

Cheat Sheet: factorial ANOVA

Measurement and Evaluation of HCC Systems

Scenario

Use factorial ANOVA if you want to test the effect of two (or more) nominal variables `varX1` and `varX2` on a continuous outcome variable `varY`. In this scenario `varX1` and `varX2` are usually orthogonally manipulated experimental manipulation with two or more conditions (but they can be any nominal variable, such as gender, race, occupation, etc). Each participant is randomly assigned to one condition.

Power analysis

- To do a power analysis for the omnibus test, use Test family “*F* tests”, “ANOVA: Fixed effects, special, main effects and interactions”.
- A power analysis has four variables: Effect size, α (usually .05), power (usually .85), and N . If you know three of these, G*Power will calculate the fourth. Select the correct type of power analysis, based on the information you have, and what you want to find out.
- “Numerator df” is the number of levels you are comparing minus 1. For the main effect of an X variable with k levels, this is $k-1$; for an interaction effect of two X variables with k and r levels, this is $(k-1)(r-1)$.
- “Number of groups” is the total number of experimental conditions; e.g. for the aforementioned two X variables it is $k*r$.
- By clicking on “Determine”, you can compute the effect size f from the partial η^2 .
- Click on “Calculate” to calculate the missing parameter.
- To conduct power analyses for contrasts or post hoc tests, refer to power analysis of the independent t test. Use a corrected α if needed.

Plotting a line plot and a box plot

- Use the `ggplot2` package to plot a bar chart with error bars.

```
ggplot(data, aes(varX1, varY, color=varX2)) + stat_summary(fun.y=mean,  
  geom="line", aes(group = varX2) + stat_summary(fun.data=mean_cl_normal,  
  geom="errorbar", width=0.2)
```
- Plot a boxplot:

```
ggplot(data, aes(varX1, varY)) + geom_boxplot() + facet_wrap(~varX2)
```

Pre-testing assumptions

- In an ANOVA, Y should be independent, continuous, and unbounded. The error variance should be normally distributed, which is true if Y is normally distributed for each level of X . The ANOVA test is fairly robust against violations of normality. Variances are assumed to be equal across groups (homoscedasticity).
- If your N is small:
 - o Test for significant skewness, kurtosis, and Shapiro-Wilk test within each group using `stat.desc` in the `pastecs` package and the `by` function.
`by(data$varY, list(data$varX1, data$varX2), stat.desc, basic=F, norm=T)`
 - o Multiply `skew.2SE` and `kurt.2SE` by 2 to get the Z-scores of skewness and kurtosis. Compare these values to typical cut-off values ($Z > \pm 1.96: p < .05$, $Z > \pm 2.58: p < .01$, $Z > \pm 3.29: p < .001$). The significance of the Shapiro-Wilk test is listed under `normtest.p`.
 - o Test the assumption of homoscedasticity, conduct Levene's test in the `car` package.
`leveneTest(varY ~ varX1*varX2, data=data, center=median)`
 - o A significant Levene's test means that the data is heteroscedastic. In this case, Welch's ANOVA should be used instead of a regular ANOVA.
- If your N is large:
 - o Subset the data into multiple groups based on `varX1` and `varX2`.
`ADdata <- subset(data, data$varX1 == "A" & data$varX2 == "D")`
`AEdata <- subset(data, data$varX1 == "A" & data$varX2 == "E")`
etc...
 - o Draw the histogram for `varY` in `ADdata`, overlaid with a normal curve (using `ggplot2`), and visually inspect whether it follows the normal distribution:
`ggplot(ADdata, aes(varY)) + geom_histogram(aes(y=..density..),
binwidth=1, color="black", fill="white") + stat_function(fun = dnorm,
args = list(mean = mean(ADdata$varY), sd = sd(ADdata$varY)))`
 - o Change the `binwidth` setting based on what is suitable for your data.
 - o Draw normal a Q-Q plot, and visually inspect whether the data follows the diagonal line:
`qplot(sample = ADdata$varY, stat="qq")`
 - o Do the same for `AEdata`, etc.
 - o Visually inspect the boxplot to assess the level of heteroscedasticity. You can also use the variance ratio test (Hartley's F_{max}). This is the ratio of the variances between the group with the biggest variance and the group with the smallest variance.
`max(by(data$varY, list(data$varX1, data$varX2), var)) /
min(by(temp$varY, list(data$varX1, data$varX2), var))`
 - o Compare this value against Figure 5.8 in the book.
- If your data has positive skew, and your data only has positive values, you can possibly fix this by transforming your Y variable:

- Log transform:
`data$varYlog <- log(data$varY + 1)`
- Or, square root transform:
`data$varYsqrt <- sqrt(data$varY)`
- In other cases of violations of assumptions (e.g. outliers), you can conduct a robust test (see below).

(optional) Preparing contrasts

- If you have specific hypotheses about where the differences between conditions exist, you can run tests as planned contrasts.
- See the cheat sheet for ANOVA on how to prepare contrast for each of your X variables.

Running the test

- Run the factorial ANOVA as follows:
`model1 <- aov(varY ~ varX1 * varX2, data = data)`
- Get the model summary using Type III sum of squares (if you expect an interaction effect; requires orthogonal contrasts) or Type II sum of squares (if you expect no interaction effect):
`Anova(model1, type=3)`
- The output gives the F -statistic and its p -value for the main effects and the interaction effect(s). A significant interaction effect means that the effect of `varX1` is different at different levels of `varX2`, and vice versa. The effects of lower-level effects are not interpretable if the (higher-level) interaction effect is significant.
- You can get the effect sizes omega-squared using the `omega_aov` function I created:
`omega_aov(model1)`

(optional) Robust versions

- You can use functions in the `WRS2` package to run trimmed and/or bootstrapped ANOVAs.
- Robust two-way ANOVA using 10% trimmed means (change the percentage if desired):
`t2way(varY ~ varX1 * varX2, data = data, tr = 0.1)`
- Robust two-way ANOVA using M-measures on a bootstrapped median (may use “mom” instead of “median”):
`pbad2way(varY ~ varX1 * varX2, data = data, est="median", nboot = 2000)`

(optional) Evaluate contrasts

- Evaluate your planned contrasts using the `summary.lm` function.
`summary.lm(model1)`
- In terms of overall model fit, this output will give you the R-squared, which is the proportion of the variance of `varY` explained by the model. In ANOVA this value is called eta-squared.

- Lower-level effects are not interpretable if the (higher-level) interaction effect is significant. If the interaction effect is not significant, then the main effects are the effects of an X on Y , given the other $X(s)$.
- Interaction effects between two contrasts represent the difference in the effect of one contrast at different levels of the other contrast (and vice versa). Refer to the ANOVA cheat sheet on how to interpret the main effect contrast coefficients.
- Each coefficient has a t test and a p -value to test if the effect is significant. Multiply the p -value by 2 if you were conducting a one-sided test (this is very likely for planned contrasts!). You can get the effect size r using the formula $r = \sqrt{t^2 / t^2 + df}$.
- You can get confidence intervals for the coefficients using the `confint` function:
`confint(model1)`
- (optional) see the ANOVA cheat sheet for how to correct your critical p -value in case you used non-orthogonal contrasts.

(optional) Simple effects

- You can test the effect of `varX2` (levels: D and E) at each level of `varX1` (levels: A, B, and C) with a simple effects test.
- First create an interaction variable:
`data$simple <- interaction(data$X1,data$X2)`
- Then create dummies for the contrasts:
`AvBC <- c(-2/3, 1/3, 1/3, -2/3, 1/3, 1/3)`
`BvC <- c(0, -1/2, 1/2, 0, -1/2, 1/2)`
`X2_A <- (-1/2, 0, 0, 0, 1/2, 0, 0)`
`X2_B <- (0, -1/2, 0, 0, 0, 1/2, 0)`
`X2_C <- (0, 0, -1/2, 0, 0, 0, 1/2)`
`contrasts(data$simple) <- cbind(AvBC, BvC, X2_A, X2_B, X2_C)`
`simpleModel <- aov(varY ~ simple, data=data)`
`summary.lm(simpleModel)`
`confint(simpleModel)`

(optional) Post-hoc tests

- You can only do post-hoc tests and robust post-hoc tests on the main effects. These are the same as for ANOVA (see the ANOVA cheat sheet).

(optional) Effect sizes of specific differences

- Get means, *sds*, and *ns* from `stat.desc` in the `pastecs` package.
`desc <- by(data$varY, list(data$varX1, data$varX2), stat.desc)`
- Plug the values into `mes`, e.g. to test the differences between conditions D and E in group A:
`mes(desc$[[“A”,“D”]][“mean”], desc$[[“A”,“E”]][“mean”],`

```
desc$[["A","D"]]["std.dev"], desc$[["A","E"]]["std.dev"],  
desc$[["A","D"]]["nbr.val"], desc$[["A","E"]]["nbr.val"])
```

- It is common to report d instead of r .

Reporting

- Use the following format to report on an ANOVA (replace the full names (not just the variable names) of A, B and varY, and replace the xx'es with the actual numbers).
- For the main effect tests: "There was a significant effect of [varX1] on [varY], $F(x, x) = x.xx$, $p = .xxx$, $\omega^2 = .xx$."
- For the interaction effect and simple effects: "There was a significant interaction effect between [varX1] and [varX2] on [varY], $F(x, x) = x.xx$, $p = .xxx$, $\omega^2 = .xx$. Specifically, there was a significant difference between [D] ($M = xx.xx$, $SD = xx.xx$) and [E] ($M = xx.xx$, $SD = xx.xx$) in condition [A], $d = x.xx$; between [D] ($M = xx.xx$, $SD = xx.xx$) and [E] ($M = xx.xx$, $SD = xx.xx$) in condition [B], $d = x.xx$; and between [D] ($M = xx.xx$, $SD = xx.xx$) and [E] ($M = xx.xx$, $SD = xx.xx$) in condition [C], $d = x.xx$."
- For reporting contrasts and post-hoc tests, see the ANOVA cheat sheet.