RSM Split-Plot Design

Introduction

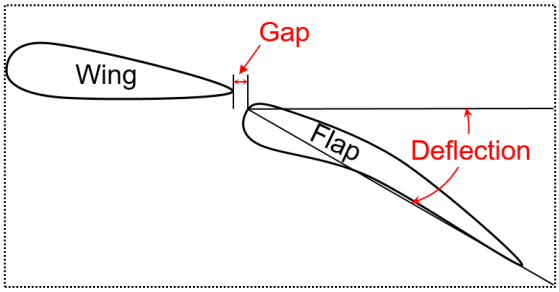
This tutorial illustrates Design-Expert® software tools for applying split-plot design to response surface method (RSM) experiments for process optimization.

In many experiments, the levels of one or more factors are hard to change. For example, heating an industrial oven from 300°F to 400°F may take hours. In this case, it will be far easier to group experimental runs by temperature rather than waiting for the oven to heat up and cool down all of the time. A design called a "split-plot" does just this. To get the proper p-values from a split-plot, specialized statistical tools such as restricted maximum likelihood (REML) must be performed. (More on REML later.)

Before getting started, note that throughout this tutorial we provide boxed-up details that are good to know but not essential for working through the program. Bypass these and press ahead if you are in a great hurry.

To illustrate how Design-Expert designs and analyzes a split plot for RSM let's follow a wind tunnel experiment on an aircraft wing—specifically the flap (Kowalski, Parker, and Vining, "Tutorial: Industrial Split-plot Experiments" example from the *Quality Engineering*, V1, #19, 2007).

If you've ever sat in the window seat of an airplane near the wing, the schematic below will look familiar to you.



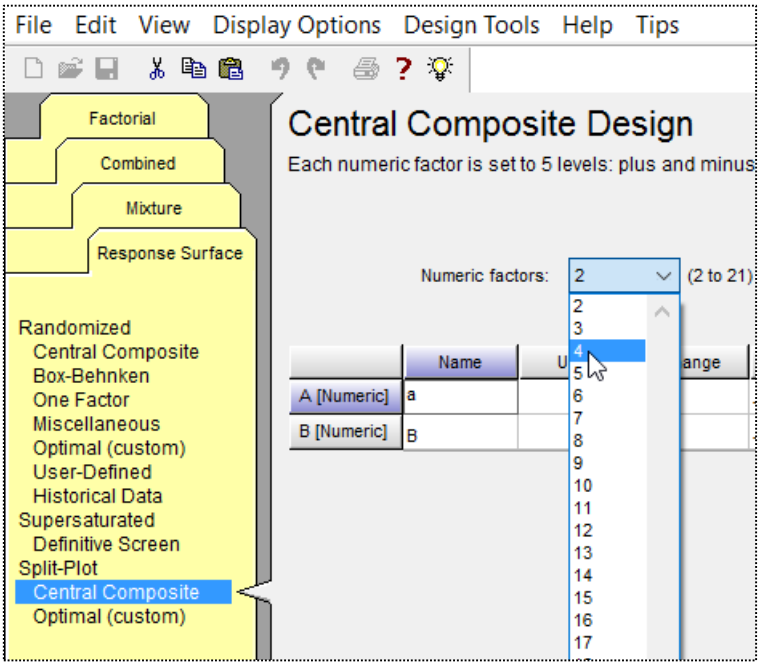
Wing and trailing flap with gap and deflection called out

The two parameters shown, gap and deflection angle, are hard-to-change (HTC) factors. Changing either variable requires opening up the environmentally controlled test chamber. This requires the process to be shut down, which is very costly both in terms of time and expense. A split-plot design minimizes the number of times these HTC factors must be changed, i.e., it restricts their randomization. Two other parameters to be tested are the angle-of-attack and the Reynold's number, both of which can be changed remotely from outside the chamber, making them easy-to-change (ETC) factors.

Design the Experiment

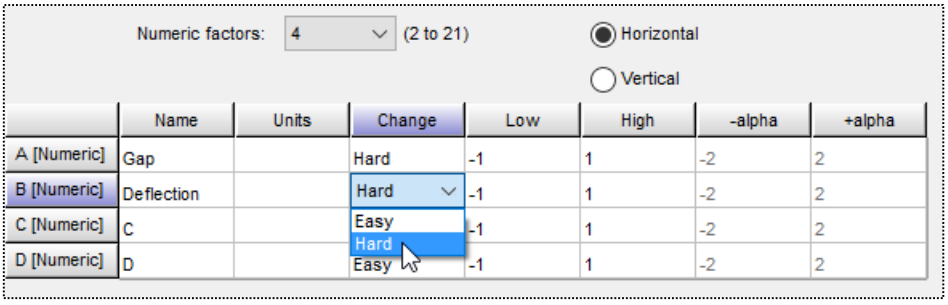
Click the blank-sheet icon 🗋 on your toolbar (or choose **File, New Design**). Then, click on the **Response Surface** tab and in the Split-Plot section near the bottom

choose **Central Composite** (CCD). Choose **4** as the number of factors from the droplist.



Setting up a four-factor Split-Plot CCD design

Enter **Gap** as your Factor a name. Notice that this factor is already labeled as HTC since a split-plot design was chosen. Click on Factor B and enter **Deflection** as the second factor. **Tab** to the **Change** column and type **h** to change this factor to **Hard**. You could also click on the droplist and change it, but typing is a nice keyboard shortcut for those who like them.



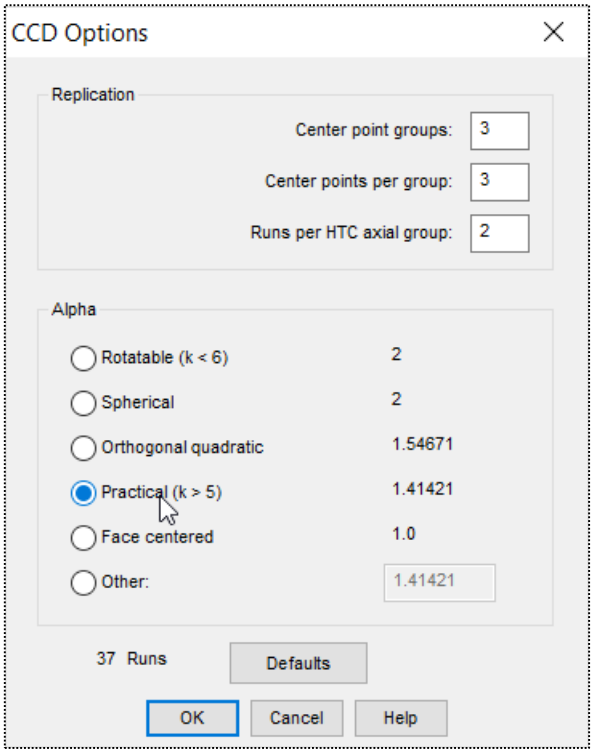
Changing Factor b to Hard- to-Change using the droplist

There is no need to enter the low and high levels of the factors because the authors of this example provided only the coded (-1 to +1) levels—not the actuals, perhaps because these needed to remain proprietary. Press ahead to Factor C and enter **Angle** as the name. Then, enter **Reynolds** as the Factor D name. Since these are ETC factors, leave the change column at the default of Easy.

	Name	Units	Change	Low	High	-alpha	+alpha
A [Numeric]	Gap		Hard	-1	1	-2	2
B [Numeric]	Deflection		Hard	-1	1	-2	2
C [Numeric]	Angle		Easy	-1	1	-2	2
D [Numeric]	Reynolds		Easy	-1	1	-2	2

All four factor names have been entered

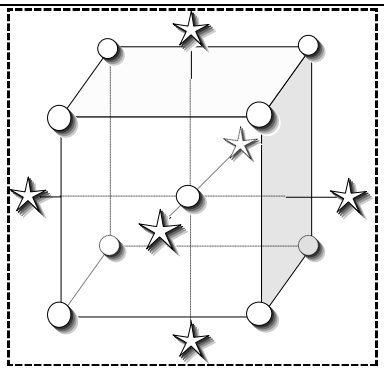
Click on the **Options** button near the bottom and middle of the screen and change the Alpha value to **Practical**.



Changing options to the Practical Alpha value.

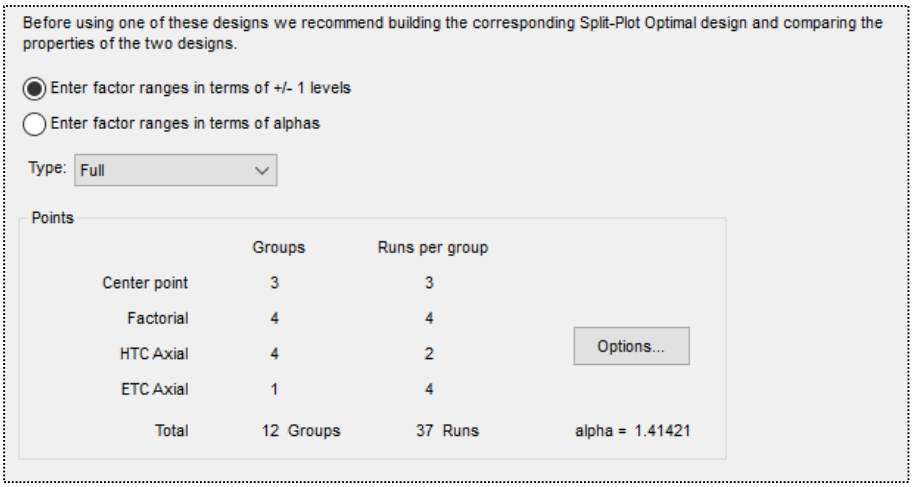
This value determines what levels the axial (star) points in the CCD will be run at. Alpha is the distance (in coded units) from the center of the design.

☒ ***Explore the makeup of a split-plot CCD***: The three-factor layout for this CCD is pictured below. It is composed of a core factorial that forms a cube with sides that are two coded units in length (from -1 to +1 as noted in the low and high levels above). The stars represent axial points. How far out from the center these points should go is a matter for much discussion between statisticians. They designate this distance "alpha" – measured in terms of coded factor levels. As you see, Design-Expert offers a variety of options for alpha. The default "rotatable" alpha gives the best distribution of the standard error (uncertainty) in your predictions, so it is desirable, though not completely necessary. The rotatable alpha requires the star points to be farther from the center (2.0 in coded units in this case), so it is not always feasible. If the levels for the star points are too far away from the center, the experiment may not produce measurable results. For this reason, the practical alpha can be used, which pulls the star points in a little bit to make sure the experiments are able to be run and produce meaningful results. Practical Alpha is recommended for more than 5 factors (k > 5), since the rotatable alpha becomes unwieldy when you get that many factors.



Central composite design for three factors

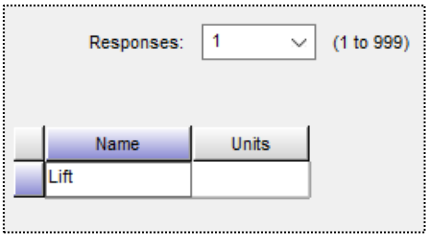
In addition to the alpha levels, in a split-plot Central Composite Design, there is also the complication of groups. At the bottom of the page, near the options button, you will see how this design was broken up into groups.



Breakdown of split-plot CCD into groups

For each type of point, it shows the number of groups and the number of runs per group. These can also be adjusted in the Options dialogue box, but we'll keep them at the defaults for this example.

Press **Continue**. For the R1 **Name**, enter **Lift** as the response.



Entering Lift as response name

Press **Finish** to complete the design-building wizard. You will get a warning to reset the factor levels between groups. Just click **OK** to bypass this warning, for now. Design-Expert (DX) now displays, in random run order grouped by Gap (Factor a) and Deflection (Factor b), the runs needed to complete the experiment. Your runs will most likely be different due to the randomizing (where not restricted) of the design. Also, note that in DX, the convention has been adopted that HTC factors will be denoted with lower case letters, as seen in the column

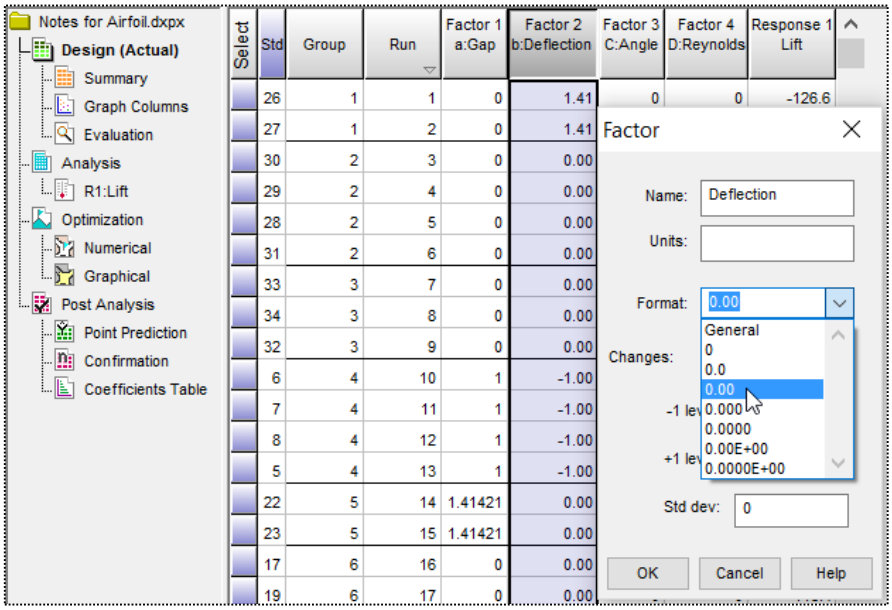
header. This sets them apart from the ETC factors, these being shown with upper case letters, as always in DX for fully randomized designs.

Select	Std	Group	Run	Factor 1 a:Gap	Factor 2 b:Deflection	Factor 3 C:Angle	Factor 4 D:Reynolds	Response 1 Lift
	26	1	1	0	1.41421	0	0	
	27	1	2	0	1.41421	0	0	
	30	2	3	0	0	0	-1.41421	
	29	2	4	0	0	1.41421	0	
	28	2	5	0	0	-1.41421	0	
	31	2	6	0	0	0	1.41421	
	33	3	7	0	0	0	0	
	34	3	8	0	0	0	0	
	32	3	9	0	0	0	0	
	6	4	10	1	-1	1	-1	

The first 10 runs in the design (yours will differ).

The Group column has been added to keep track of HTC factor combinations, which form the "whole plots". Note that the groups are separated by a horizontal black line. All of the runs in any given group maintain a constant level of the HTC factors. For example, in group 1, Gap is set to 0 and Deflection is set to 1.41421 for both runs. For all four of the runs in Group 2, the gap is set to the mid-level (0) and factor B is also at the mid-level (0).

To reduce the number of significant digits, **right-click** on the column header for Deflection and choose **Edit Info**. Click on the Format droplist and choose **0.00** and you will be left with 2 decimal points. Do this for all four of the factors to make the design cleaner to look at and run.



Changing the Factor settings to two decimal points results in the factors above.

The grouping of HTC factors makes the experiment much easier to run. As noted in the Introduction, the sealed experimental chamber would have to be opened every run to change the Gap and Deflection angle. With the split-plot design, this

disruption occurs only between the groups. Along the way, the aeronautic engineers sometimes go as many as four runs without having to take the time-consuming step of opening, resealing and, finally, re-pressurizing the chamber. One thing to remember, however, is that all the factors should be reset between groups, even if the HTC factor levels don't change. You may recall the warning that we clicked Yes on earlier, alluding to this. For instance, in groups 2 and 3 in this design, factor "a" and "b" stay constant at 0. To capture the variation that really occurs between groups, the experimenter should open the chamber and adjust these factors to some nominal value and back to 0 before running the third group. That way, the noise caused by resetting these levels is properly accounted for.

Analyze the Results

To avoid potential data-entry errors and ensure you get the same layout shown, simply read in the responses via the folder icon or by **File, Open Design** from the main menu. Select the file named **Airfoil.dxpx** and click **Open**.

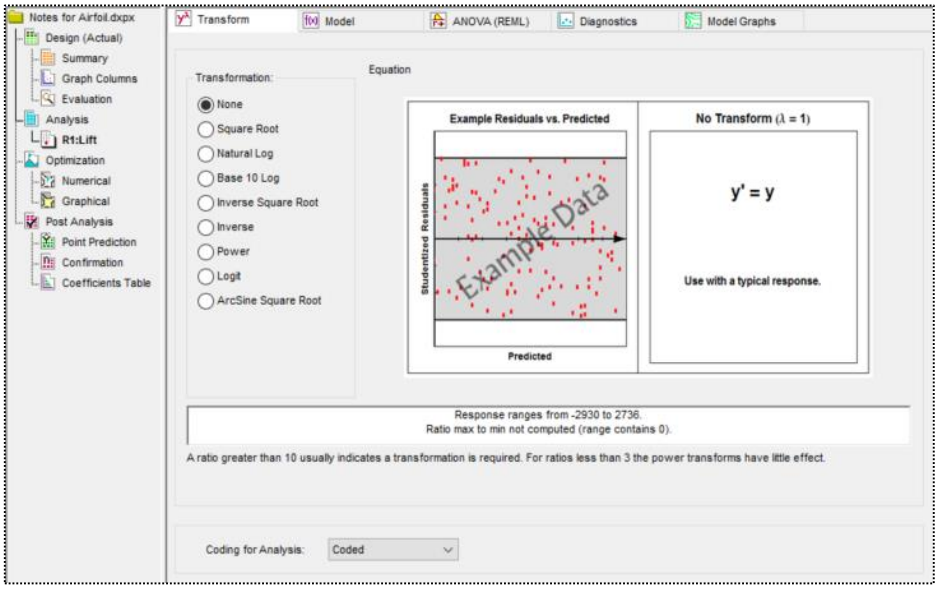
You should now see the data for the Lift Response as depicted in the screenshot. Only the first 22 of the 37 runs are shown, to save space.

Select	Std	Group	Run	Factor 1 a:Gap	Factor 2 b:Deflection	Factor 3 C:Angle	Factor 4 D:Reynolds	Response 1 Lift
	26	1	1	0.00	1.41	0.00	0.00	-126.6
	27	1	2	0.00	1.41	0.00	0.00	-245.5
	30	2	3	0.00	0.00	0.00	-1.41	-720.3
	29	2	4	0.00	0.00	1.41	0.00	2614
	28	2	5	0.00	0.00	-1.41	0.00	-2709
	31	2	6	0.00	0.00	0.00	1.41	199.3
	33	3	7	0.00	0.00	0.00	0.00	-168.5
	34	3	8	0.00	0.00	0.00	0.00	-36.61
	32	3	9	0.00	0.00	0.00	0.00	140.8
	6	4	10	1.00	-1.00	1.00	-1.00	1027
	7	4	11	1.00	-1.00	-1.00	1.00	-2392
	8	4	12	1.00	-1.00	1.00	1.00	1831
	5	4	13	1.00	-1.00	-1.00	-1.00	-2112
	22	5	14	1.41	0.00	0.00	0.00	-1058
	23	5	15	1.41	0.00	0.00	0.00	-1241
	17	6	16	0.00	0.00	0.00	0.00	78.67
	19	6	17	0.00	0.00	0.00	0.00	-118.4
	18	6	18	0.00	0.00	0.00	0.00	-38.03
	14	7	19	1.00	1.00	1.00	-1.00	589.2
	15	7	20	1.00	1.00	-1.00	1.00	-2737
	16	7	21	1.00	1.00	1.00	1.00	1233
	13	7	22	1.00	1.00	-1.00	-1.00	-2815

First 22 runs in the design with the response data entered

The response has been multiplied by a factor of 10,000 to convert the lift coefficient to integers for the most part. Positive values indicate an increase in lift relative to average and negative values indicate a decrease.

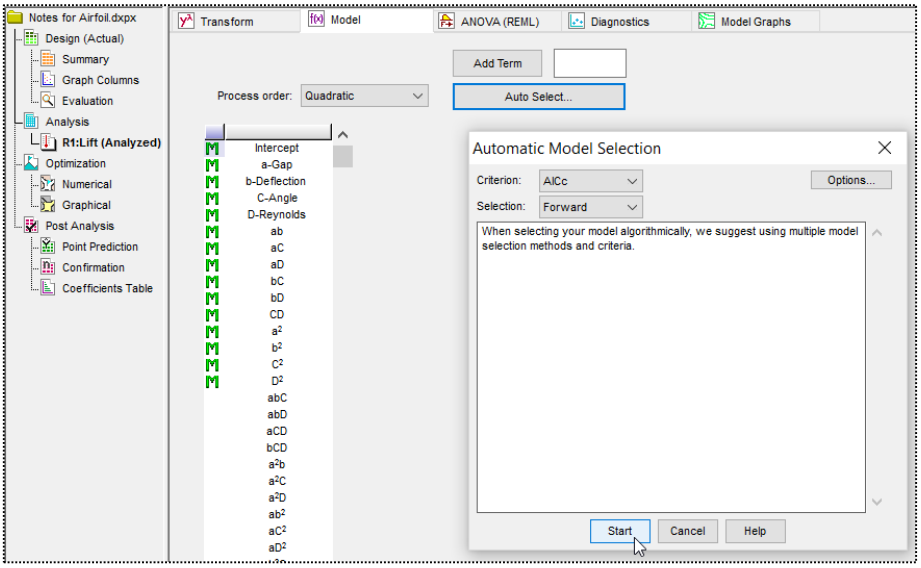
To get started with the analysis, click the node labeled **R1: Lift** under the **Analysis** branch. As with a normal RSM analysis, a new set of tabs appears at the top of your screen and they are arranged from left to right in the order needed to complete the analysis.



Begin analysis of Lift

There are a variety of Transforms that can be applied on this page. Not knowing if they will help at this point, click ahead to the ***Model*** tab. There are diagnostics that are checked later that can determine if a transform will help.

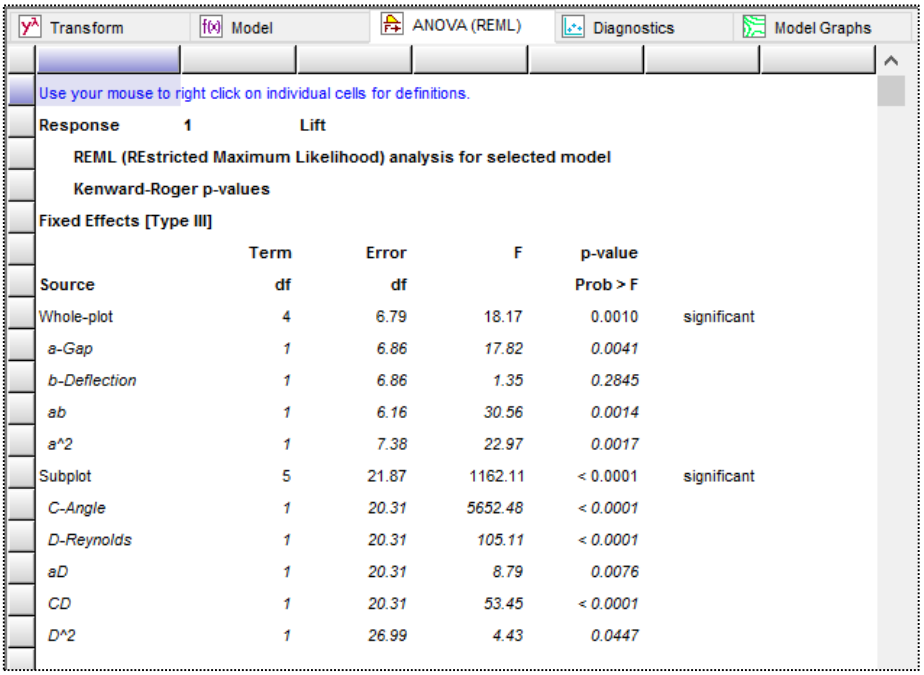
On the model tab, the quadratic model is presented for consideration (denoted by the green "**M**" by each term from the linear terms to the squared terms). Clicking ahead to the ANOVA (REML) screen at this point will evaluate the full quadratic model. However, it's best to do some analysis to select the best model from among the quadratic terms. To allow the computer to do this automatically, click on the **Auto Select…** button. Accept the defaults for AICc criterion and forward selection and click the **Start** button to run the analysis. DX will go through the terms in the quadratic model and select which ones improve the AICc criterion the most and add them to the model one at a time until adding terms will no longer improve the criterion.



Using Auto Select… to determine the best model

DX shows you the terms added in selecting the model, showing AICc criterion for each step. Click the Help button for more details on algorithmic model selection and the criterion used. Otherwise, click **Accept** to continue and evaluate the resulting model. You will get a warning that the model you have selected is not hierarchical. Be sure to click **Yes** to correct for hierarchy. This will give you a more statistically sound model, ensuring lower order terms (in this case b) are present to support higher order terms (like ab), even if they are insignificant. This is good statistical practice. Click on the help button in the warning box for more information.

The model statistics are shown on the ANOVA (REML) tab.



Source	Term df	Error df	F	p-value Prob > F	
Whole-plot	4	6.79	18.17	0.0010	significant
a-Gap	*1*	*6.86*	*17.82*	*0.0041*	
b-Deflection	*1*	*6.86*	*1.35*	*0.2845*	
ab	*1*	*6.16*	*30.56*	*0.0014*	
a^2	*1*	*7.38*	*22.97*	*0.0017*	
Subplot	5	21.87	1162.11	< 0.0001	significant
C-Angle	*1*	*20.31*	*5652.48*	*< 0.0001*	
D-Reynolds	*1*	*20.31*	*105.11*	*< 0.0001*	
aD	*1*	*20.31*	*8.79*	*0.0076*	
CD	*1*	*20.31*	*53.45*	*< 0.0001*	
D^2	*1*	*26.99*	*4.43*	*0.0447*	

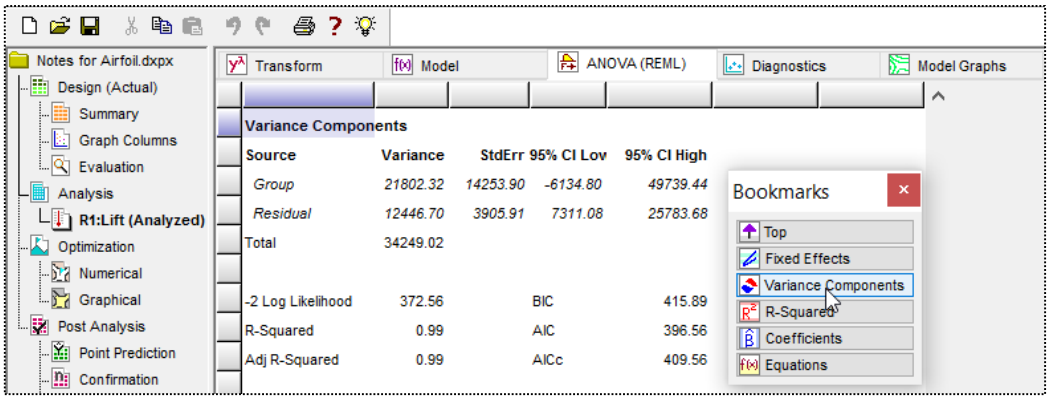
Model statistics from the REML analysis

This is not your standard ANOVA analysis, which relies on randomization for validity. The analysis done for split-plot designs in DX is a form of maximum likelihood estimation, more specifically, restricted maximum likelihood (REML), as noted at the top of the results table.

☒ ***Details on split-plot analysis***: The aim of maximum likelihood estimation is to find the parameter value(s) that makes the observed data most likely. Restricted maximum likelihood estimation, which is generally used unless you click on the Analysis menu available on the Model screen to change the method, is another way to estimate variances. In the split-plot case, REML estimates the Group variance for the whole plot factors and the residual variance for the subplot factors. Once the variances are estimated, Generalized Least Squares (GLS) is used to estimate the factor effects. The Kenward-Roger's method is then used to produce F-tests and the corresponding p-values. You can learn even more by clicking on the lightbulb icon for screen tips and following the links.

The big difference between the statistics on this table and a normal ANOVA is the grouping of variance terms into a Whole-plot section for the HTC factors and a subplot section for ETC factors. Look at the p-value (Prob>F) for the whole-plot. It is highly significant with a p-value much less than the generally acceptable 0.05 alpha level. That is for all the terms making up the whole-plot (HTC) portion of the model. You can also look at the p-values for the individual whole-plot terms, which look fine except for the *b-Deflection* factor, which has a p-value of 0.28. Remember, though, that this term was included for hierarchy to support the significant *ab* interaction.

The subplot terms as a whole are significant as are all the individual subplot terms. Next, press the **Variance Components** on the floating **Bookmarks** tool to see that DX presents various statistics to augment the REML analysis.

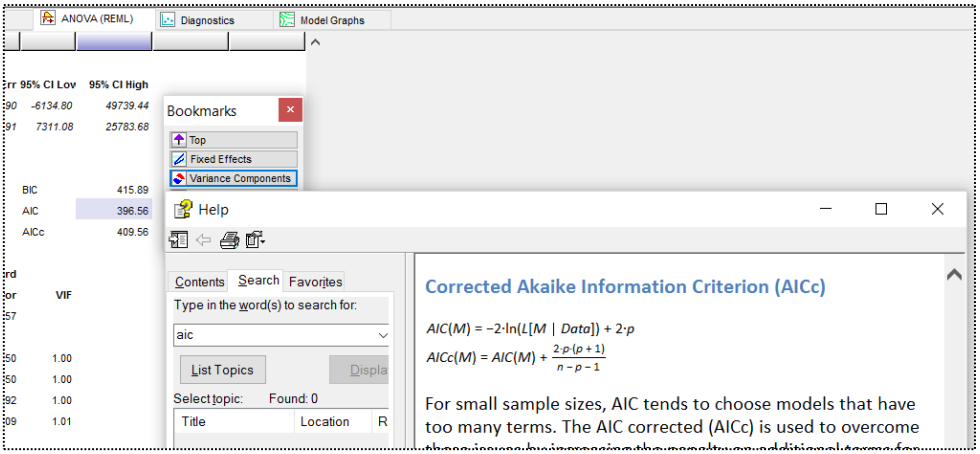


Variance Components and model criterion

Here, you will see more details on the variance components and likelihood ratios for the selected model, including the information criteria (AIC, BIC, and AICc) that score how good the model is. Just remember that when comparing models with these criteria, lower numbers are better, and you should only compare models from within one design that are nested to one another. For example, the model A+B can be compared to the model A+B+AB as all the terms in the larger model are also in the smaller model.

Another good statistic to study is the Adj. (adjusted) R-squared. This number goes from 0 to 1, with 1 being the best. In this case, the 0.99 Adj. R-Squared shows that we've captured most of the variation in the data (~99%) with the selected model.

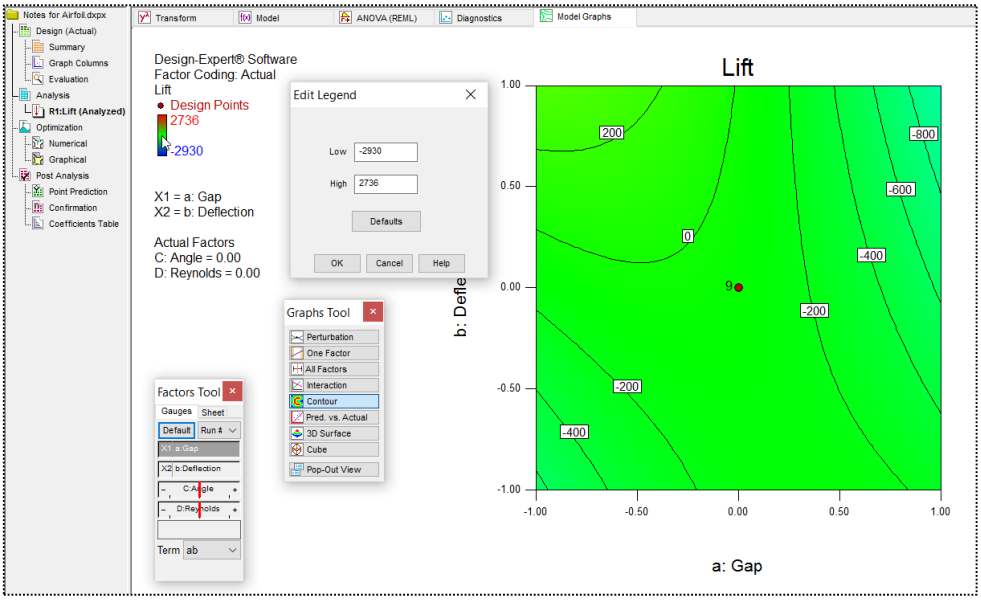
To learn more about the various model criterion, this would be a great place to exercise DX's context sensitive help. Just click on a number you are interested in to highlight it and then press the **F1** key (or right-click and select Help). For example, look at the information obtained about the AIC criterion and its counterpart AICc (you may have to scroll down in the help to see this).



Using the context sensitive Help (F1) to see more details

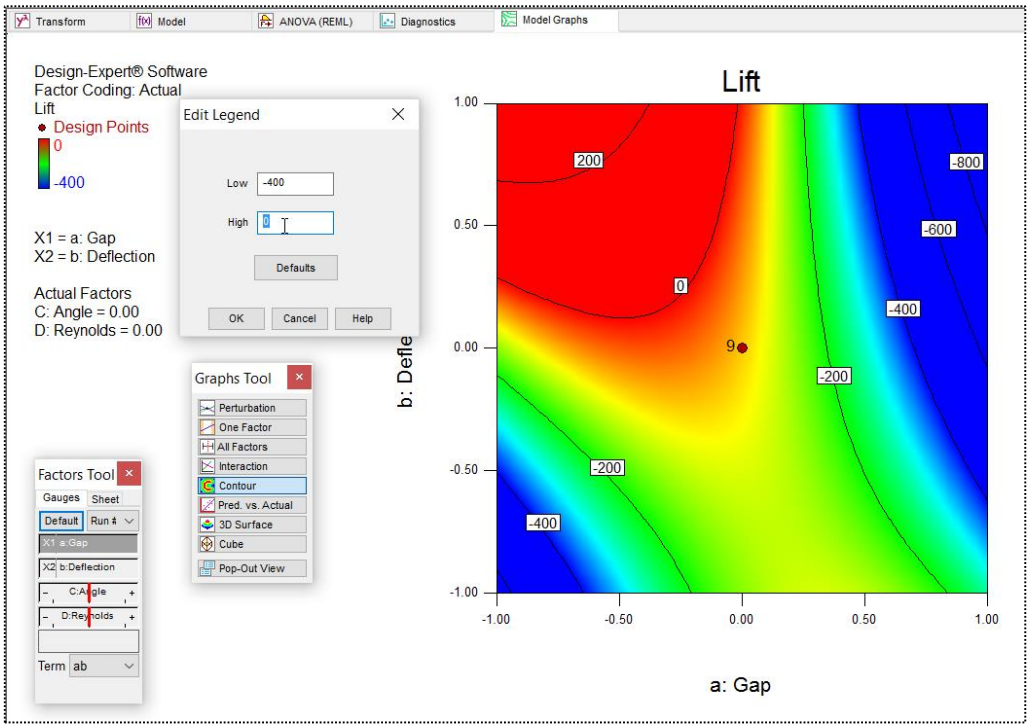
That's enough on the model statistics. It seems we have come up with quite a strong model. The next tab over is the "Diagnostics" tab, but those have been covered extensively in other tutorials. For instance, see the *Multi-Factor RSM* tutorial available on our tutorials page: http://www.statease.com/software/dx10-tut.html. The diagnostics used in a standard RSM design are calculated and can be applied here. They all pass with flying colors in this case so we'll move on. The focus of this tutorial is on the differences caused by this being a split-plot and the experimental results achieved with less effort

Jump ahead to the **Model Graphs** to explore graphical presentations of the predictions stemming from the selected model.



The initial graph is one color. Click to change the color scale via Edit Legend.

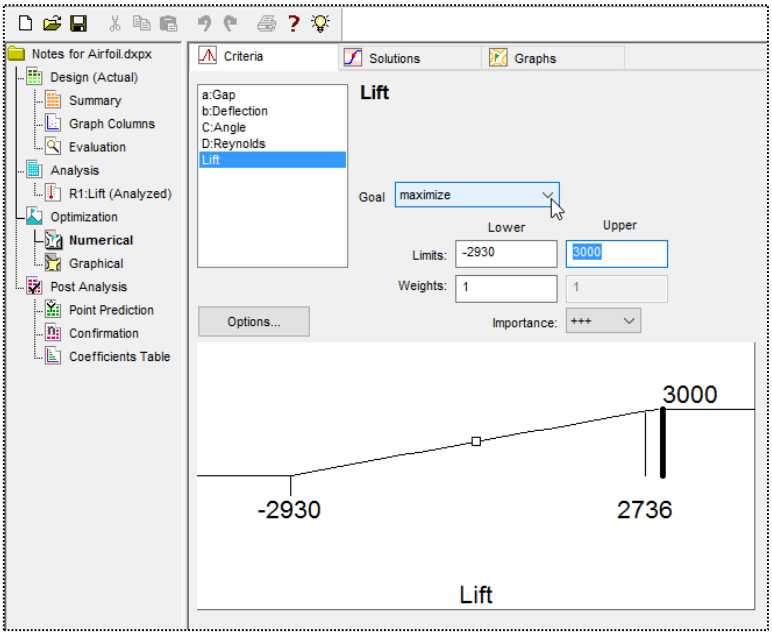
The first contour plot leaves a lot to be desired. The color scales being used are not highlighting anything, with the whole plot being green. Normally going "green" is a good thing, but some color variation will make this plot far more useful. Do this via a **Left-click on the color scale** in the legend to pull up the **Edit Legend** dialogue box. Change the **low** value to **-400**, which is one of the lowest contour values on the graph and the **high** value to **0**, the highest contour on the graph.



Changing the low value to -400 and high value to 0 gives a much better plot.

This gives a much more colorful, and thus, informative plot, which highlights the higher values of Lift achieved in the upper left corner at low values of Gap and high values of Deflection.

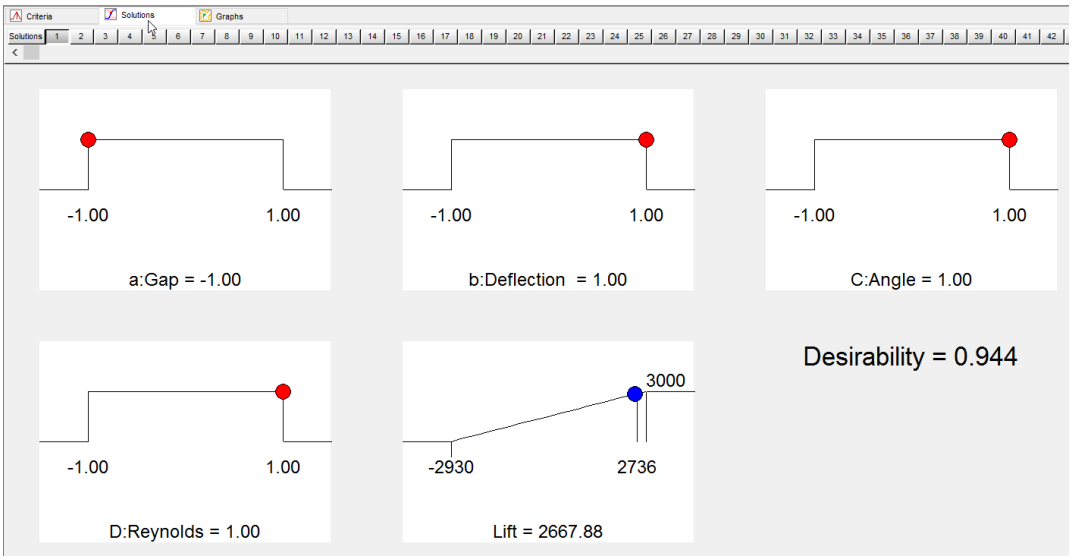
To see if these trends hold elsewhere and how Factor C (Angle) and Factor D-(Reynolds #) affect the results, use numerical optimization. Click on the **Numerical** node under the **Optimization** branch in DX. Click on **Lift** and set a goal to **Maximize** with an **Upper limit** of **3000**.



Setting up the Criteria for numerical optimization.

When working with quadratic models, it's good to put in a little stretch goal like this that exceeds the values achieved in the experiment. This ensures achieving the highest possible value.

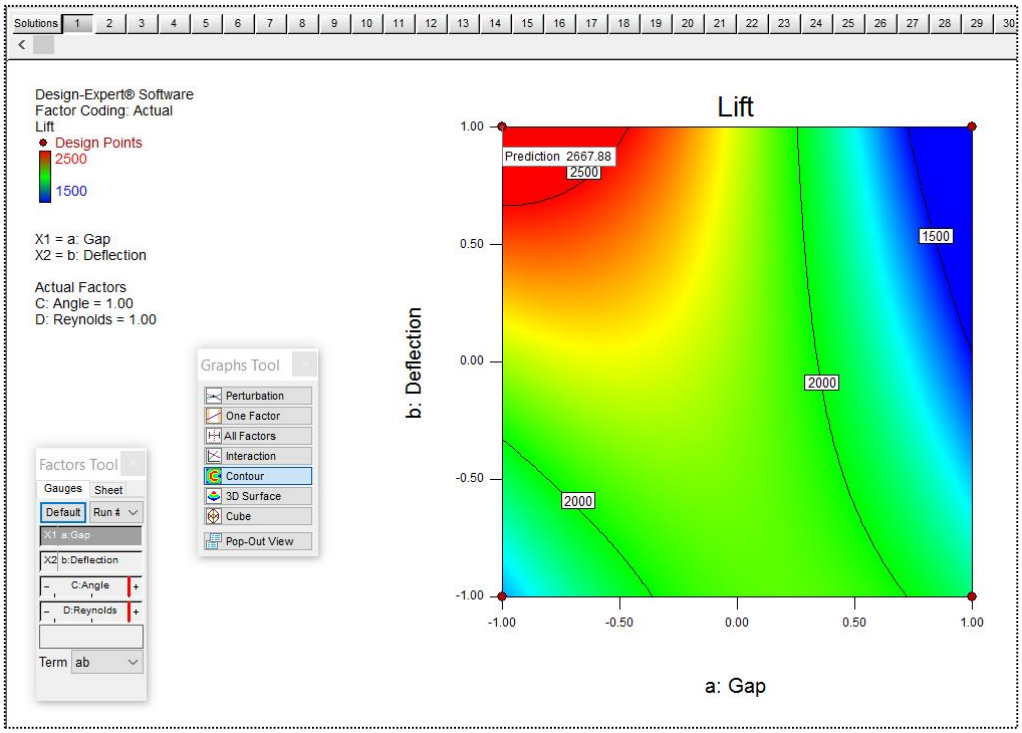
To see the optimal solutions, click on the **Solutions** tab.



The number one solution is automatically loaded in the ramps view

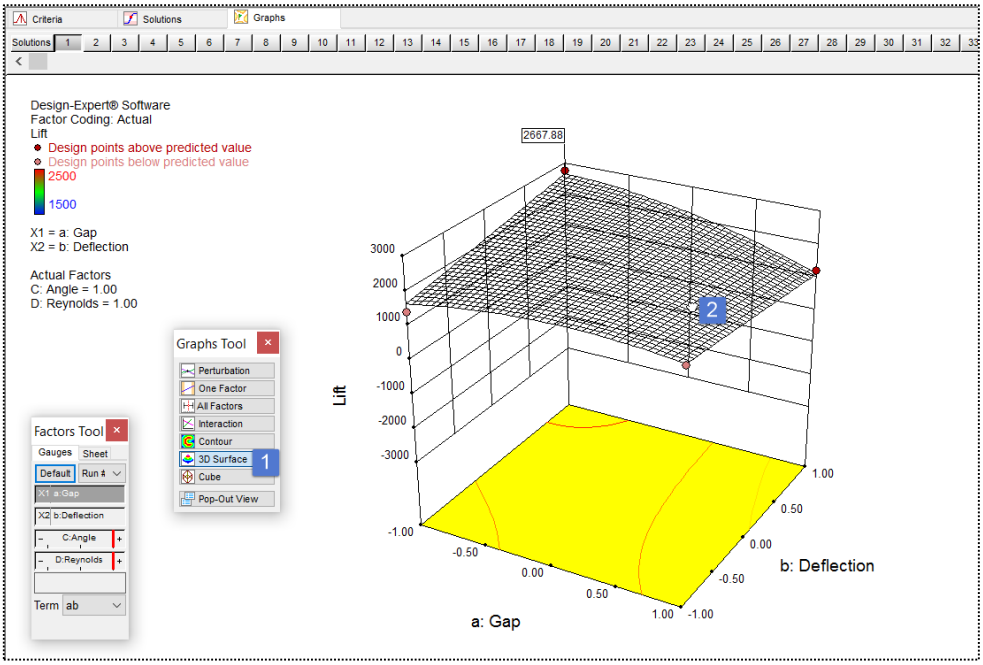
The best result occurs at a low level of gap (a) with high deflection (b), high Angle (C), and high Reynolds (D). This agrees with what we saw on the contour plot between gap (a) and deflection (b). With these settings, a Lift of 2667.9 is achieved. All of the other solutions, which can be explored by clicking on the other numbers across the top, give less lift.

To see what the contour plot looks like at this recommended number 1 solution, click on the **Graphs** tab. Note that solution 1 button is automatically depressed in the upper solutions bar. By default, All Responses is selected on the Response droplist. This shows the Desirability plot, which is used to find the optimum, alongside the Lift plot. To focus on the Lift plot, click on the Response droplist and select **Lift**. Now, there is again little stratification of colors on the plot, so left-click on the color scale and change the low to 1500 and high to 2500. This results in a better plot, highlighting the region where the high lift optimum is achieved in red. DX also plants a flag at the optimal solution (the upper left corner) found by the numerical optimization.



Optimal solution is flagged on the contour plot and highlighted in red region after changing the color scaling (1500-2500)

For a better view of the response surface, click on the **3D Surface** plot from the floating Graphs tool (1) and then click and drag the graph to a new orientation (2) with the optimum solution in the back to better highlight the slope of the surface.



Examining the optimum in 3D. Click on the 3D surface and drag the plot to rotate.

Summary and Conclusion

In this experiment, there were HTC factors (gap and deflection) that required the costly step of opening the environmental chamber. By using a split-plot design, the number of times the chamber has to be opened is reduced by sorting the design into whole-plot groups. Since randomization was restricted by sorting, a REML analysis was then applied to get the proper p-values for the model. Numerical optimization then resulted in the conditions that will optimize the Lift of an airplane wing. Using a split-plot design allowed the experimenters to save time and money and run a design which might well have been infeasible if fully randomized. If you ever find yourself tempted to sort your randomized design to make it easier to run, try a split-plot design instead. This will give you the proper p-values and do the sorting for you.

Split-plot designs are the epitome of efficiency when it comes to the experimental effort needed. Just be aware that the easier experimenting comes at a cost, less precision and power in determining the HTC factor's effects. Be sure to use power or fraction of design space (FDS) sizing tools to make sure you have enough experimental runs to overcome the loss in power and precision that come with restrictions in randomization. These tools are discussed in other tutorials and in webinars you can find here: http://www.statease.com/training/webinar.html. Then, just follow the steps outlined above to design and analyze your experiment.