602$NewProportion
> summary(case0602)

      Proportion      Pair      Length      NewProportion
Min.   :10.0   Pair1:16   Min.    :28.0   Min.    :10.0
1st Qu.:53.0   Pair2:14   1st Qu.:31.0   1st Qu.:53.0
Median :61.5   Pair3:17   Median :34.0   Median :61.5
Mean   :62.1   Pair4:14   Mean   :32.8   Mean   :62.1
3rd Qu.:72.0   Pair5: 9   3rd Qu.:34.0   3rd Qu.:72.0
Max.   :92.0   Pair6:14   Max.    :35.0   Max.    :92.0

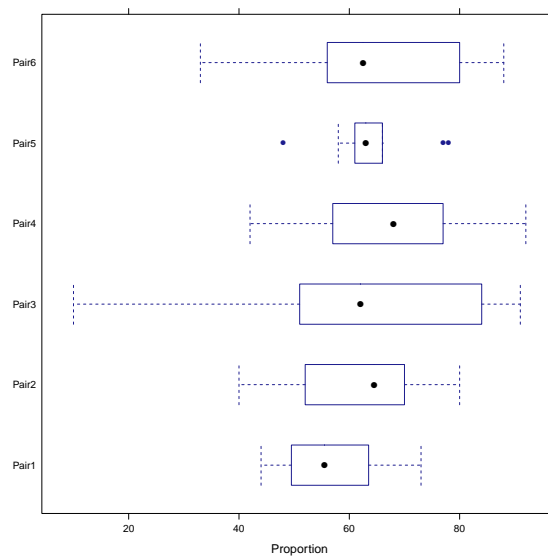
> favstats(Proportion ~ Pair, data = case0602)

```

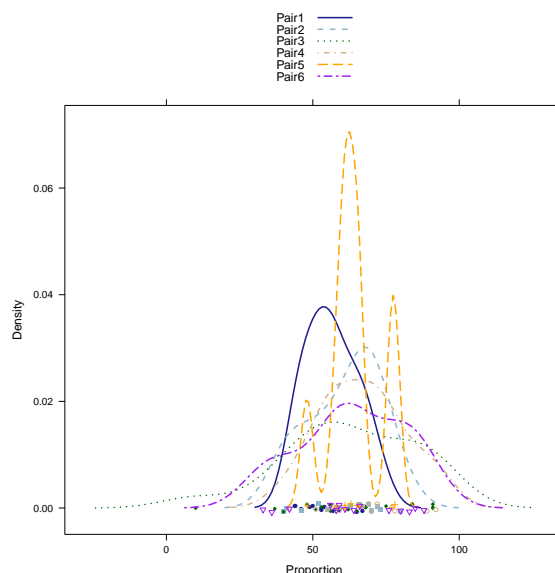
	min	Q1	median	Q3	max	mean	sd	n	missing
Pair1	44	49.8	55.5	62.8	73	56.4	8.86	16	0
Pair2	40	52.8	64.5	69.8	80	60.9	12.47	14	0
Pair3	10	51.0	62.0	84.0	91	62.5	22.34	17	0
Pair4	42	57.2	68.0	75.8	92	66.9	14.24	14	0
Pair5	48	61.0	63.0	66.0	78	64.1	9.23	9	0
Pair6	33	56.5	62.5	79.0	88	63.4	17.84	14	0

A total of 84 female fish were involved in this experiment, which is shown on page 153.

```
> bwplot(Pair ~ Proportion, data = case0602)
```



```
> densityplot(~Proportion, groups = Pair, auto.key = TRUE, data = case0602)
```



Besides the distribution of pair 5, all distributions of other pairs are approximately normally distributed.

3.2 One-way ANOVA

First we fit the one way analysis of variance (ANOVA) model, using all of the groups:

```
> anova(lm(Proportion ~ Pair, data = case0602))
```

Analysis of Variance Table

Response: Proportion

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Pair	5	921	184	0.77	0.57
Residuals	78	18643	239		

The p-value is not small, and does not provide much evidence that the mean percentage of time with the yellow-sword male significantly differed from one male pair to another back in the population.

By default, the use of the linear model (regression) function displays the pairwise differences between the first group and each of the other groups. Note that the overall test of the model is the same.

```
> summary(lm(Proportion ~ Pair, data = case0602))
```

Call:

```
lm(formula = Proportion ~ Pair, data = case0602)
```

Residuals:

```

      Min      1Q  Median      3Q      Max
-52.47  -8.88   0.38  10.74  28.53

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)    56.38      3.86    14.59  <2e-16 ***
PairPair2       4.55      5.66     0.80   0.423
PairPair3       6.10      5.38     1.13   0.261
PairPair4      10.48      5.66     1.85   0.068 .
PairPair5       7.74      6.44     1.20   0.233
PairPair6       6.98      5.66     1.23   0.221
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 15.5 on 78 degrees of freedom
Multiple R-squared:  0.0471, Adjusted R-squared: -0.014
F-statistic: 0.771 on 5 and 78 DF,  p-value: 0.574

```

The reference group here is pair 1, followed by pairs 2-6. Another way of viewing these results is through a model table, which displays the differences between the grand mean and the group means.

```
> model.tables(aov(Proportion ~ Pair, data = case0602))
```

Tables of effects

```

Pair
  Pair1 Pair2 Pair3 Pair4 Pair5 Pair6
-5.732 -1.179 0.3634 4.75 2.004 1.25
rep 16.000 14.000 17.0000 14.00 9.000 14.00

```

Or by:

```

> mean(Proportion ~ Pair, data = case0602) - mean(~Proportion, data = case0602)

Pair1 Pair2 Pair3 Pair4 Pair5 Pair6
-5.732 -1.179 0.363 4.750 2.004 1.250

```

3.3 Contrasts and linear combination

We can calculate the values on page 152 and Display 6.5 on page 158 using contrasts.

```

> require(gmodels)
> lc = fit.contrast(lm(Proportion ~ Pair, data = case0602), "Pair", c(5, -3, 1,
+   3, -9, 3), conf.int = 0.95)
> lc

```

```

              Estimate Std. Error t value Pr(>|t|) lower CI
Pair c=( 5 -3 1 3 -9 3 )    -24.8      54.8  -0.453    0.652    -134
              upper CI
Pair c=( 5 -3 1 3 -9 3 )     84.2

> t = round(lc[, "t value"], 2)
> t

[1] -0.45

> pt(t, 78, lower.tail = TRUE)

[1] 0.327

```

The t -value is -0.45 and the one-sided p -value is 0.33.

```

> mean(mean(Proportion ~ Pair, data = case0602))

[1] 62.3

> t.test(mean(Proportion ~ Pair, data = case0602))

```

One Sample t-test

```

data: mean(Proportion ~ Pair, data = case0602)
t = 43.3, df = 5, p-value = 1.238e-07
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 58.6 66.1
sample estimates:
mean of x
 62.3

```

The estimated mean percentage of time spent with the yellow-sword male is 62.35%. The one-sided p -value < 0.0001 , and the 95% confidence interval is (58.65%, 66.05%).