

Lab Tutorial 9: Repeated-Measures Factorial ANOVA & Split-Plot ANOVA

Chris M. Fiacconi

Overview

To this point we have discussed analysis of designs that include two (or more) *crossed* independent variables where all variables were manipulated between-subjects. As you might imagine, it is also possible to manipulate two (or more) independent variables in a fully repeated-measures design, where every participant is tested in all possible cells of the design. In fact, such designs are extremely commonplace in human (and animal) cognition research. In this lab tutorial, we will discuss analysis techniques for *repeated-measures factorial designs*, as well as for designs in which at least one of the independent variables is manipulated between-subjects. These latter designs consist of a mix of between- and within-subject factors and are known as *split-plot* or *mixed* designs. As you'll see, we can apply many of the principles already discussed in the context of between-subjects factorial designs and one-way repeated measures designs in our analysis of the present experimental designs.

Repeated-Measures Factorial ANOVA

Model Comparisons

Recall that in our discussion of between-subject factorial ANOVA, we compared a full model with a series of restricted models to evaluate the main effects of factors A, B, and their interaction. The full model had the following form:

$$Y_{ijk} = \mu + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \epsilon_{ijk}$$

where μ again represents the grand mean, α_j represents the effect of being in the j th level of Factor A, β_k represents the effect of being in the k th level of Factor B, $(\alpha\beta)_{jk}$ represents the *interaction* effect of being in level j of Factor A and level k of Factor B in combination with one another. This model captures the fact that we can examine the main effects of factor A (α_j 's) and factor B (β_k 's), along with their interaction ($\alpha\beta_{jk}$'s).

Similarly, for a repeated-measures design with one independent variable we compared a full model with a single restricted model to evaluate the main effect of factor A. The full model had the following form:

$$Y_{ij} = \mu + \alpha_j + \pi_i + (\alpha\pi)_{ij} + \epsilon_{ij}$$

where μ is the grand mean, α_j is the effect of being in condition j , π_i is the effect associated with the i th subject (or participant), $(\alpha\pi)_{ij}$ is the interaction between the i th subject in condition j , and ϵ_{ij} is the prediction error associated with subject i in condition j .

For a two-way repeated-measures factorial design, we essentially combine the two models shown above to create the following full model:

$$Y_{ijk} = \mu + \alpha_j + \beta_k + \pi_i + (\alpha\beta)_{jk} + (\alpha\pi)_{ji} + (\beta\pi)_{ki} + (\alpha\beta\pi)_{jki} + \epsilon_{ijk}$$

where where μ is the grand mean, α_j is the effect of being in condition j , π_i is the effect associated with the i th subject (or participant), $(\alpha\pi)_{ij}$ is the interaction between the i th subject in condition j , $(\alpha\beta)_{jk}$ represents the interaction effect of being in level j of A and level k of B, $(\alpha\pi)_{ji}$ represents the interaction effect for the i th subject in condition j , $(\beta\pi)_{ki}$ represents the interaction effect for the i th subject in condition k , $(\alpha\beta\pi)_{jki}$

represents the *three-way* interaction effect for the i th subject in the j th level of A, and k th level of B, and ϵ_{ij} is the prediction error associated with subject i in condition jk . Note that α , β , and $\alpha\beta$ are fixed effects whereas π and all the interactions with π are random effects.

As was the case in a two-way between-subjects factorial design, we can still evaluate the main effects of A, B and the A x B interaction, but the models being compared in the evaluation of each effect differ somewhat due to the repeated-measures nature of the design. Once again, we can evaluate the fixed effects in the model by comparing the full model to a restricted model that omits the parameters associated with the effect being tested. As per usual, these models can be compared using the (hopefully familiar by now) formula:

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

When testing the the main effect of A the F -statistic has the following form:

$$F_A = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_A}{MS_{A \times S}}$$

and the df are calculated as follows:

$$\begin{aligned}(df_R - df_F) &= a - 1 \\ df_F &= (a - 1)(n - 1)\end{aligned}$$

Similarly, when testing the the main effect of B the F -statistic has the following form:

$$F_B = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_B}{MS_{B \times S}}$$

and the df are calculated as follows:

$$\begin{aligned}(df_R - df_F) &= b - 1 \\ df_F &= (b - 1)(n - 1)\end{aligned}$$

For the A x B interaction, the the F -statistic has the following form:

$$F_{A \times B} = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_{A \times B}}{MS_{A \times B \times S}}$$

and the df are:

$$\begin{aligned}(df_R - df_F) &= (a - 1)(b - 1) \\ df_F &= (a - 1)(b - 1)(n - 1)\end{aligned}$$

Note that similar to a one-way repeated-measures ANOVA, the error term used to evaluate each effect (e.g., A main effect) consists of the corresponding effect \times subjects interaction (i.e., $A \times S$). Recall that the need for this error term stems from the fact that MS_A is influenced by the variance associated with the $A \times S$ interaction because S is a random effect (see Lab Tutorial 8). By using $MS_{A \times S}$ as the error term in our F -statistic, we effectively cancel out $A \times S$ interaction variance in both the numerator and the denominator, thus isolating the variance that is solely attributable to A . The same basic logic holds when evaluating the main effect of B and the $A \times B$ interaction.

Example #1

A perceptual psychologist is interested in how the visual system detects target stimuli amongst visual noise. Participants are told to identify a letter as either a 'T' or an 'I' while reaction time (RT) is recorded on each trial. On half of the trials, the target letter is presented clearly on its own, and on the other half of trials it is presented embedded in white visual noise. Moreover, on a third of the trials in both conditions, the target stimulus can appear at 0, 4, or 8 degrees of visual angle. Thus, we have a design with two crossed factors, the first of which is Noise (present, absent) and the second of which is Angle (0-, 4-, 8-degrees) yielding a 2 x 3 factorial design. Both factors are manipulated within-subjects meaning that we have a repeated-measures factorial design. Let's load in the data file adapted from MDK Chapter 12 Table 1.

```
options(contrasts=c("contr.sum","contr.poly")) # Set sum-to-zero contrasts

angle.data<-read.table(file="Ch12T1.txt",header=T)
head(angle.data)
```

```
##   Absent0 Absent4 Absent8 Present0 Present4 Present8
## 1     420     420     480     480     600     780
## 2     420     480     480     360     480     600
## 3     480     480     540     660     780     780
## 4     420     540     540     480     780     900
## 5     540     660     540     480     660     720
## 6     360     420     360     360     480     540
```

Notice that the data are in wide format which is less than ideal for performing our analyses. Let's first convert the data to long format using the **gather** and **separate** functions from the *dplyr* package within *tidyverse*.

```
library(tidyverse)

angle.data$Subject<-as.factor(1:nrow(angle.data)) # add Subject ID column to angle.data
angle.long<-gather(data=angle.data,key=Noise.Angle,value=RT,-Subject) # Gather columns except Subject
head(angle.long)
```

```
##   Subject Noise.Angle  RT
## 1       1     Absent0 420
## 2       2     Absent0 420
## 3       3     Absent0 480
## 4       4     Absent0 420
## 5       5     Absent0 540
## 6       6     Absent0 360
```

```
# Break up Noise.Angle column into 2 separate columns
angle.long<-separate(angle.long,col=Noise.Angle,into=c("Noise","Angle"),sep=-1)

angle.long$Noise<-as.factor(angle.long$Noise) # Ensure Noise is treated as a factor!!
angle.long$Angle<-as.factor(angle.long$Angle) # Ensure Angle is treated as a factor!!

angle.long<-arrange(angle.long,Subject,Noise,Angle) # Sort data by Subject, Noise, Angle

head(angle.long,n=12)
```

```
##   Subject  Noise Angle  RT
## 1      1 Absent   0 420
## 2      1 Absent   4 420
## 3      1 Absent   8 480
## 4      1 Present  0 480
## 5      1 Present  4 600
## 6      1 Present  8 780
## 7      2 Absent   0 420
## 8      2 Absent   4 480
## 9      2 Absent   8 480
## 10     2 Present  0 360
## 11     2 Present  4 480
## 12     2 Present  8 600
```

Now that the data are in long format, we're ready to obtain descriptive statistics, plot the data, as well as perform a 2 x 3 repeated-measures factorial ANOVA:

```
# Get condition means
angle.means<-angle.long %>% group_by(Noise,Angle) %>% summarize(Mean=mean(RT))
print(angle.means)
```

```
## # A tibble: 6 x 3
## # Groups:   Noise [?]
##   Noise  Angle Mean
##   <fct> <fct> <dbl>
## 1 Absent  0     462
## 2 Absent  4     510
## 3 Absent  8     528
## 4 Present 0     492
## 5 Present 4     660
## 6 Present 8     762
```

```
# Plot data
with(angle.long,interaction.plot(x.factor=Angle,trace.factor=Noise,response=RT,las=1,
                                cex=1.5,bty="n",type="b",pch=c(17,19),
                                xlab="Degrees of Visual Angle",ylab="RT (ms)"))

axis(side=1,at=c(1,2,3),labels=c("0","4","8")) # add in x-axis with ticks and labels
```

```
# Do RM ANOVA
angle.aov<-aov(RT~Noise*Angle+Error(Subject/(Noise*Angle)),data=angle.long)
print(summary(angle.aov))
```

```
##
## Error: Subject
##           Df Sum Sq Mean Sq F value Pr(>F)
## Residuals  9 292140   32460
##
## Error: Subject:Noise
##           Df Sum Sq Mean Sq F value  Pr(>F)
## Noise      1 285660  285660  33.77 0.000256 ***
## Residuals  9  76140    8460
```




Figure 1: Plot depicting data taken from MDK Chapter 12 Table 1

```
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Error: Subject:Angle
##           Df Sum Sq Mean Sq F value    Pr(>F)
## Angle      2 289920  144960  40.72 2.09e-07 ***
## Residuals 18  64080    3560
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Error: Subject:Noise:Angle
##           Df Sum Sq Mean Sq F value    Pr(>F)
## Noise:Angle 2 105120   52560  45.31 9.42e-08 ***
## Residuals  18  20880    1160
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Although it appears as though the main effects of Noise, $F(1,9) = 33.77, p < .001$, Angle, $F(1,18) = 40.72, p < .001$, and the Noise x Angle interaction, $F(2,18) = 45.31, p < .001$, are all significant, before drawing any conclusions we need to assess whether the sphericity assumption is upheld, particularly for the effect of Angle (3+ levels) and the Noise x Angle interaction.

As we saw for one-way repeated-measures ANOVA, the `aov` function does not provide this information for us. Thus, once again we'll turn to the `ezANOVA` function in the `ez` package:

```
library(ez)
```

```
# Notice the .() when specifying the within-subject factors
ezANOVA(data=angle.long,dv=RT,wid=Subject,within=(Noise,Angle),detailed=TRUE)
```

```
## $ANOVA
##      Effect DFn DFd      SSn      SSd      F      p p<.05
## 1 (Intercept)  1   9 19425660 292140 598.44917 1.526600e-09 *
## 2      Noise   1   9  285660  76140 33.76596 2.559737e-04 *
## 3      Angle   2  18  289920  64080 40.71910 2.086763e-07 *
## 4 Noise:Angle  2  18  105120  20880 45.31034 9.424093e-08 *
##      ges
## 1 0.9771999
## 2 0.3866017
## 3 0.3901179
## 4 0.1882656
##
## $`Mauchly's Test for Sphericity`
##      Effect      W      p p<.05
## 3      Angle 0.9601060 0.8497219
## 4 Noise:Angle 0.8937772 0.6381418
##
## $`Sphericity Corrections`
##      Effect      GGe      p[GG] p[GG]<.05      HFe      p[HF]
## 3      Angle 0.9616365 3.401705e-07 * 1.217564 2.086763e-07
## 4 Noise:Angle 0.9039771 3.453931e-07 * 1.117870 9.424093e-08
##      p[HF]<.05
## 3      *
## 4      *
```

In this case, the sphericity assumption does not appear to be violated because Mauchley's test is non-significant. Therefore, we can assume sphericity and go ahead using the raw F - and p -values.

Further Investigation of Main Effects

Since there are only 2 levels of the Noise variable, the significant main effect tells us that RTs are longer when noise is present vs. absent. However, given that Angle has 3 levels, a significant main effect does not inform us as to which Angle condition means differ, necessitating follow-up tests.

One approach here would be to conduct a series of pairwise t -tests on the marginal means for the Angle conditions. In principle there is nothing wrong with this approach, but it can unnecessarily inflate the probability of committing a Type I Error since multiple t -tests would need to be conducted. Moreover, this approach does not necessarily capture the overall pattern of mean differences in the data.

A better approach would be to perform a linear contrast to test a hypothesized pattern of means. In this situation, we might want to test whether there is a significant linear trend such that RT increases linearly with degrees of visual angle. Moreover, using this approach we can quantify how much variability among condition means is due to this linear trend. Such analyses present an informative overall picture of the pattern of condition means.

```
# Collapse across Noise factor for each participant
by_subj.angle<-group_by(angle.long,Subject,Angle)
angle.subj.marg<-summarize(by_subj.angle,meanRT=mean(RT))

# Convert from long to wide format
```

```
angle.marg.wide<-spread(angle.subj.marg,key=Angle,value=meanRT)
angle.marg.wide<-angle.marg.wide[,-1] # Drop Subject column
print(angle.marg.wide)
```

```
## # A tibble: 10 x 3
##   `0`   `4`   `8`
##   <dbl> <dbl> <dbl>
## 1 450   510   630
## 2 390   480   540
## 3 570   630   660
## 4 450   660   720
## 5 510   660   630
## 6 360   450   450
## 7 510   600   720
## 8 510   660   780
## 9 510   660   660
## 10 510   540   660
```

```
# Perform linear trend analysis
lin.weights<-c(-1,0,1)
lin.mat<-as.matrix(angle.marg.wide)%*%lin.weights # Make sure data matrix is first!
print(lin.mat)
```

```
##      [,1]
## [1,] 180
## [2,] 150
## [3,]  90
## [4,] 270
## [5,] 120
## [6,]  90
## [7,] 210
## [8,] 270
## [9,] 150
## [10,] 150
```

```
# Test whether subject-level composite scores differ from 0 using one-sample t-test
t.test(lin.mat)
```

```
##
## One Sample t-test
##
## data:  lin.mat
## t = 8.1588, df = 9, p-value = 1.891e-05
## alternative hypothesis: true mean is not equal to 0
## 95 percent confidence interval:
## 121.4193 214.5807
## sample estimates:
## mean of x
##      168
```

The results indicate that there is indeed a highly-significant linear trend among the condition means such that RT increases linearly with degrees of visual angle. How much variability among the condition means is

accounted for by this linear trend? To answer this question, we need to compute $R_{Alerting}^2$ using the formula:

$$R_{Alerting}^2 = SS_{\psi} / SS_{Angle}$$

where SS_{ψ} is calculated using the formula:

$$SS_{\psi} = nb(\psi^2) / \sum_{j=1}^a c_j^2$$

and ψ is equal to:

$$\psi = \sum c_j \bar{Y}_j$$

```
n<-10 # Number of participants
b<-2 # Levels of factor B
angle.marg.mean<-colMeans(angle.marg.wide)
print(angle.marg.mean)
```

```
## 0 4 8
## 477 585 645
```

```
psi<-(-477)+0*585+645
ss.contrast<-n*b*(psi^2)/sum(lin.weights^2) # See equation 12.9 in MDK
ss.angle<-289920 # Taken from ANOVA table
```

```
r2.alerting<-ss.contrast/ss.angle
print(round(r2.alerting,digits=2))
```

```
## [1] 0.97
```

Therefore the linear trend accounts for 97% of the variance among the marginal means of Angle!

Further Investigation of Significant Interactions

The significant interaction informs us that the effect of one IV is not the same at each level of the other IV. However, to better specify these differences, we need to conduct *simple main effect* tests in which the effect of one IV is examined at each level of the other IV. Let's test the simple main effect of Angle at each level of the Noise variable. This amounts to conducting two one-way repeated measures ANOVAs for Angle - one with noise present and one with noise absent.

```
nz.present<-filter(angle.long,Noise=="Present") # get Noise Present data
nz.absent<-filter(angle.long,Noise=="Absent") # get Noise Absent data
```

```
ezANOVA(data=nz.present,dv=RT,wid=Subject,
         within=Angle,detailed=TRUE) # Do ANOVA looking at effect of Angle for Noise Present
```

```
## $ANOVA
##      Effect DFn DFd      SSn      SSd      F      p p<.05
## 1 (Intercept) 1 9 12211320 257880 426.1745 6.864246e-09 *
## 2 Angle 2 18 371760 43440 77.0221 1.502077e-09 *
## ges
## 1 0.9759188
```



```
## 2 0.5523266
##
## $`Mauchly's Test for Sphericity`
##   Effect      W      p p<.05
## 2 Angle 0.710601 0.2549783
##
## $`Sphericity Corrections`
##   Effect      GGe      p[GG] p[GG]<.05      HFe      p[HF] p[HF]<.05
## 2 Angle 0.7755551 7.853739e-08      * 0.9069206 7.731464e-09      *
```

```
ezANOVA(data=nz.absent,dv=RT,wid=Subject,
         within=Angle,detailed=TRUE) # Do ANOVA looking at effect of Angle for Noise Absent
```

```
## $ANOVA
##      Effect DFn DFd      SSn      SSd      F      p p<.05
## 1 (Intercept)  1   9 7500000 110400 611.413043 1.388034e-09      *
## 2 Angle  2  18  23280  41520  5.046243 1.820292e-02      *
##      ges
## 1 0.9801462
## 2 0.1328767
##
## $`Mauchly's Test for Sphericity`
##   Effect      W      p p<.05
## 2 Angle 0.8419927 0.5026126
##
## $`Sphericity Corrections`
##   Effect      GGe      p[GG] p[GG]<.05      HFe      p[HF] p[HF]<.05
## 2 Angle 0.8635524 0.02408953      * 1.04986 0.01820292      *
```

Therefore, each simple main effect test was statistically significant. At this point, we could proceed to do a series of paired *t*-tests using the `pairwise.t.test` function to compare the cell means of the Angle variable separately for each level of Noise. Alternatively, we could instead perform linear contrasts to test for a specific pattern of Angle cell means within each Noise condition. Let's use the latter approach by specifying two linear contrasts - one for each level of the Noise variable.

```
nz.present.weights<-c(-1,0,1)
nz.absent.weights<-c(-1,.5,.5)

nz.present.wide<-angle.data[,4:6] # Get columns 4-6 from angle.data
nz.absent.wide<-angle.data[,1:3] # Get columns 1-3 from angle.data

# Test linear contrast for Noise Present
nz.present.linscores<-as.matrix(nz.present.wide)%*%nz.present.weights
t.test(nz.present.linscores)
```

```
##
## One Sample t-test
##
## data:  nz.present.linscores
## t = 9.9254, df = 9, p-value = 3.808e-06
## alternative hypothesis: true mean is not equal to 0
## 95 percent confidence interval:
```

```
## 208.4627 331.5373
## sample estimates:
## mean of x
##      270

# Test linear contrast for Noise Absent
nz.absent.linscores<-as.matrix(nz.absent.wide)%*%nz.absent.weights
t.test(nz.absent.linscores)

##
## One Sample t-test
##
## data:  nz.absent.linscores
## t = 3.9428, df = 9, p-value = 0.003392
## alternative hypothesis: true mean is not equal to 0
## 95 percent confidence interval:
##  24.29637 89.70363
## sample estimates:
## mean of x
##      57
```

Note that the last contrast (for the `nz.absent` data) was decided after looking at the data. In this case, we could use Scheffe's F to control $\alpha_{FW} = .05$. The key thing to note here is that for single degree of freedom contrasts, $F = t^2$, so we'll need to convert our t -value to an F -statistic, and then compare this F -statistic to Scheffe's F . The following code will illustrate how to do this:

```
n.cond<-3
n.subj<-10

t.obs<-3.9428 # Taken from output above
F.obs<-t.obs^2 # Transfrom t-statistic to F-statistic
print(F.obs)
```

```
## [1] 15.54567
```

```
F.crit<-qf(p=.95,df1=n.cond-1,df2=(n.cond-1)*(n.subj-1)) # Get critical F-value
F.scheffe<-(n.cond-1)*F.crit # Calculate F-scheffe
F.obs>F.scheffe
```

```
## [1] TRUE
```

Our observed F -value is $>$ F -scheffe, so we can safely reject the null even when controlling $\alpha_{FW} = .05$.

Split-Plot ANOVA

Model Comparisons

As alluded to earlier, split-plot designs involve the mixture of between- and within-subjects factors usually in a factorial combination. Such designs are fairly commonplace in experimental psychology so we will cover

them here. For a design that contains one between-subjects (A) and one within-subjects factor (B), the full model has the following form:

$$Y_{ijk} = \mu + \alpha_j + \beta_k + \pi_{i/j} + (\alpha\beta)_{jk} + (\beta\pi)_{ki/j} + \epsilon_{ijk}$$

where μ is the grand mean parameter, α_j is the effect of being in group j of the between-subjects factor (A), β_k is the effect of being in condition k of the within-subjects factor (B), $(\alpha\beta)_{jk}$ is the effect of the interaction between the j -th level of A and the k -th level of B , $(\beta\pi)_{ki/j}$ is the effect of the interaction of the k -th level of B and the i -th subject in the j -th level of A , and ϵ_{ijk} is the error for the i -th subject in the j -th level of A and the k -th level of B . The i/j notation signifies that subjects are nested within the levels of factor A (the between-subjects factor).

Importantly, the α_j , β_k , and $(\alpha\beta)_{jk}$ terms represent fixed effects, whereas the $\pi_{i/j}$ and $(\beta\pi)_{ki/j}$ terms are random effects. In a split-plot design, we are primarily interested in evaluating the fixed effect terms within the model, which can be done by comparing the full model with a restricted model that omits the term of interest. As per usual, this comparison is achieved using the familiar formula for an F -test:

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

To evaluate the main effect of factor A (α_j 's), the F -statistic has the following form:

$$F_A = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_A}{MS_{S/A}}$$

where $MS_{S/A}$ is the mean-square term associated with subjects nested within group j and corresponds to the $\pi_{i/j}$ term in the full model.

The appropriate degrees of freedom are as follows:

$$\begin{aligned}(df_R - df_F) &= a - 1 \\ df_F &= N - a\end{aligned}$$

To evaluate the main effect of factor B (β_k 's), the F -statistic has the following form:

$$F_B = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_B}{MS_{B \times S/A}}$$

where $MS_{B \times S/A}$ represents the mean square for the interaction between factor B and subject i nested within group j of factor A , and corresponds to the $(\beta\pi)_{ki/j}$ term in the full model.

The appropriate degrees of freedom are as follows:

$$\begin{aligned}(df_R - df_F) &= b - 1 \\ df_F &= (N - a)(b - 1)\end{aligned}$$

Finally, to evaluate the interaction between factors A and B ($\alpha\beta_{jk}$'s), the F -statistic has the following form:

$$F_{A \times B} = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_{A \times B}}{MS_{B \times S/A}}$$

The appropriate degrees of freedom are as follows:

$$\begin{aligned}(df_R - df_F) &= (a - 1)(b - 1) \\ df_F &= (N - a)(b - 1)\end{aligned}$$

Example #2

This example is similar to Example #1 above, but now we are interested in how Age impacts the time required to detect visual targets at different degrees of visual angle. To address this question, the experimenter recruits 10 young adults and 10 older adults and presents visual targets at either 0, 4, or 8 degrees of visual angle. This experiment constitutes a 2×3 split-plot design with one between-subjects factor (Age), and one within-subjects factor (Angle).

```
age.angle.data<-read.table(file="Ch12T15.txt",header=T)
str(age.angle.data)

## 'data.frame': 60 obs. of 4 variables:
## $ Subject: int 1 2 3 4 5 6 7 8 9 10 ...
## $ Age : Factor w/ 2 levels "Old","Young": 2 2 2 2 2 2 2 2 2 2 ...
## $ Angle : int 0 0 0 0 0 0 0 0 0 0 ...
## $ RT : int 450 390 570 450 510 360 510 510 510 510 ...

# Convert Subject and Angle to a factor
age.angle.data$Subject<-as.factor(age.angle.data$Subject)
age.angle.data$Angle<-as.factor(age.angle.data$Angle)

# Compute condition means
age.angle.means<-age.angle.data %>% group_by(Age,Angle) %>% summarize(Mean=mean(RT))
print(age.angle.means)

## # A tibble: 6 x 3
## # Groups: Age [?]
## Age Angle Mean
## <fct> <fct> <dbl>
## 1 Old 0 543
## 2 Old 4 654
## 3 Old 8 792
## 4 Young 0 477
## 5 Young 4 585
## 6 Young 8 645

# Plot the data
with(age.angle.data,interaction.plot(x.factor=Angle,trace.factor=Age,response=RT,las=1,
cex=1.5,bty="n",type="b",pch=c(17,19),
xlab="Degrees of Visual Angle",ylab="RT (ms)"))

axis(side=1,at=c(1,2,3),labels=c("0","4","8")) # add in x-axis with ticks and labels
```

To conduct a split-plot ANOVA, we could again use the `aov` function. Note that we specify which of the two factors is within-subjects using the `Error(Subject/Angle)` argument.

```
# Compute 2 X 3 Split-Plot ANOVA using aov
age.angle.aov<-aov(RT~Age*Angle+Error(Subject/Angle),data=age.angle.data)
summary(age.angle.aov)

##
## Error: Subject
```

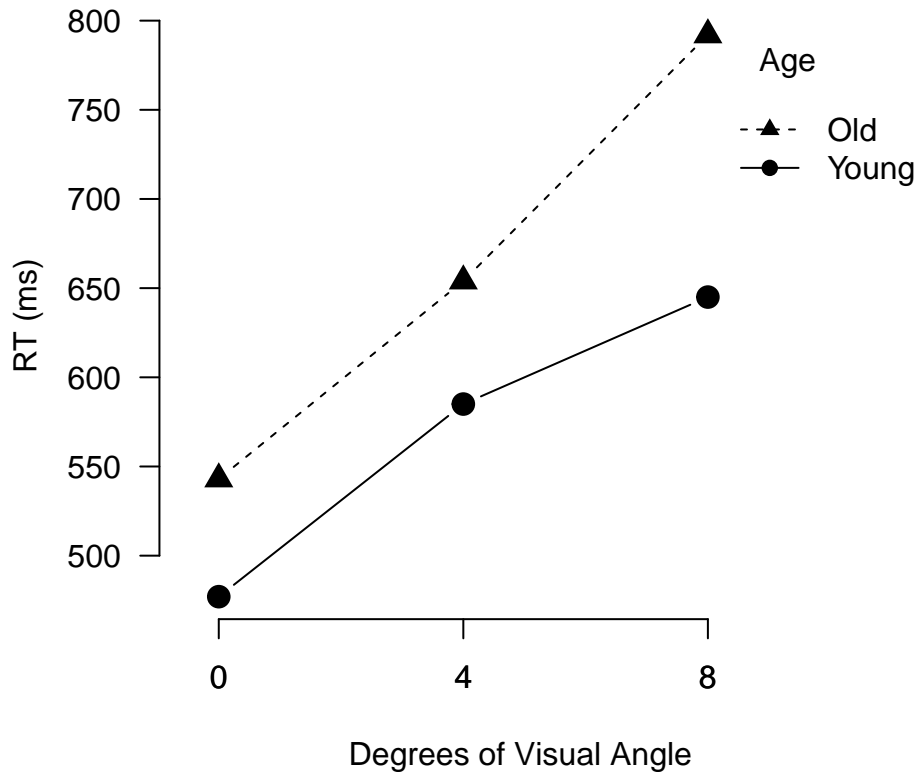


Figure 2: Plot of data from MDK Chapter 12 Table 15

```
##           Df Sum Sq Mean Sq F value Pr(>F)
## Age       1 132540 132540   7.276 0.0147 *
## Residuals 18 327900 18217
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Error: Subject:Angle
##           Df Sum Sq Mean Sq F value Pr(>F)
## Angle     2 435090 217545 143.911 < 2e-16 ***
## Age:Angle  2  21090  10545   6.976 0.00275 **
## Residuals 36  54420   1512
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

As you can see, the analysis revealed significant main effects of Age, $F(1, 18) = 7.28, p = .015$ and Angle, $F(2, 36) = 143.91, p < .001$, as well as a significant Age \times Angle interaction, $F(2, 36) = 6.98, p = .0028$.

However, given that Angle has three levels, we need to consider the possibility that the sphericity assumption may be violated. Let's redo the ANOVA using the `ezANOVA` function from the `ez` package as it provides us with a test of the sphericity assumption.

```
# Compute 2 X 3 Split-Plot ANOVA using ezANOVA function
ezANOVA(data=age.angle.data, dv=RT, wid=Subject, between=Age, within=Angle, detailed=TRUE)
```

```
## $ANOVA
##      Effect DFn DFd      SSn      SSD      F      p p<.05
```

```

## 1 (Intercept) 1 18 22767360 327900 1249.809332 4.375573e-18 *
## 2 Age 1 18 132540 327900 7.275755 1.473401e-02 *
## 3 Angle 2 36 435090 54420 143.910695 6.728921e-18 *
## 4 Age:Angle 2 36 21090 54420 6.975744 2.751523e-03 *
## ges
## 1 0.98348487
## 2 0.25742920
## 3 0.53227878
## 4 0.05227932
##
## $`Mauchly's Test for Sphericity`
## Effect W p p<.05
## 3 Angle 0.936837 0.5743071
## 4 Age:Angle 0.936837 0.5743071
##
## $`Sphericity Corrections`
## Effect GGe p[GG] p[GG]<.05 HFe p[HF]
## 3 Angle 0.9405895 5.769113e-17 * 1.046677 6.728921e-18
## 4 Age:Angle 0.9405895 3.399421e-03 * 1.046677 2.751523e-03
## p[HF]<.05
## 3 *
## 4 *

```

Since the Mauchly's test is not significant, we can interpret the raw F - and p -values. As you can see, there was a significant Age \times Angle interaction, implying that we should be conducting follow-up simple main effects tests.

Simple Main Effects

Within-Subjects Effects at Each Level of Between-Subjects Factor

Suppose we wished to test the simple main effect of Angle for both levels of Age. This would amount to conducting two one-way repeated-measures ANOVAs that examine the effect of Angle separately for younger and older adults. To perform this analysis in *R*, we first need to divide the data into two separate variables using the `filter` command based on the level of Age. Then, we simply conduct a one-way repeated-measures ANOVA on each subset.

```

young.angle<-filter(age.angle.data,Age=="Young") # Create subset of data for young adults
old.angle<-filter(age.angle.data,Age=="Old") # Create subset of data for older adults

# Perform one-way RM-ANOVA for young adults
ezANOVA(data=young.angle,dv=RT,wid=Subject,within=Angle,detailed=TRUE)

```

```

## Warning: You have removed one or more Ss from the analysis. Refactoring
## "Subject" for ANOVA.

```

```

## $ANOVA
## Effect DFn DFd SSn SSd F p p<.05 ges
## 1 (Intercept) 1 9 9712830 146070 598.4492 1.526600e-09 * 0.9819926
## 2 Angle 2 18 144960 32040 40.7191 2.086763e-07 * 0.4486953
##
## $`Mauchly's Test for Sphericity`

```

```
## Effect W p p<.05
## 2 Angle 0.960106 0.8497219
##
## $`Sphericity Corrections`
## Effect GGe p[GG] p[GG]<.05 HFe p[HF] p[HF]<.05
## 2 Angle 0.9616365 3.401705e-07 * 1.217564 2.086763e-07 *
```

```
# Perform one-way RM-ANOVA for older adults
ezANOVA(data=old.angle,dv=RT,wid=Subject,within=Angle,detailed=TRUE)
```

```
## Warning: You have removed one or more Ss from the analysis. Refactoring
## "Subject" for ANOVA.
```

```
## $ANOVA
## Effect DFn DFd SSn SSd F p p<.05
## 1 (Intercept) 1 9 13187070 181830 652.7175 1.038185e-09 *
## 2 Angle 2 18 311220 22380 125.1555 2.752367e-11 *
## ges
## 1 0.9847505
## 2 0.6038065
##
## $`Mauchly's Test for Sphericity`
## Effect W p p<.05
## 2 Angle 0.8142084 0.4394832
##
## $`Sphericity Corrections`
## Effect GGe p[GG] p[GG]<.05 HFe p[HF] p[HF]<.05
## 2 Angle 0.8433185 8.091608e-10 * 1.016384 2.752367e-11 *
```

Because the test for each simple main effect above is simply a one-way repeated measures ANOVA, subsequent comparisons among individual means could use the procedures outlined in the Lab 8 Tutorial.

Between-Subjects Effects at Each Level of Within-Subjects Factor

Alternatively, we could test the simple main effect of Age at each level of the Angle variable. Indeed from Fig. 2 it appears as though the RT difference between young and older adults is greatest at 8°. Again, let's proceed by dividing the data into groups based on levels of Angle again using the **filter** command.

```
a0<-filter(age.angle.data,Angle=="0")
a4<-filter(age.angle.data,Angle=="4")
a8<-filter(age.angle.data,Angle=="8")

a0.aov<-aov(RT~Age,data=a0)
summary(a0.aov)
```

```
## Df Sum Sq Mean Sq F value Pr(>F)
## Age 1 21780 21780 3.161 0.0923 .
## Residuals 18 124020 6890
## ---
## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
a4.aov<-aov(RT~Age,data=a4)
summary(a4.aov)
```

```
##           Df Sum Sq Mean Sq F value Pr(>F)
## Age         1  23805   23805    3.35 0.0838 .
## Residuals   18 127890    7105
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
a8.aov<-aov(RT~Age,data=a8)
summary(a8.aov)
```

```
##           Df Sum Sq Mean Sq F value Pr(>F)
## Age         1 108045  108045   14.91 0.00114 **
## Residuals   18 130410    7245
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

It is clear from these analyses that the simple main effect of Age is significant, $F(1,18) = 14.91, p = .001$, only when targets are presented at 8° of visual angle.