Cheat Sheet: repeated and mixed ANOVA

Measurement and Evaluation of HCC Systems

Scenario

Use repeated or mixed ANOVA if you want to test the effect of one or more nominal variables `varX1`, `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). In repeated and mixed ANOVA, one or more of the *X* variables can be manipulated within-subjects (i.e. each participant receives *all* conditions of that *X* variable). This cheat sheet assumes that `data` is in the long format, with a user `id` to tie together the repeated measures. (if it is not, use the `melt` function in the `reshape2` package to fix that).

Power analysis

- The power analysis for a repeated or mixed ANOVA is the same as for a factorial ANOVA, except the statistical test depends on the effect to be tested: "ANOVA: Repeated measures, between factors" for any between-subjects effects; "ANOVA: Repeated measures, within factors" for any within-subjects effects; and "ANOVA: Repeated measures, within-between interaction" for any interaction effects that cross both levels.

Plotting a line plot and a box plot

- Reshape the data into wide format (add additional repeated variables to `v.names` if needed):
  ```
  dataWide <- reshape(data, idvar="id", timevar="varX", direction="wide",
  v.names=c("varY"))
  ```
- Remove the between-subjects differences from the data (include `varY.C`, `varY.D`, etc if you have more than 2 within-subjects categories):
  ```
  dataAdjusted <- dataWide
  dataAdjusted$varY.A <- dataWide$varY.A - (dataWide$varY.A +
  dataWide$varY.B)/2 + mean((dataWide$varY.A + dataWide$varY.B)/2)
  dataAdjusted$varY.B <- dataWide$varY.B - (dataWide$varY.A +
  dataWide$varY.B)/2 + mean((dataWide$varY.A + dataWide$varY.B)/2)
  ```

- Stack the data back into long format:
  `dataStack <- stack(dataAdjusted)`
  `names(dataStack) <- c("varY","varX")`
- Refer to the ANOVA and factorial ANOVA cheat sheets to create plots for `dataStack`.

Pre-testing assumptions

- For between-subjects factors, refer to Pre-testing assumptions in the ANOVA and/or factorial ANOVA cheat sheets.
- For within-subjects factors, the main assumption is sphericity. Mauchly's test of sphericity is included as part of the `ezANOVA` output (see below).
- If you have sphericity, you can report on the regular ANOVA results, and conduct any post-hoc tests you seem fit.
- If you do not have sphericity, you have to report the Greenhouse-Geisser or Huynh-Feldt corrected ANOVA results (the former is a bit more conservative), and you can only conduct Bonferonni corrected post-hoc tests.

(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 ANOVA using the `ez` package (list more factors as `within` or `between` if needed, remove `between` if all factors are within; use `type=2` if you expect no interaction effect):
  `ezANOVA(data = data, dv=.(varY), wid=.(id), within=.(varX1), between=.(varX2), detailed=T, type=3)`
- Interpreting the results: first check to see if Mauchly's Test of Sphericity is significant.
- For effects for which the sphericity test is not significant, report the *F* value, with `DFn` and `DFd` degrees of freedom, the *p*-value, and the generalized η^2 (`ges`).
- For effects for which it is significant, report the regular *F* value but divide `DFn` and `DFd` by `GGe`. Report `p[GG]` as the *p*-value. Report the regular generalized η^2 (`ges`).

(optional) Robust versions

- You can use functions in the `WRS2` package to run trimmed and/or bootstrapped repeated and mixed ANOVAs.
- Load WRS: (see installation instructions at https://github.com/nicebread/WRS)

- Load the "source" version of WRS2:
  `install.packages("WRS2", type="source")`
- Robust repeated ANOVA using 10% trimmed means (change the percentage if desired):
  `rmanova(data$varY, data$varX, data$id, tr = 0.1)`
- Robust repeated ANOVA using 10% trimmed means and 2000 bootstrap samples:
  `rmanova(data$varY, data$varX, data$id, tr = 0.1, nboot = 2000)`
- Robust factorial repeated or mixed ANOVA using 10% trimmed means (change the percentage if desired):
  `tsplit(varY ~ varX1 * varX2, id=id, data=data, tr = 0.1)`
- Robust mixed ANOVA using 10% trimmed means and 2000 bootstrap samples:
  Factor A: `sppba(varY ~ varX1 * varX2, id=id, data=data, nboot=2000)`
  Factor B: `sppbb(varY ~ varX1 * varX2, id=id, data=data, nboot=2000)`
  Interaction: `sppbi(varY ~ varX1 * varX2, id=id, data=data, nboot=2000)`
  This only works for at most 2 independent variables.

(optional) Evaluate contrasts

- Evaluate your planned contrasts by running an `lme` (using the `nlme` package).
- First run the baseline model (assuming `varX1` is the only within-subjects variable, otherwise you should use `random = ~1|id/varX1/varX2` etc.):
  `baseline <- lme(varY ~ 1, random = ~1|id/varX1, data=data, method="ML")`
- Add the main effect of the first variable:
  `model1 <- update(baseline, .~. + varX1)`
- If factorial or mixed, add the second variable (and more if needed, one by one):
  `model2 <- update(model1, .~. + varX2)`
- Finally, if factorial or mixed, add the interaction(s):
  `model3 <- update(model2, .~. + varX1*varX2)`
- Conduct an ANOVA comparison of all models:
  `anova(baseline, model1, model2, model3)`
- Get the summary of the final model:
  `summary(model3)`
- For interpretation of the contrasts, refer to "Evaluate contrast" in the ANOVA and factorial ANOVA cheat sheets.

(optional) Post-hoc tests

- If you do not have specific hypotheses about where specific differences between conditions exist, you can test all of them using post-hoc tests. Note that when you do not have sphericity, you can only do the Bonferroni-corrected tests.

- Bonferroni-corrected tests (use "holm" for Holm, and "BH" for Bejamini-Hochberg):
```
pairwise.t.test(data$varY, data$varX, paired=T, p.adjust.method =
"Bonferroni")
```
- Tukey post-hoc tests and confidence intervals using the package `multicmp` (can also conduct "Dunnett" post-hoc test); this only works on `lme` models, so you should construct one of those first (see Evaluate contrasts):
```
post <- glht(model1, linfct=mcp(varX = "Tukey"))
summary(post)
confint(post)
```

(optional) Robust post-hoc tests

- Post-hoc tests using 10% trimmed means (confidence intervals are corrected; p-values are not):
```
rmmcp(data$varY, data$varX, data$id, tr = 0.1)
```
- Post-hoc tests using 10% trimmed bootstrapped means, using 2000 bootstrap samples:
```
pairdepb(data$varY, data$varX, data$id, tr = 0.1, nboot = 2000)
```

Reporting

- Use the following format to report on an ANOVA (replace the full names (not just the variable names) of varX1, varX2, and varY, and replace the xx'es with the actual numbers).
- ezANOVA Mauchly's test: "Mauchly's test indicated that the assumption of sphericity was violated for [varX1], W = x.xxx, p = .xxx, ε = .xx, and [varX2] W = x.xxx, p = .xxx, ε = .xx. The degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity."
- ezANOVA main results: "There was a significant main effect of [varX1] on [varY], F(x.xx, x.xx) = x.xx, p = .xxx, generalized η^2 = .xxx, a main effect of [varX2], F(x.xx, x.xx) = x.xx, p = .xxx, generalized η^2 = .xxx, and an interaction effect between [varX1] and [varX2], F(x.xx, x.xx) = x.xx, p = .xxx, generalized η^2 = .xxx. The latter indicates that [varX1] had different effects on [varY] depending on [varX2]."
- lme main results: "[varX1] had a significant effect on [varY] χ^2(x) = x.xx, p = .xxx, as did the effect of [varX2] χ^2(x) = x.xx, p < .xxx. The interaction between [varX1] and [varX2] was significant, χ^2(x) = x.xx, p < .xxx.
- lme contrast results: "Planned contrasts revealed that [group Q of contrast 1] [increased / decreased] [varY], compared to [group P of contrast 1], b = x.xx, t(xx) = x.xx, $p_{\text{one-tailed}}$ = .xxx, r = .xxx."
- lme contrast interaction results: "The effect of [varX2] on [varY] is significantly [smaller/larger] in [group Q of contrast 1] than in [group P of contrast 1], b = x.xx, t(xx) = x.xx, $p_{\text{one-tailed}}$ = .xxx, r = .xxx."