GLOSSARY OF SOME STATS AND EXPERIMENTAL DESIGN TERMS

Here are some terms we've used that might start to blend together in your head, so maybe it'll help to keep them straight:

Complete vs. Partial Effect Size for eta² (η^2, or R^2) and ω^2

A complete effect size measure and a partial effect size measure both refer to the "proportion of variance accounted for" type of effect size measure, whether R^2 or ω^2. The difference is what variance they're trying to account for, i.e., what it's the proportion OF that they're trying to account for, or most simply, what the denominator is.
Using η^2 as an illustration:
Complete $\eta^2 = SS_A / SS_{TOT}$
Partial η^2 is the same idea but using a slightly different denominator: it's $SS_A / (SS_A + SS_{ERROR})$. With just one factor in the design, this is exactly the same thing as complete η^2 above, because $SS_{ERROR} = SS_{S/A}$ and $SS_{TOT} = SS_A + SS_{S/A}$ making the denominators the same.
But with more than one factor in the design, the effect being measured could be factor A, or B, or the interaction term A*B. For instance, looking at factor B as the factor of interest, partial η^2 would be $SS_B / (SS_B + SS_{S/AB})$, where $SS_{S/AB} = SS_{ERROR}$. The reason it's a better measure of effect size in that case is that we wouldn't expect factor B to explain ANY of the variation due to factor A or the interaction AB, so we might as well leave those two factors out of the denominator. In the example just mentioned, $SS_{TOT} = (SS_A + SS_B + SS_{AB} + SS_{S/AB})$, so leaving out the irrelevant parts of the denominator just leaves $(SS_B + SS_{S/AB})$.
A similar logic applies to ω^2, though the formula for estimating complete and partial ω^2 from the data isn't as transparent.
See K&W top of p. 235 for a good discussion of when it might be appropriate to use complete vs. partial effect size measures.

Subscripts ijk

Y_{ijk} is the score of the i-th subject in the j-th level of factor A and the k-th level of factor B (and so on to l and m and n etc. if there are factors C, D, and E etc.)

Between Subjects / CR (completely randomized) design vs. Within Subjects / RM (repeated measures) design vs. Mixed Design

(not to be confused with Between Groups and Within Groups variance, e.g. MS_A and $MS_{S/A}$)
Between Subjects factors have different subjects in each level, while Within Subjects factors have the same subjects getting more than one level, i.e., they're Repeated Measures. A design including only Between Subjects factors, so each subject contributes only one data point, is also called a Completely Randomized design. A design including only Within Subjects factors is also called a Repeated Measures design. It's possible to mix together one or more Between Subjects factors with one or more Within Subjects factors, and that's called a Mixed Design and is one of the most popular designs in psychology (two or more groups, each getting multiple treatments of a repeated factor).

Crossed Factor vs. Nested Factor
With Crossed factors, every level of one factor occurs in combination with every level of the other; with Nested factors, the levels of one factor occur only within one level of the other factor. The most obvious example is Subjects (you can think of them as a factor, with the levels being the various individuals), who are nested within levels of the independent variable(s) in a Between Subjects design (e.g., S/A, or S/AB, or S/ABC etc.) and only get one treatment or combination of treatments, while other treatments are given to different subjects. But other factors can be nested, especially when those factors are Random effects.

Fixed Effects vs. Random Effects (or Fixed vs. Random Factors)
Fixed factors have levels that were deliberately chosen, while Random factors have levels that are arbitrary or chosen randomly so they can be generalized to any other potential levels that could have been used -- the most obvious example of which is Subjects, since subjects are selected randomly so you can generalize to all other subjects you might have chosen. See the beginning of the Expected Mean Squares handout for some examples other than subjects (http://web9.uits.uconn.edu/lundquis/EMS.pdf).

Omnibus Analysis or Omnibus F
"Omnibus" is Latin "all" or "for all" and in English means "relating to or containing multiple items". The omnibus analysis or omnibus F is the overall ANOVA test of the null hypothesis that all group means are equal in the population. If that null hypothesis is rejected, more tests are called for to see which means differ from which other means. Even if the null hypothesis is not rejected (i.e., the omnibus F is not significant), some of the specific comparisons among the individual group means may be significant.

Comparisons vs. Contrasts
Comparisons and contrasts refer to the same thing in this context: A contrast is the difference between two means or sets of means (i.e., a linear combination of means with coefficients adding up to 0) that can be tested in order to compare those means or sets of means and see if they differ significantly (see http://web9.uits.uconn.edu/lundquis/ANALYTIC CONTRASTSf16.doc).

Planned Comparisons vs. Post Hoc Comparisons
Planned comparisons are those differences between means that were intended to be analyzed when the experiment was designed, corresponding to questions the research is intended to answer; they may be an orthogonal set, but don't have to be. They should be limited in number (see K&W pp. 115-117) and are generally not corrected for potential increased Type I error rates. Planned comparisons don't actually require that the omnibus analysis be done at all, let alone that it be significant (though sources differ on that).
Post hoc comparisons are those that would be interesting or useful to make based on the patterns observed in the data, thus are indicated "after the fact," i.e. after data is collected. They may be exhaustive or may be restricted to smaller "families" of comparisons, and

they require correction for increased Type I error rates since there are potentially very many that could have suggested themselves based on different ways the data might have turned out. Some comparison techniques require that the omnibus F be significant before proceeding (called "protected" tests, e.g., Fisher or Fisher-Hayter); others don't require a significant omnibus F to proceed (e.g., all the rest).

Analytic Contrasts: Coefficients, Single df Contrasts, SS Formula, Global Error Term, F Formula and df, Orthogonality

Analytic contrasts refer to contrasts set up to analyze the differences among specfic means or combinations of means and usually have the connotation of applying to planned comparisons rather than post hoc, though the procedure is completely general. (If used for post hoc tests a Type I error correction can be employed, such as the Bonferroni adjustment.) The procedure involves using a set of coefficients to make a linear combination of the means, contrasting one mean or set of means with another (thus using 1 df), calculating a SS and MS for the contrast, making an F ratio using the omnibus MS error term as the denominator along with that omnibus error term df, and checking that the set of such contrasts being tested are orthogonal if desirable and possible (see again http://web9.uits.uconn.edu/lundquis/ANALYTIC CONTRASTSf16.doc).

Simple Contrasts vs. Complex Contrasts

A simple contrast is the difference between the mean of one group and the mean of another group (e.g., the mean of group 1 vs. the mean of group 2). A complex comparison is the difference between a combination of two or more means (e.g., the mean of groups 3 and 4 combined) and another mean or combination of means (e.g., the mean of group 1 or the mean of groups 1 and 2 combined). In every case they refer to one mean vs. another so that the numerator df of the F ratio will be 1, and are therefore called "single-df" contrasts.

Orthogonal Comparisons or Contrasts

Orthogonal comparisons or contrasts are contrasts (single-df, as always) that are independent of one another, or put another way, are uncorrelated with one another -- i.e., they don't re-use any information, so the result of one contrast's test has no bearing on the result of another. For instance, the contrast "M1 vs. M2" is orthogonal to the contrast "M3 vs. mean of M1 and M2" because knowing the mean of M1 and M2 doesn't give any information about what M1 and M2 are or whether they're significantly different (e.g. if their mean is 7, are they 6 and 8? 4 and 10? -23 and 37?). For a given number of means being compared, there are "the number of means minus 1" possible orthogonal contrasts to make; no more than that can be mutually orthogonal. There are various choices of which set of orthogonal contrasts to make, but each set will have "number of means minus 1" comparisons in it. The set of orthogonal tests is usually considered to be dividing up the omnibus F's Type I error rate (α) and therefore doesn't require correction for increasing α due to multiple tests. This is not mathematically established but is a common convention. Since the "number of means minus 1" also describes the between-groups numerator df for an ANOVA (a-1), each contrast in an orthogonal set represents one of the between-groups

df, and in fact it's also true that the SS's for the orthogonal contrasts add up to the between-groups SS (SS_A) too.
The criteria for whether contrasts are orthogonal involve the contrasts' coefficients. The contrasts are orthogonal if the sum of the products of their respective coefficients for each mean is 0 -- that is, multiply the first contrast's coefficient for mean 1 by the second contrast's coefficient for mean 1, do that for the rest of the means, then add those products and they should sum to 0. This is less intuitive but more reliable than thinking through whether the questions are independent of one another.
If the set of planned contrasts that are theoretically interesting are orthogonal, asking separable independent questions, so much the better. But if not, the contrasts that are theoretically interesting should be the ones that get analyzed; orthogonality is neat, but not necessary. Usually the number of these non-orthogonal planned contrasts is limited to how many orthogonal contrasts there would have been -- "number of means minus 1" -- and though the lack of independence means the significance tests may be correlated (e.g., if one is significant, others may be more likely to be), the convention of not correcting for increased Type I errors is pretty well-recognized even without orthogonality.

Post Hoc Multiple Comparisons Type I Error Rates and Correction Techniques for α_{PC} , α_{FW} , and α_{EW}
PC: per comparison error rate - the α that would be used if only making one comparison, as in a t-test; typically .05
FW: family-wise error rate - limiting the Type I errors in a particular set ("family") of simple (or pairwise) comparisons to a family-wise error rate such as .05, or more liberally but still plausibly .10; Bonferroni (also called Dunn), Sidak, Dunnett, Tukey
FW protected: same as FW, but requiring a significant omnibus F first; Fisher, Fisher-Hayter
EW: experiment-wise error rate - limiting the Type I errors among any and all simple or complex comparisons that could be made, allowing for exploratory findings after data are collected, constrained by a very conservative significance criterion; Scheffé

Main Effects vs. Interaction Effects
General information about two-way interactions:
http://web9.uits.uconn.edu/lundquis/2way.doc
The "Main effect" of factor A is the difference in its levels' means when averaged over all the levels of factor B (and any other factors in the design), as if factor B was omitted and the analysis was done like a single factor (one-way) ANOVA on just factor A: "is the mean of all the scores from level a1 different from the mean of all the scores from level a2?".
"Interaction" means the effect of one factor depends on the level of the other factor. You probably know what interaction means but see K&W 201-202 and then the bottom of 209 for some right and wrong definitions, to sharpen your intuitions about it. A useful way of thinking about interaction is: A one-way ANOVA, or a test of a main effect in a factorial ANOVA, asks whether there is a difference between the means, e.g., "is there a difference between the means of conditions a1 and a2?" A test of an interaction in a factorial ANOVA

asks whether there is a difference between the differences between the means -- that is, "is the difference between the means of groups a1 and a2 different for condition b1 than for condition b2?" And though it becomes syntactically unwieldy, it is true that a test of a three-way interaction between factors A, B, and C asks whether there is a difference between the differences between the differences between the means (or, whether there's a difference between the AB interactions in c1 vs. c2)... and so on.
(General information about three-way interactions:
http://web9.uits.uconn.edu/lundquis/3way.doc)

Simple Effects and Interaction Contrasts: Techniques for Analyzing Interactions
The Simple effect of factor A "at b1" is the difference in means for the levels of A looking only at the subjects that got level 1 of factor B -- similarly there's a Simple effect of A "at b2" where you look at the A effect only for the subjects that got level 2 of factor B. See the 2way interactions handout linked above.
"Interaction contrasts" are 2x2 subsets of a larger interaction, for instance if factor A with 3 levels interacts significantly with factor B with 4 levels, and you want to test whether that interaction is due specifically to the effects at levels 1&2 of A and levels 2&3 of B. You've reduced it to a 2x2 interaction (with (a-1)(b-1) = 1 df, like any contrast) and that gives you a way of being precise about what's affecting what in the experiment.

See http://web9.uits.uconn.edu/lundquis/misc brief stats explications.htm for more miscellaneous terms and concepts (explained at a slightly more rudimentary level for the undergraduate research methods course).