STAT 430/830 Lecture Blog Spring 2012

Lecture Blog

I'll post brief summaries of the material we cover in lectures here.

Lecture 32 - July 25

Diana Katherine Skrzydlo - Aug 3, 2012 11:44 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today was a review class, and we went over chapters 6, 7, 8, 11, 12, 13, and 14.

Lecture 31 - July 13

Diana Katherine Skrzydlo - Aug 3, 2012 11:44 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today (the last lecture before project presentations!) we talked about split-plot designs. These are factorial designs (unlike nested designs), but with one of the factors not completely randomized due to practical reasons. Instead of comparing all the MS values to the MSE, we have two different error structures. If factor A is not completely randomized (it's called the whole plot) then we compare it to the interaction between A and the blocking/replication term. But anything that involves the other factors (called sub plot) we can compare to the MSE. If we want, we can be particular and compare MSB to MSB-block, MSAB to MSAB-block, etc but often we just assume block-treatment interactions within the subplots are negligible and group them together as MSE. We also did an example.

Lecture 30 - July 11

Diana Katherine Skrzydlo - Aug 3, 2012 11:40 PM

0 unread of 1 messages - 1 author(s)

Mark Unread
[Reply]
More actions...

Today we talked about nested designs, which are designs where the levels of one factor (say B) are contained within the levels of another factor (say A). Because of that, there is no meaning to the interaction term between A and B. Instead, our parameter estimates and ANOVA table include factor A effects and factor "B within A" (which we denote B(A)) effects. Luckily, it's easy to obtain the SSB(A) by just adding up SSB + SSAB if you had fit the model as if it were a factorial model. We did an example with machines each having two spindles.

Lecture 29 - July 9

Diana Katherine Skrzydlo - Aug 3, 2012 11:37 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we talked about the random effects model where there are two random factors, as well as the mixed model where there is one random factor and one fixed factor. I did an example (which turns out was slightly incorrect, I used the assumption that the tau's sum to 0, which they don't have to) of calculating the estimates of variance components, ICC, and a confidence interval for it.

Lectures 27 & 28 - July 6

Diana Katherine Skrzydlo - Jul 6, 2012 4:28 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we finished up chapter 12. We discussed what to do once we have our mean response and variance of response models (both of which are functions of the controllable factors **only**) - we can solve for the optimum analytically or if we have not very many variables, graphically with contour plots. We looked at an example (bottle filling, based on the data in chapter 6 question 18) and the results of that are posted.

Then we started chapter 13 - random effects. This is something that can be used in many of your projects without really changing your design at all, just how you think about the analysis. With a random factor, the levels you study are still specific levels, but we think of them as being randomly selected from a large population of potential levels. The parameters associated with the effect of each level on response (the τ_i's) are no longer unknown fixed parameters but random variables ~ $N(0, \sigma_\tau^2)$ which are independent of the ε's.

Because of this, instead of testing whether the τ_i's are all 0, we test whether σ_τ^2 is 0. To do this, we compare MSTreat to MSE (since under the null hypothesis both are estimates of the same thing). We can also get estimates of σ_τ^2 and σ^2 (the variance components) as well as the proportion of variability in y caused by the treatment factor, $\sigma_\tau^2/(\sigma^2+\sigma_\tau^2)$.

Lecture 26 - July 4

Diana Katherine Skrzydlo - Jul 6, 2012 4:15 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we started chapter 12 - robust parameter designs.

The goals of a robust parameter design problem are usually to get the best possible response (max, min, in a certain range, with low variance, or a combination) even though many factors that can be controlled in the lab will not be controllable in the real world when the product/process is used.

Robustness is not the same thing as optimality but we use many of the same techniques to find a robust solution. RPD problems will only arise when there is a significant interaction between at least one controllable and uncontrollable factor.

We briefly discussed the crossed array design procedure (which you will do in question 3abc on assignment 4) which is not the most efficient, and then talked in more detail about the combined array design method (which is 3d). We derived models for the expected response and variance of response in terms of the controllable factors (and the variance of the uncontrollable factors σ_z^2 as well as σ^2.)

Lecture 25 - June 29

Diana Katherine Skrzydlo - Jul 6, 2012 4:09 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we finished off chapter 11. We talked about multiple responses and what to do to optimize two or more components of a system at once.

We also discussed the types of models that are used in RSM in the two stages. For the initial model (to determine the direction to move in) we can use a 2^k or 2^{k-p} with or without centre points. Once we are near the optimum we want a more detailed model, so we can use a Cenral Composite Design (CCD) which includes factorial runs, centre point runs, and axial runs (runs where all factors are set to 0 except for one which is set to α). The choice of α is determined by the shape of the experimental region of interest. If it's a cube, $\alpha = 1$, and if it's a sphere, $\alpha = \sqrt{k}$ (not what I wrote, that was wrong).

Finally, we talked about blocking orthogonally in designs for RSM, which is important since we often will do some runs, then add more later (and the new runs can be considered a new block since they might be on a different day).

Lecture 24 - June 27

[Diana Katherine Skrzydlo](#) - Jun 29, 2012 11:14 AM

0 unread of 1 messages - 1 author(s)

[Mark Unread](#)
[[Reply]](#)
[More actions...](#)

Today we reviewed the first part of the algorithm for the method of steepest ascent. Then we talked about what happens when we actually get near the optimum.

We can tell we are there because a first order model will not fit well - there will be some curvature present in at least one of the factors. We want to find the maximum (or minimum) y with respect to the x's, so we fit a second order model and solve for the critical point by taking partial derivatives wrt $\underline{x}$ and setting to $\underline{0}$.

The solution (which we call the stationary point $\underline{x}_s$) may or may not be within the experimental region, and may or may not actually be a maximum/minimum for y (it could be a saddle point). So we can use canonical analysis to transform the variables into a more useful form that will tell us the nature of the response surface at the stationary point.

Lecture 23 - June 25

[Diana Katherine Skrzydlo](#) - Jun 29, 2012 11:06 AM

0 unread of 1 messages - 1 author(s)

[Mark Unread](#)
[[Reply]](#)
[More actions...](#)

Today we started chapter 11 - response surface methods.

We are aiming to find the optimal choices of the explanatory variables (x's) that give us the optimal response (y), where optimal could be maximum, minimum, or perhaps within a target range.

The technique we use is called the method of steepest ascent, and we discussed the first part of the algorithm today. We start by doing a small simple experiment at the current operating conditions, which may be quite far from optimal (so we don't want to waste too many runs there). We fit a first order model and use it to determine the direction of the optimum (the beta-hat vector gives us the direction to go, visually it's perpendicular to the lines of the contour plot) and we choose our step sizes based on process knowledge.

We then complete runs at each "step" along the path of steepest ascent and observe the response at each step. Eventually the response will stop getting better, and then we stop and conduct another simple experiment (with re-coded variables so the X matrix is still orthogonal) and repeat the process again.

Lecture 22 - June 22

Diana Katherine Skrzydlo - Jun 22, 2012 2:53 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we finished up chapter 8 by talking about saturated designs and full/partial foldovers.

A saturated resolution III design allows us to investigate 2^k-1 factors in only 2^k runs, as long as we don't care about two-way interactions. We start with a 2^k design matrix, and then alias each additional factor with one of the interaction columns. We looked at an example of investigating 7 factors in 8 runs, giving us a 2^{7-4}_{III} design.

If we want to investigate one factor (and its interactions) more closely, we can do a partial foldover, which is another 2^k runs with the factor of interest set to the opposite level as before. By combining the two experiments, we get an estimate of that factor along, and all its two-way interactions with other factors alone. The defining relation of the combined experiment omits all the words in which that factor appears.

If instead we just want better information about all the factors (and don't care about interactions), we can do a full foldover. Any aliasing with a even column is set to the opposite level, which means any words of odd length in the defining relation become negative.* Then, in the combined experiment, all words of odd length are now absent from the defining relation, so that means the resolution is now IV (since the shortest word is length 4.)

* someone asked why in the example we didn't make G = -ABC. If we did, the defining relation for the second experiment would be I = -ABD = -ACE = -BCF = -ABCG = ... = +CDG = ... = +BEG = ... = +AFG = ... Some of the words in the combined defining relation

of length 3 would still be around (since they are positive in both experiments), so the resolution would still be III.

Lecture 21 - June 20

Diana Katherine Skrzydlo - Jun 20, 2012 5:01 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we looked at the one quarter fraction, and also generalized to a $1/2^p$ FFD.

Our one quarter fraction of a 2^5 model requires us to alias two main effects with columns of the 2^3 design matrix, and we end up with a resolution III design, meaning many main effects are aliased with two-way interactions. The defining relation has 4 words in it (including I), and looking at the complete aliasing structure we see main&two-way and some two-way&two-way are aliased. Usually we ignore all 3-way and higher interactions so we can see the important aliasing more clearly.

We could also have chosen one of three complementary quarter fractions by setting one or both of our new factors' levels to the negative of the column chosen. Doing another quarter fraction means we can combine the results and get a clearer picture of some effects. Doing all three complementary quarter fractions would give us complete information with no aliasing.

In general, we can investigate k factors in 2^{k-p} runs with a $1/2^p$ FFD. The defining relation will have 2^p words and each effect will be aliased with 2^p-1 other effects! We want to choose the aliasing so we have the highest resolution possible, and also the fewest words of minimum length in the defining relation.

Lecture 20 - June 18

Diana Katherine Skrzydlo - Jun 18, 2012 3:21 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we looked at the complete aliasing structure of the 24-1 fractional factorial design. We notice that main effects are aliased with 3-way interactions and 2-way interactions are aliased with other 2-way interactions.

If we want more information than the principal one half fraction (the half fraction that includes the (1) run) provides, we can do the rest of the runs as their own FFD, called the complementary half fraction. The defining relation for that is I = -ABCD. Then we can combine the results from both experiments to obtain unaliased estimates of all the main effects and 2- and 3-way interactions.

We defined the resolution of a design, which is the length of the shortest "word" in the defining relation, and what resolution III, IV, and V designs mean in terms of which types of aliasing exist.

Finally, we looked at the projection property in a little more detail, and talked about confirmation runs.

Don't forget to choose a presentation date for your group project by following the instructions on the main page. Only one member per group can choose a day.

Lecture 19 - June 13 (my birthday!)

Diana Katherine Skrzydlo - Jun 14, 2012 2:51 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we finished up chapter 7. We discussed why when splitting the 2^3 design into 4 blocks, we chose the AB and AC columns to alias with our blocking factors, instead of using the ABC column and another 2-way interaction. (If we do that, we end up aliasing a main effect with a block effect, which is bad.)

We then started chapter 8 - fractional factorial designs. This uses many of the same techniques and principles as chapter 7, but instead of using those techniques to split runs across blocks, we use them to decide which subset of the possible treatments to actually perform - we don't do all treatment combinations in a FFD.

For example, if we have a 2^3 design (with 8 runs), instead of aliasing ABC with a blocking factor, we can set the levels of a fourth factor D to be equal to ABC. Then we can examine 4 factors in only 8 runs instead of 16. This is called a one half fraction because we do half as many runs as there are possible treatments. Graphically, we showed what those 8 runs look like, spread out across 2 cubes.

Lecture 18 - June 11

Diana Katherine Skrzydlo - Jun 11, 2012 2:33 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we talked about how we can use the group properties of the columns of the X matrix to determine what effects are aliased with each other in an experiment with blocking.

If we have chosen 2 blocks and aliased block = ABC, then the defining relation of the design is I = ABCblock. Then we can obtain the complete aliasing structure by multiplying through by every column of the X matrix. A is aliased with BCblock, B is aliased with ACblock, etc. In a design without replication, we have to use the graphical approach since we have no df for error. With replication, we can do partial confounding if we want.

If we would like to split the runs up into 4 blocks instead of just 2, we need to choose 2 columns of the X matrix to alias with 2 blocking factors. That leads to a more complicated aliasing structure, with each effect aliased with 3 other effects.

Lecture 17 - June 8

Diana Katherine Skrzydlo - Jun 8, 2012 2:33 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we started chapter 7 - introducing blocking to the 2^k factorial design.

We want a strategic way to divide our 2^k runs over 2 blocks that minimizes the confounding with effects we might be interested in. From the 2^2 and 2^3 examples, we derived the rule that we take the highest order interaction column and use that to select which runs go in each block.

The result is that the block effect is aliased with the highest order interaction effect (we cannot tell them apart), but usually that is OK since high interactions are negligible.

Finally, we looked at some interesting proerties about the columns of the X matrix in the 2^k design - they form a group where every element is its own inverse (sometimes called a Boolean group). This relationship will be useful in determining other aliased factors in models with blocking and also in chapter 8 when we do fractional factorial designs.

Lecture 16 - June 6

Diana Katherine Skrzydlo - Jun 6, 2012 2:47 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we talked about the second major disadvantage, inability to test for polynomial effects.

We can remedy this by adding centre points to our model, which are runs where all the factors are simultaneously set to 0 (as coded, which would be the average of the high and low levels in natual units). These do not change the estimates of the effect parameters, but we can use them to estimate the error and test for linearity.

We find the average response for the centre point runs and the factorial runs, and compute the SSPQ (SS pure quadratic) and compare it to the MSE (estimated from the centre point runs since the other df are used up estimating the parameters), and compare the result to an F distribution. If we fail to reject the hypothesis, it means the linear model is appropriate.

Lecture 15 - June 4

[Diana Katherine Skrzydlo](#) - Jun 6, 2012 2:41 PM
[Mark Unread](#)
[[Reply]](#)
[More actions...](#)

0 unread of 1 messages - 1 author(s)

Today we talked about one of the major disadvantages of the 2^k factorial model and how to get around it.

We cannot estimate all the interaction effects and the SSE in an unreplicated design. However, if the null hypothesis (that all parameters other than β_0 are 0) is true, then the estimates of the β_j's should all be independent and identically distributed normal random variables. So we can use a qq-plot to identify which, if any, differ from the qq-line.

We looked at the example with the springs, and identified the same effects as our earlier analysis as being significant.

Lectures 13 & 14 - June 1

[Diana Katherine Skrzydlo](#) - Jun 1, 2012 5:56 PM
[Mark Unread](#)
[[Reply]](#)
[More actions...](#)

0 unread of 1 messages - 1 author(s)

Today we started Chapter 6 by talking about the 2^k factorial model, its advantages and disadvantages. We code the low and high levels as -1 and +1 respectively (just a linear transformation of the natural variables), which makes the columns of X orthogonal and also means our regression model is consistent with the effects model.

We looked at an example of a 2^3 factorial model where we had 8 runs. We include all main effects and two-factor interactions in our model, but not the three-way interaction, because we would run out of df if we did.

The X^TX matrix just ends up being a multiple of the identity matrix, and our parameter estimates end up being **half** the difference between the average response with the factor at high and at low. Informally, this makes sense if you consider low to high being a 2-unit change in x, and the parameter is the change in y per unit change in x.

Formally, we can derive the variance of the parameter estimates (they all have the same variance) and do t-tests, F-tests, CIs, or build an ANOVA table. If we remove some terms from the model, the rest of the model (except for the SSE and the df associated with it) are unchanged, which means it's easy to construct a new ANOVA table for a reduced model. We looked at confidence intervals for mean response and also for the effect of a factor on the response.

Lecture 12 - May 25

Diana Katherine Skrzydlo - May 28, 2012 5:01 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we looked at an example in a lot of detail. The R code and data (and some of the analysis) is posted in Content.

We examined the data as a replicated design with one factor, as a randomized complete block design (which improved the residual analysis slightly), and as a two-factor design with replication. In the third case, we were able to examine the SS for polynomial effects both with the factors individually and the interaction terms. Interaction polynomial effects are rare, and in this case none were significant, but there was a linear z and a linear, quadratic, and cubic x in the functional relationship.

Throughout, we looked at residual analysis, which is examining for evidence to support our assumptions that the errors are normal, mean 0, constant variance, and independent. All the techniques are somewhat subjective, and if you want more techniques for residual analysis I'm sure there are some in the textbook or other courses.

Lecture 11 - May 23

0 unread of 1 messages - 1 author(s)

Diana Katherine Skrzydlo - May 28, 2012 4:56 PM
Mark Unread
[Reply]
More actions...

Today we talked about the a x b factorial model with replication. With n replications of every treatment, we can estimate interaction effects.

The interaction effect parameter $(\tau\beta)_{ij}$ represents the additional effect on response of being at level i of factor A and level j of factor B, beyond what the individual effects would contribute. We can test whether all interaction terms are zero simultaneously using an F-test, as well as test for factor A and factor B effects. Our ANOVA table performs all three tests.

We started talking about how we can apply polynomial contrasts to split the SS from each component (factor A, factor B, and AB interaction) into pieces to examine the functional relationship.

Lecture 10 - May 22

Diana Katherine Skrzydlo - May 22, 2012 9:35 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we finished off chapter 4 by talking about latin square designs, which allow us to block on two factors (each with p levels) and to examine one treatment factor in only p^2 runs. A related model is the graeco-latin square design, where we can actually have three blocking factors and one treatment factor in only p^2 runs.

Then we began chapter 5 by talking about the a x b factorial design with no replication. Luckily, the estimation of the parameters and construction of the anova table is identical to the randomized complete block design (with one treatment factor and one blocking factor) we looked at in chapter 4. We use an additive model (i.e. assuming the treatment factors have no combined effect on response other than their individual effects added together). Next time we'll add replication and interaction terms to the model.

Lecture 9 - May 18

0 unread of 1 messages - 1 author(s)

Diana Katherine Skrzydlo - May 18, 2012 4:28 PM
Mark Unread
[Reply]
More actions...

Today we started chapter 4 with a simple example of a randomized complete block design - the paired comparison experiment.

Then we generalized to a levels of the treatment factor and b levels of the blocking factor. We defined the matrices and vectors needed, found the parameter estimates, and derived the ANOVA table. Now we have components for treatment, block, and error (and the error component is smaller since the block term removes some of the unexplained variability), and we can get a better result when we test for the treatment factor's significance.

Lecture 8 - May 16

Diana Katherine Skrzydlo - May 18, 2012 4:26 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we talked about orthogonal contrasts.

One application of orthogonal contrasts is to test a model (if the explanatory variable is numerical) for polynomial effects in the relationship between x and y. We treat x as a categorical variable and fit the means model (or effects model), then use orthogonal contrasts to split the SSR into pieces based on the contribution of linear, quadratic, etc terms to the response. If we have a levels, we can fit up to a-1 polynomial effects. The coefficients to use if we have equally spaced levels were given. R can actually determine the coefficients for your contrasts automatically even if the levels are not equally spaced.

Lecture 7 - May 14

Diana Katherine Skrzydlo - May 18, 2012 4:22 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we continued the general model with a single factor with a levels.

We examined in detail the means model and the effects model, and found the parameter estimates for each. The estimates end up being independent because the columns of X are orthogonal.

We discussed the use of contrasts (a linear combination of treatment means where the coefficients add up to 0) to find out more information about how the null hypothesis is violated, if it is. We can get the estimate of a contrast, the variance of it, and the SS (sum of squares) of it, and with those quantities perform a t-test, f-test, or confidence interval.

Lectures 5 & 6 - May 11

Diana Katherine Skrzydlo - May 18, 2012 4:18 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we continued the general linear regression model by talking about confidence intervals, a simple comparative experiment (comparing the means of two groups), the paired comparison model.

We also began chapter 3 by looking at the analysis of variance (ANOVA) table and the f-test that goes along with it. We can also test several hypotheses at once using a nested f-test by fitting two models and using the results from each ANOVA table.

We looked at the difference between a confidence interval and a prediction interval for a specific set of x values.

Finally, we generalized the simple comparative experiment to a single factor with a levels (instead of just 2), and set up two possible models for this situation (the means model and the effects model).

Lecture 4 - May 9

Diana Katherine Skrzydlo - May 18, 2012 4:13 PM
Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we talked about the process for the general linear regression model, including setting up the model in vector/matrix form, solving for the least squares estimate of the parameters, deriving the distribution of the parameter estimates, and how to conduct hypothesis tests on a single parameter.

Lecture 3 - May 7

Diana Katherine Skrzydlo - May 9, 2012 11:04 AM

Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

Today we talked about the Data, Analysis, and Conclusions steps of PPDAC, including common things to avoid and things to keep in mind.

We did an example of a fishbone diagram which we use to identify all the factors that might have an effect on response. Ideally, for each factor identified, we must decide whether we will hold it constant, allow it to vary, use it as a design factor, or treat it as a nuisance factor (and either block for it or recognize it might have an effect we cannot measure)

Finally we discussed the 3 basic principles of experimental design - randomization, replication, and blocking. We will be seeing a lot of these as the course goes on!

Lecture 2 - May 4

Diana Katherine Skrzydlo - May 7, 2012 1:05 PM

Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

We discussed the possible goals of an experiment (which can sometimes conflict with each other), strategies for experimentation (best guess, one factor at a time, and factorial approaches), and overviewed the Problem and Plan steps of PPDAC including some definitions (treatment, run, and various kinds of factors)

Lecture 1 - May 2

Diana Katherine Skrzydlo - May 7, 2012 1:01 PM

Mark Unread
[Reply]
More actions...

0 unread of 1 messages - 1 author(s)

We went over the course syllabus.

We discussed the definition of an experiment and compared it to an observational study.
