Helicopter ANOVAs II

Finally (!) we get around to analyzing our copter data *the correct way*. No more dinking around with t-tests and wrong ANOVAs – today we apply the full analysis the way we designed the experiment. We'll do that in a couple of steps, so you can see how those steps matter.

Reminder: We evaluated two treatments that potentially affect helicopter flight times: wing length and whether wings were folded or not. We conducted that experiment as a factorial design, where treatments were combined in all possible combinations, so we could explore interactive effects. Replicate copters were organized by Groups, and we accounted for the effect of stairs as a covariate.

1. Import, remove NAs, and attach our Copter data from
https://sciences.ucf.edu/biology/d4lab/wp-content/uploads/sites/125/2019/09/helicopter-data.txt

2. Note that I refer to this file as "data" hereafter.
3. Make factor variables (i.e., fwl, fbw, ffold, and fgroup – those names are assumed below) for WL, BW, Fold, and Group. Because we are doing ANOVAs, which assume categorical treatments, and ensure that these are correctly IDed.

Here we save time by skipping worries about homogeneity of variance and normality– just remember that you should always check those assumptions before your own ANOVAs.

3. Let's squint at potential interactions first:

```
interaction.plot(fbw, trace.factor=fwl, response=Time, fun=mean, type="b")
interaction.plot(ffold, trace.factor=fwl, response=Time, fun=mean, type="b")
interaction.plot(ffold, trace.factor=fgroup, response=Time, fun=mean,
type="b")
```

What do you see? What should we expect in the analyses?

An important note on the "block" effect: Each Group had 5 replicates per treatment combination. A simple randomized complete block experiment would have only one replicate of each treatment combination in a Group, but then could not evaluate a Group x Treatments interaction. So Group here was more than a *simple* block – we *can* evaluate how Group interacts with a treatment because we have replicates in a Group. But *should* we? Here we stick to the design (Group is a block), but the replicates allow us to also evaluate Group as a random effect, below.

A note about Steps: This is a covariate. We get to Analysis of Covariance later in the semester. Suffice to say now that this analysis is actually that – a combination of regression (steps effect) and anova (other treatments).

4. Run this code to analyze the experiment as we designed it:

```
copters1 <- lm(Time ~ Step + fwl*fbw*ffold + fgroup)
summary(copters1)
summary.aov(copters1)
```

Notice the code shortcut above? You ran a lm model and then asked for the aov table that comes from it. In only 3 lines of code. And remember that using * above actually means [fwl + fbw + ffold + fwl X fbw + fwl X ffold + fbw X ffold + fwl X fbw X ffold].

What can you conclude? Would you report this complete model in a paper?

There are (at least) two philosophies about this:

1. Standard: Report the most complete model as-is (e.g., `copters1` above). That is the study design, even if multiple terms are not significant. Truth in advertising, so to speak.
2. Crawley, Michael (2007) The R Book: Iteratively simplify by eliminating nonsignificant terms to report the most parsimonious model that represents significant effects. Crawley advocates simplification, in the spirit of AIC-based model selection - to report the most efficient model that works. If you do this, you should also explain which hypothesized terms were not significant, and that the final model and its adjusted R^2 are simplified.

   NOTE: not everyone agrees with this philosophy. For example it could report a partial view of a study, and emphasize effects based on p values rather than the actual experimental design.

But let's simplify to see what happens. Edit the `copters1` model (call it `copters2`) to *remove the weakest, nonsignificant term*. Run `copters2`. Still have nonsignificant terms? Repeat until only significant terms remain.

Did the Adjusted R^2 drop because we removed terms, or did it go up because we removed noise? In other words, did you retain the same signal, but do so more efficiently?

And now for a more advanced / appropriate approach that is a hint of things to come. Do we (a) *really* care about the effect of Group, or (b) was it just a noise factor we wanted to account for?

If (b), we can treat Group as a *random factor because you have replicates*. And what about Step? Is that a similar matter? A mixture of fixed (i.e., prescribed) and random factors is a mixed model. Here we will use the package lme4.

6. Install and load the `lme4` package if it is not already in your computer.
7. Then run this set of commands:

```
fstep <- factor(Step) # now we can treat this as a categorical group
copters3 <- lmer(Time ~ Step + fwl*fbw*ffold + (1|fgroup))
copters4 <- lmer(Time ~ fwl*fbw*ffold + GROUP + (1|fstep))
copters5 <- lmer(Time ~ fwl*fbw*ffold + (1|fstep) + (1|fgroup))
copters6 <- lmer(Time ~ 1 + (1|fgroup))
```

```
anova(copters3,copters4,copters5,copters6)
```
This set of commands runs linear mixed effect models, and compares them in an ANOVA based on AIC. The model with the lowest AICc value is the most plausible model.

Does the outcome make sense to you?

One last important step: we have to evaluate the residuals of our model to test for normality of the residuals and homogeneity of variance. *Because the error terms (residuals) are key for assumptions.* Here is a way to do that for one model:

```
plot(resid(copters3)~predict(copters3))
abline(0,0)
```

This plots the residuals (i.e., differences between observed and predicted) as a function of the predicted values.
An even scatter above/below the line indicates normality.
The same breadth of scatter from left to right indicates homogeneity of variance.
What do you think?