

Lecture 19: Repeated-Measures ANOVA Notation and Equations

Notation:

k : The number of measurements made of each subject; the number of levels of the treatment

M_i : The mean score for the i th measurement (i represents any number 1 through k)

M_s : The mean score for Subject s (s represents any number 1 through n)

SS_{subject} : The variability accounted for by individual differences in a repeated-measures ANOVA

Partitioning variability. Repeated-measures ANOVA works like simple ANOVA, in that we break the total variability into parts that are explainable and parts that are unexplainable. The difference is that there are now two explainable parts: variability explainable by the treatment (e.g., the independent variable) and variability explainable by individual differences (i.e., by the subjects).

$$SS_{\text{total}} = SS_{\text{subject}} + SS_{\text{treatment}} + SS_{\text{residual}} \quad (1)$$

The total variability is computed in the same way as for simple ANOVA. First, we calculate the grand mean as the average of all the data, treated as a single group. Then we find the differences between all the data and the grand mean, square those differences, and add them up.

$$SS_{\text{total}} = \sum_{\text{all scores}} (X - \bar{M})^2 \quad (2)$$

The variability due to the treatment is also computed in the same way as for simple ANOVA. First, we calculate the average, M_i , for each level of the treatment. Then we find the differences between these averages and the grand mean, and square those differences. Finally, we multiply each square by n (because it represents n raw scores, one from each subject) and add them up.

$$SS_{\text{treatment}} = \sum_{\text{levels } i} n(M_i - \bar{M})^2 \quad (3)$$

The variability due to individual differences is computed in essentially the same way as $SS_{\text{treatment}}$. In this case, we find the average score for each subject (instead of the average for each measurement), M_s , and calculate the squared differences from the mean. We multiply each square by k because it represents k raw scores (all the measurements from a single subject), and then add them up.

$$SS_{\text{subject}} = \sum_{\text{subjects } s} k(M_s - \bar{M})^2 \quad (4)$$

If you think of the whole dataset as a matrix, with rows representing subjects and columns representing measurements, then you can see how the formulas for $SS_{\text{treatment}}$ and SS_{subject}

work in exactly the same way. The only difference is that $SS_{\text{treatment}}$ is based on the columns (treatment levels) and SS_{subject} is based on the rows (subjects).

The residual variability is the variability that's left over after we remove $SS_{\text{treatment}}$ and SS_{subject} . There are ways to compute SS_{residual} directly, but the formulas are complicated and don't provide much insight into what's going on. So, the best way to think about how SS_{residual} is computed is by subtracting $SS_{\text{treatment}}$ and SS_{subject} from SS_{total} . Notice that this formula is essentially the same as Equation 1.

$$SS_{\text{residual}} = SS_{\text{total}} - SS_{\text{subject}} - SS_{\text{treatment}}$$

Hypothesis testing. Hypothesis testing in repeated-measures ANOVA works in the same way as in simple ANOVA. We convert $SS_{\text{treatment}}$ and SS_{residual} into mean squares, and then compute their ratio to get an F statistic. The only difference between the two tests is that, in repeated-measures ANOVA, we first remove the individual differences from the residuals. This was taken care of when we computed the sums of squares, so at this point there's nothing new. The equations for MS , F , F_{crit} , and p are shown here, but they are identical to the equations for simple ANOVA (see the 11/8 notes for explanations).

$$MS_{\text{treatment}} = \frac{SS_{\text{treatment}}}{df_{\text{treatment}}} \qquad MS_{\text{residual}} = \frac{SS_{\text{residual}}}{df_{\text{residual}}}$$

$$F = \frac{MS_{\text{treatment}}}{MS_{\text{residual}}}$$

$$p(F_{df_{\text{treatment}}, df_{\text{residual}}} \geq F_{\text{crit}}) = \alpha \qquad p = p(F_{df_{\text{treatment}}, df_{\text{residual}}} \geq F)$$

Degrees of Freedom. This section is optional but may be useful for understanding how the sums of squares fit together. Since there are n subjects each being measured k times, there are a total of nk raw scores. Therefore there are nk squares of the form $(X - \bar{M})^2$.

However, when we add them up to get SS_{total} , one of them disappears because \bar{M} depends on the X s. Therefore the total degrees of freedom in the data equals $nk - 1$.

$$df_{\text{total}} = nk - 1$$

$SS_{\text{treatment}}$ and SS_{subject} also appear to be sums of nk squares, but in this case there are actually far fewer squares that are independent of each other. For $SS_{\text{treatment}}$, every raw score contributes a square of the form $(M_i - \bar{M})^2$, where M_i is the mean of whichever treatment level that score came from. However, since there are only k measurements, meaning only k different M_i s, there are only k unrelated squares that go into $SS_{\text{treatment}}$ (each one gets repeated n times). Furthermore, because the grand mean is in the formula and it depends on the M_i s, one of those degrees of freedom disappears. Therefore $SS_{\text{treatment}}$ has $k - 1$ degrees of freedom. Another way to see this is to notice that $SS_{\text{treatment}}$ is just a multiple of the variance of the treatment means (i.e., of the M_i s), and that variance has $k - 1$ degrees of freedom because there are k treatment means.

$$df_{\text{treatment}} = k - 1$$

The degrees of freedom for SS_{subject} works in the same way. There are n different subject means, so there are n different squares $(M_s - \bar{M})^2$. One of them disappears because of the grand mean, so there are $n - 1$ degrees of freedom for SS_{subject} . This fits with the fact that SS_{subject} is a multiple of the variance of the M_s s, which has $n - 1$ degrees of freedom.

$$df_{\text{subject}} = n - 1$$

The residual degrees of freedom are what's left over from the total degrees of freedom after $SS_{\text{treatment}}$ and SS_{subject} are removed. Therefore, $df_{\text{residual}} = df_{\text{total}} - df_{\text{treatment}} - df_{\text{subject}}$.

$$df_{\text{residual}} = nk - 1 - (k-1) - (n-1) = nk - n - k + 1$$