Top N Pareto Front Search Add-In (Version 1) Guide

Authored by:
Sarah Burke, Ph.D., STAT COE
Christine Anderson-Cook, Ph.D., Los Alamos National Laboratory
Lu Lu, Ph.D., University of South Florida

27 February 2017

Revised 10 October 2018

The goal of the STAT COE is to assist in developing rigorous, defensible test strategies to more effectively quantify and characterize system performance and provide information that reduces risk. This and other COE products are available at www.AFIT.edu/STAT.

Approved for Public Release LAUR-17-20624

Table of Contents

Executive Summary 3
Objectives 3
Layered Pareto Front (PF) 4
Example Application 5
Launch Window 6
TopN-PFS Add-In Options and Settings 7
Phase 1 Options 8
Phase 2 Options 8
Scaling 8
Desirability Functions (DF) 9
Phase 1 Output 10
Data Table: Pareto Front Results All Data 10
Scatterplot Matrix 11
Phase 2 Output 12
Mixture Plot 12
Proportion Plot 12
Parallel Plot 13
Synthesized Efficiency Plot 14
N Comparison Plot 14
Additional Output Tables 15
Miscellaneous Notes 15
References 16

Revision 1, 10 Oct 2018: Formatting and minor typographical/grammatical edits.

Executive Summary

This guide is for the Top N Pareto Front Search (TopN-PFS) Add-In (Version 1) for JMP. We first present the overall Objectives of this Add-In. We then describe each of the settings currently available through a short example application to prioritize stockpiles based on several criteria.

Direct any comments, questions, or issues to Sarah Burke at sarah.burke@theperducogroup.com. This Add-In can be obtained by emailing the STAT COE at CoE@afit.edu. The Add-In is also available on the Science of Test website (http://www.testscience.org/wordpress/resources/tools/).

Keywords: Pareto front, Define-Measure-Reduce-Combine-Select, reliability, decision-making, multiple criteria

Objectives

In a decision-making process, relying on only one objective or criterion can lead to oversimplified sub-optimal decisions which ignore important considerations. Incorporating multiple, and likely competing, objectives is critical during the decision-making process in order to balance the tradeoffs of all potential solutions. There are methods available (Lu et al., 2011) to identify the best or optimal solution for a specified set of priorities. There are many situations, however, where we are interested in several optimal solutions, not just one. These types of decisions may fall into one of two scenarios: 1) decision-makers want to identify the best N solutions to accomplish a goal or specific task, or 2) a decision is evaluated based on several primary, quantitative objectives along with secondary, qualitative priorities. In this second scenario, we identify several contending solutions (and eliminate non-contenders) using the primary, quantitative objectives. We then use the secondary, qualitative objectives to make the final decision.

The application presented here involves a complex decision-making problem for budget allocation among stockpile programs to enhance stockpile performance. Each year, choices are made about which stockpiles should receive additional funding to ensure the units are available and reliable for use. In this case, only 4 stockpiles will receive funds due to limited resources.

This Add-In supports the structured decision-making strategy proposed in Anderson-Cook and Lu (2015): Define-Measure-Reduce-Combine-Select (DMRCS).

The Define step focuses on selecting the best characteristics over which to optimize. Three criteria were selected in this problem: measures of reliability, urgency, and consequence. These three aspects of each stockpile represent different dimensions of the decision and are important to consider.

The Measure step focuses on the importance of having appropriate data on which to base the decision and to ensure that all the important factors of the decision are represented with appropriate metrics. Each of the three criteria, reliability, urgency, and consequence, were measured on a 0-10 scale where 0 = least critical and 10 = most critical. These measures were determined for each stockpile using historical data and multiple expert opinions.

The Reduce step simplifies the set of choices from which to choose in two ways. First, it prioritizes the criteria of the decision to focus on a small, manageable number of dimensions. Second, it reduces the possible number of solutions by eliminating non-contending choices from further consideration. A Layered Pareto Front (PF) is an objective and efficient way to achieve this goal, and is discussed below.

The Combine step considers how to examine the trade-offs between different criteria. Scaling the criteria and then using a desirability function (DF) to combine each criterion into an overall score allows each solution to be ranked based on different priorities on each criterion. Flexible choices of weights are considered to identify robust solutions across different prioritizations. Go to the Scaling and Desirability Functions sections for more details.

Finally, in the Select step, the decision-makers identify the top choices best suited to their priorities and the Add-in provides several graphical tools for making comparisons between close contenders, discussed in Phase 2 Output.

Layered Pareto Front (PF)

A Pareto front (PF) (Lu et al., 2011) is a tool that identifies multiple dominant solutions that are best for all of the possible weights for combining the criteria that could be chosen. One choice dominates another if it has all criteria with at least as good values as any other solution and at least one strictly better value among the criteria. The PF contains all dominant solutions with maximum overall desirability function values for the additive and multiplicative desirability functions (refer to the section on Desirability Functions). Identification of the PF does not require any input on the preference of one criterion over another. Burke et al. (2016) extend this approach to a layered PF approach for the top N solutions. The layered PF considers potential solutions that lie just behind the PF as additional potential candidates. While these solutions are not the top choice for any criteria weighting, they could be highly competitive choices for particular regions of weights on the criteria.

The layered PF divides the sets of all choices into multiple layers with ranked solutions. The top layer PF has choices that are strictly better than the second layer PF and so on. Because the top N solutions for any choice of weights must necessarily be included in the top N layers of PFs (Burke et al., 2016), the top N layers of PFs identify an objective set of superior choices prior to considering subjective weighting choices.

An example of a layered PF with five layers is shown in **Figure 1**. Both criteria in this example are optimized by maximizing; therefore, solutions in the upper right corner are most desirable. The triangle in **Figure 1** represents the utopia point, the ideal solution which has the best possible values for all criteria simultaneously. Note that this ideal solution rarely exists in practice, since there are generally tradeoffs among the criteria.

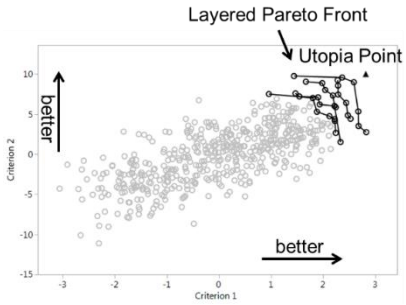

Figure 1. Layered PF with 5 layers where both criteria are optimized by maximization

Example Application

We consider a group of 42 stockpiles which come from four families (A, B, C, and D) in **Figure 2**. Because of the proprietary nature of the data, we present a dataset similar in characteristics to the original data. The budget only allows four stockpiles to receive additional funding, so the goal is to identify the top four stockpiles that have the most dire needs based on three criteria: overall reliability, overall urgency, and consequence (OR, OU, C). Each of these criteria are on a scale from 0 to 10 (0 = least critical, 10 = most critical). Note that the dataset must have a column with an ID variable so that the solutions can be compared in Phase 2 of the algorithm. The ID column could be character or numeric as shown in the first two columns in **Figure 2**.

Stockpile	Stockpile Number	Overall Reliability	Overall Urgency	Consequence
A1	1	9	6	2.5
A2	2	6.5	8.75	5.5
A3	3	9.5	8.75	9
A4	4	9	8.75	9
A5	5	7	8.5	4.5
A6	6	7.25	8	3
A7	7	8	7.75	3.5
A8	8	8.5	7.25	3.5
A9	9	7	6.5	3
A10	10	7.75	8.75	5
A11	11	8.5	7	3.5
A12	12	6.5	6.75	8.5
A13	13	7.75	7	5.5
A14	14	8.5	8	4
A15	15	7.25	8.25	3
B1	16	7.25	7.25	7
B2	17	8.25	6.25	7
B3	18	6.5	6.25	8
B4	19	8.75	9	9

Figure 2. Screenshot of stockpile data table

Launch Window

Figure 3 displays the launch window for the TopN-PFS Add-In. The columns representing the criteria of interest are selected and placed in the 'criteria' window. At least two columns are required for the criteria, which must be numeric. The ID column should be selected and placed in the 'ID' window. **Figure 4** shows the launch window populated for the stockpile example.

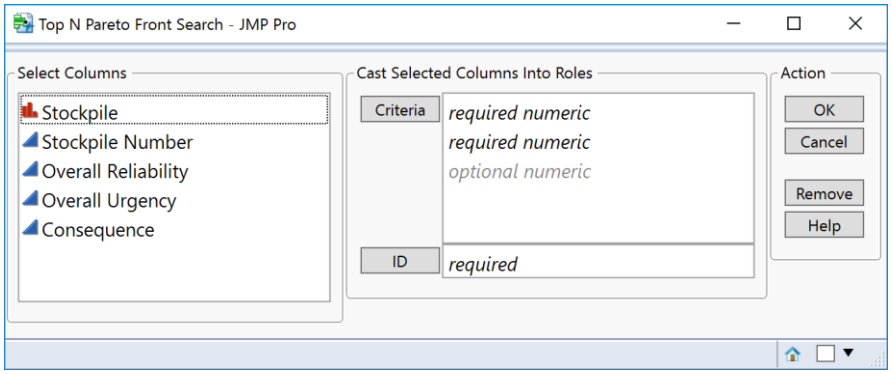

Figure 3. Launch Window of TopN-PFS Add-In

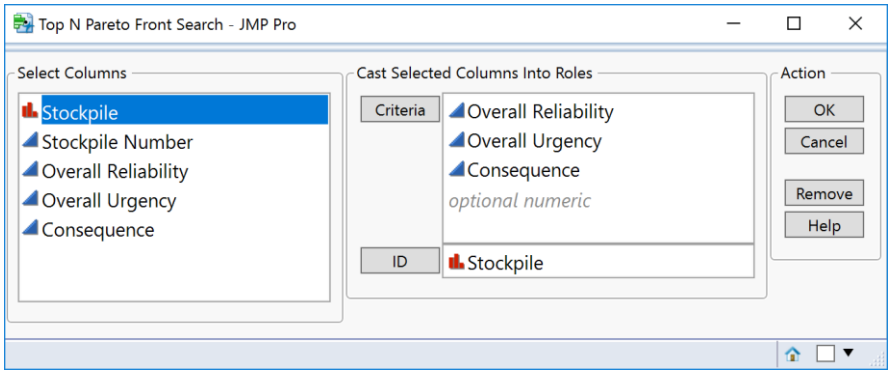

Figure 4. Launch Window of TopN-PFS Add-In (Filled in)

TopN-PFS Add-In Options and Settings

After selecting OK in the launch window, a window with several options opens (**Figure 5**). The TopN-PFS algorithm is broken into two phases. Phase 1 is the objective phase where the layered PFs are identified. The Add-In allows you to perform Phase 1 Options only (the Reduce step) if the goal is just to identify the solutions in the layered PFs.

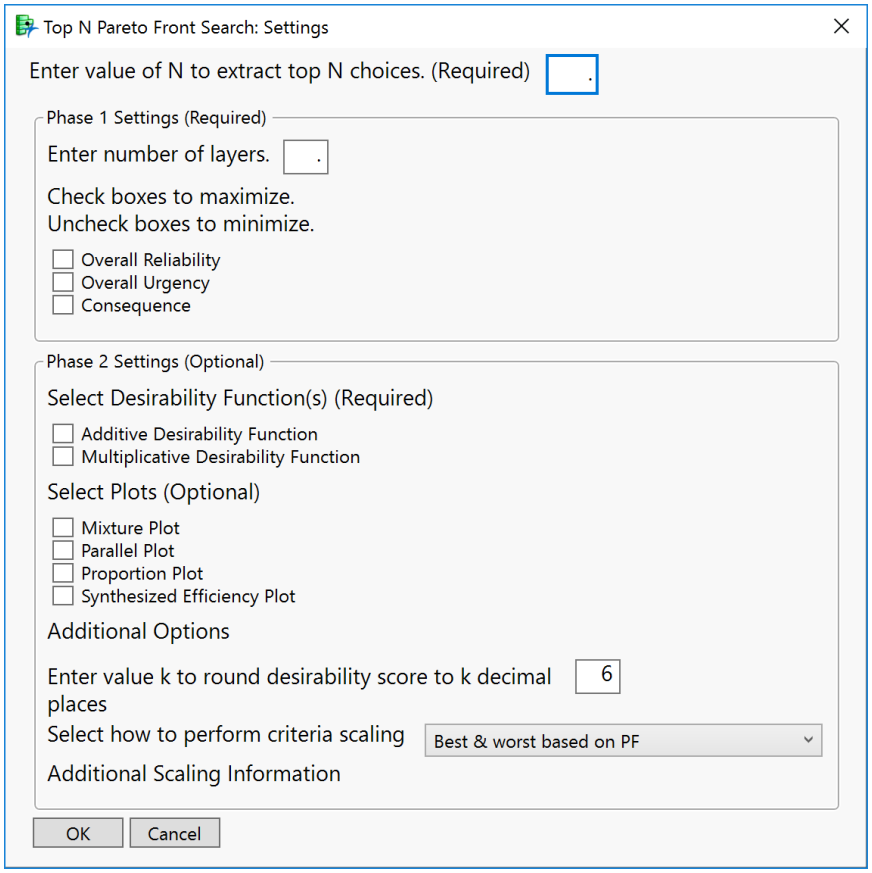

Figure 5. TopN-PFS Settings Window

Phase 1 Options

- Enter the value of N, the number of solutions you would like to identify.
- Enter the value of m, the number of PF layers. Usually, this value equals N (refer to Burke et al., 2016 for details).
- Check the box next to the criterion if it is optimized by maximizing; if it is optimized by minimizing, leave the box unchecked. If it optimized to a target value, prior to using the Add-In, use the formula option within a JMP data table to create a column to quantify the distance from the target in the data table and then minimize the new column.
- To get the Phase 1 analysis only, click OK and analyze the output.

Phase 2 Options

Phase 2 is the subjective phase that allows the user to incorporate their priorities and analyze the tradeoffs between the criteria (the Combine and Select steps of DMRCS). The criteria are scaled to values between 0 and 1. The scaled criteria are then combined into one desirability function score (Derringer and Suich, 1980) for different combinations of weights on the criteria.

- Select which desirability function to use. You must select at least one of these (if you select both, results for each DF will be in separate windows). See Desirability Functions section for details.
- Select which plots to create. The Phase 2 Output section explains all graphical options in detail.
- Select which criteria scaling matches your needs by selecting one of the four options (see Scaling section for details).
- If you would like to round the desirability function scores to k decimal places, enter in that value (k >= 0). The default value is to round to 6 decimal places.
- Click OK and analyze both Phase 1 and Phase 2 output.

Scaling

In the Combine step, the original criteria values x_{ij}, the jth value for the ith criterion, are scaled to z_{ij} such that:

$$z_{ij} = \frac{x_{ij} - worst_i}{best_i - worst_i}, \text{if criterion i is maximized}$$

$$z_{ij} = \frac{worst_i - x_{ij}}{worst_i - best_i}, \text{if criterion i is minimized}$$

There are four options available for scaling:

- "Best & worst based on PF" – the best and worst values are defined from the points from Phase 1 of the analysis (only the points in all of the layered PFs).
- "Best & worst based on all data" – the best and worst values are defined from all the points in the original data set.
- "Best based on PF, user-defined worst values" – the best values for each criterion are defined from the points from Phase 1 of the analysis (only the points in the layered PFs); the user defines the worst values for each criterion in the settings window once this option is selected.

- "User-defined best & worst values" – the best and worst values are defined by the user by entering in values in the settings window once this option is selected. If there are values outside of this specified range, these values are scaled to 0 and 1.

A table to specify best and/or worst values will appear when the third or fourth scaling options are selected from the drop-down menu. Figure 6 and Figure 7 show how to select among the scaling options and where to enter in best and/or worst values for the third and fourth scaling options.

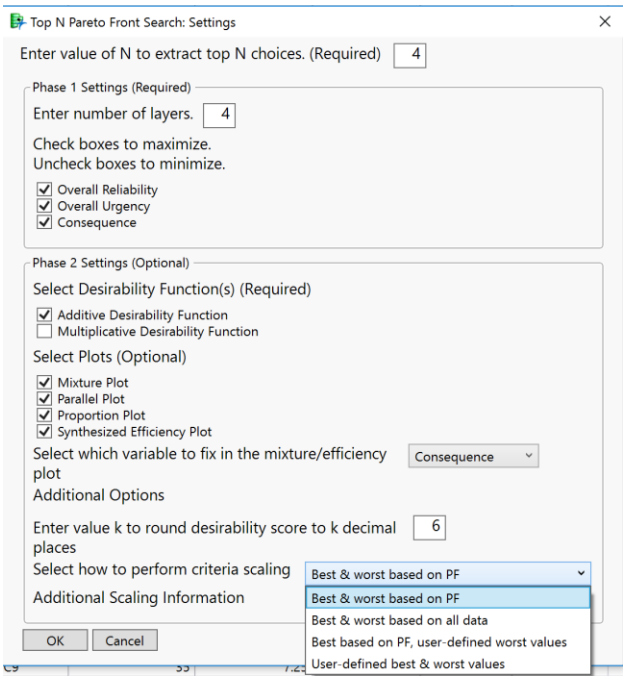

Figure 6. TopN-PFS Settings window: Scaling Options

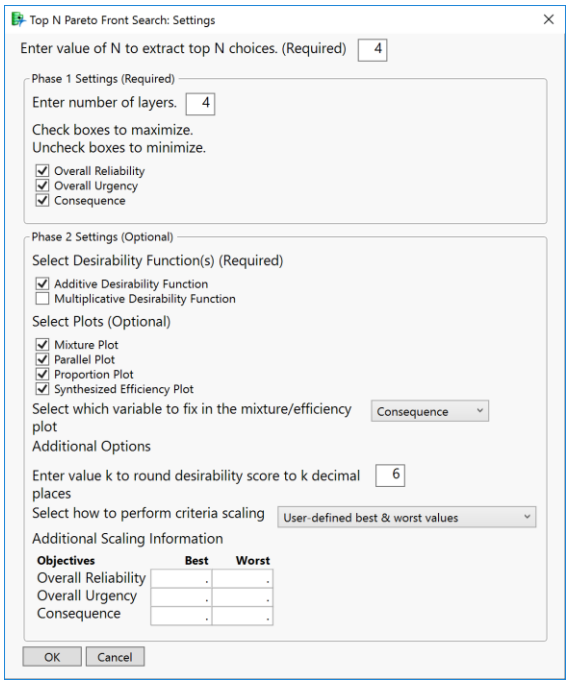

Figure 7. TopN-PFS Settings Window: Scaling Option 4 selected

Desirability Functions (DF)

Additive DF:

$$DF\ Score_j = \sum_{i=1}^{n} w_i z_{ij}$$

Multiplicative DF:

$$DF\ Score_j = \prod_{i=1}^{n} z_{ij}^{w_i},$$

where n is in the number of criteria, z_{ij} are the scaled criteria values, and $\sum_{i=1}^{n} w_i = 1$.

The choice of DF depends on the scenario and priorities of the decision-maker. The multiplicative DF penalizes low criterion values more severely than the additive DF. The choice of DF, therefore, can impact which solutions are highlighted as best during Phase 2 of the algorithm.

Figure 8 shows the settings window filled in for the stockpile example. Each criterion indicates the stockpile is more critical with large values, so they are optimized by maximizing as indicated in the figure. Note that under the list of plots, a drop-down menu appears listing the criteria. The mixture plot and synthesized efficiency plots must fix one of the variables to a value when there are three criteria. This decision is made in the settings window. When there are only two criteria, this step is not necessary.

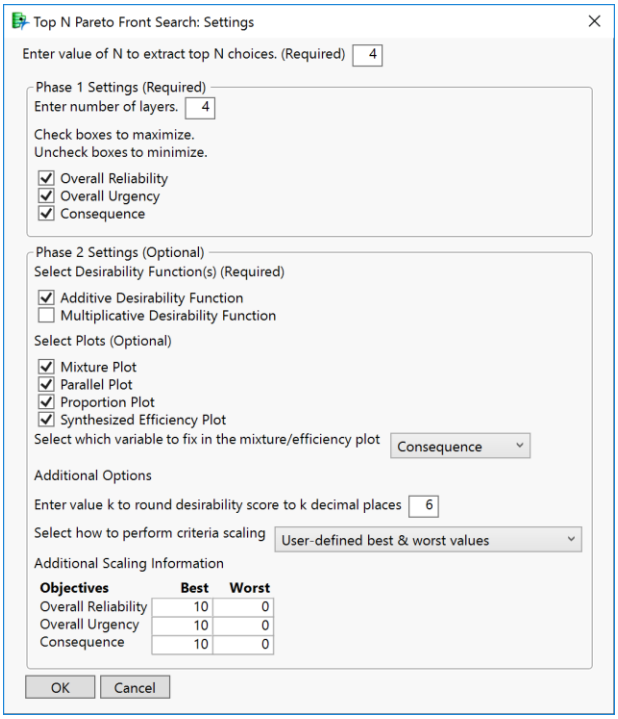

Figure 8. TopN-PFS settings window filled in for stockpile prioritization example

Phase 1 Output

Phase 1 is the objective stage of the algorithm to eliminate non-competitive solutions from further consideration. In this phase, the points in the layered PFs are identified. There are two outputs from this phase, a data table, and a scatterplot matrix.

Data Table: Pareto Front Results All Data

"Pareto Front Results All Data" is a data table with all the original data and columns with an additional column identifying the PF layer for each point. If the point is not on a PF layer, there is no value (listed as missing). The original data table is unchanged. Figure 9 shows a portion of the resulting data table for the stockpile data. Note that the column "Layer" indicates which PF layer each solution is on. Stockpiles which are not on a PF layer have a missing value in the column (for example, stockpiles A11, A12, and A13). Each PF layer also has a unique symbol as indicated on the left side of Figure 9.

6/0 Cols	Stockpile	Overall Reliability	Overall Urgency	Consequence	Layer	Stockpile Number
1	A1	9	6	2.5	4	1
2	A10	7.75	8.75	5	3	10
3	A11	8.5	7	3.5	•	11
4	A12	6.5	6.75	8.5	•	12
5	A13	7.75	7	5.5	•	13
6	A14	8.5	8	4	3	14
7	A15	7.25	8.25	3	4	15
8	A2	6.5	8.75	5.5	3	2
9	A3	9.5	8.75	9	1	3
10	A4	9	8.75	9	2	4
11	A5	7	8.5	4.5	4	5
12	A6	7.25	8	3	•	6
13	A7	8	7.75	3.5	4	7
14	A8	8.5	7.25	3.5	4	8
15	A9	7	6.5	3	•	9
16	B1	7.25	7.25	7	•	16
17	B10	6.25	8.25	6	3	25
18	B11	8.25	7.5	9	3	26
19	B2	8.25	6.25	7	4	17
20	B3	6.5	6.25	8	•	18

Figure 9. Data Table "Pareto Front Results All Data" (output from Phase 1)

Scatterplot Matrix

The scatterplot matrix is an interactive plot used to visualize the points in the layered PF. Each layer is labeled with a unique symbol. You can select any subset of the PF layers using the legend. Points shown with small black dots are not on the top N layers. Figure 10 shows the scatterplot matrix for the stockpile example. Sixteen of the 42 stockpiles are not included on any of the top 4 layers and therefore can be objectively eliminated from further consideration.

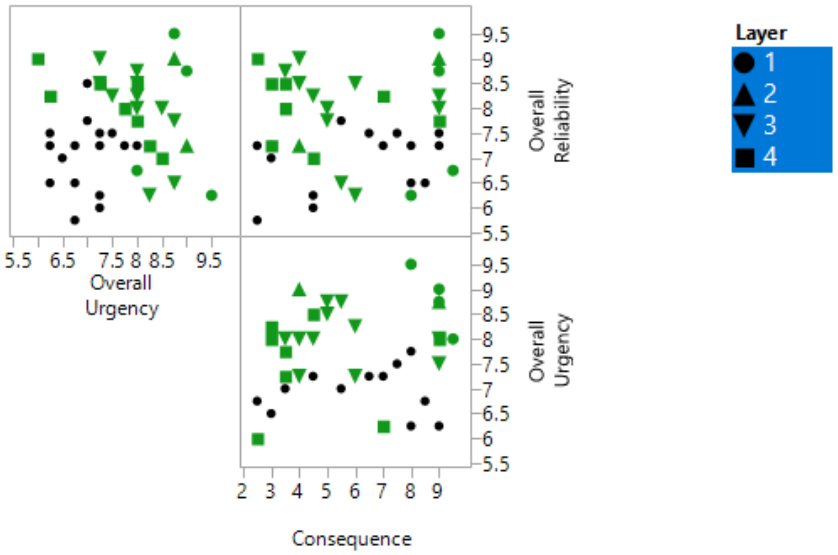

Figure 10. Top 4 Pareto front layers for the stockpile prioritization example

Phase 2 Output

Phase 2 of the algorithm is the subjective stage, which allows user priorities to be incorporated, and represents the Combine and Select steps of the DMRCS strategy. This stage analyzes the tradeoffs between identified solutions in Phase 1. Each criterion is scaled (Scaling) as defined in the settings window. A Desirability Functions score is calculated for different combinations of weights of the criteria. This phase has 5 possible graphical tools to guide the decision-maker to a final decision.

Mixture Plot

The mixture plot uses gray-scale to show the top N solution across all evaluated weight combinations. The goal of this plot is to allow the decision-maker to identify the best solutions for a given set of priorities when making a final decision. This plot allows the experimenter to visually examine tradeoffs as priorities in the criteria vary. The horizontal axis of the plot is the weight placed on one criterion. When there are three criteria, one must be fixed. The local data filter (left side of Figure 11) allows the user to change the weight on this criterion. Because the criteria weights sum to one, the weight placed on the remaining criterion is one minus the weight of the other criteria. In the example shown in Figure 11, Overall Reliability and Overall Urgency sum to 0.7, and Consequence is fixed at 0.3. The vertical axis shows the ID of the different solutions, in this example, stockpiles. Black squares denote the top choice (highest DF score) for a given weight combination, and the lightest gray denotes the Nth choice. White space indicates that a particular option was not a top N solution for that weight combination. Top solutions with tied scores have the same color shading, clearly identifying ties in DF scores. For the stockpile example, at a weight combination of (OR, OU, C) = (0.4, 0.3, 0.3), the most critical stockpile is A3, followed by A4, B4, and D4.

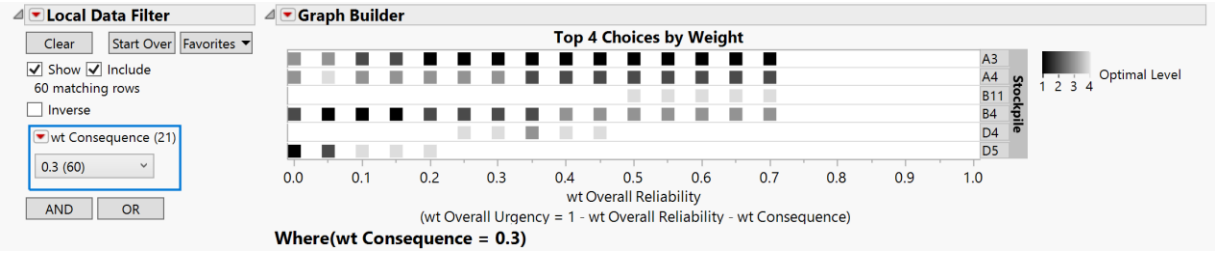

Figure 11. Mixture plot for stockpile prioritization example

Proportion Plot

This plot is a stacked bar chart that indicates how frequently each solution is a top N solution across all of the evaluated weight combinations of the criteria. This plot indicates how robust a particular choice is across the range of criteria weightings. The solutions are sorted from most frequently in the top N to least, and all stockpiles that appear anywhere in the top N for any weight combination are shown. In the stockpile example, 14 stockpiles are in the top 4 across all weight combinations. The method has reduced the number of choices from 42 down to 26 in the top 4 layered PFs, and now down to 14 using

this scaling and DF choice. Stockpiles A3, B4, and A4 are in the top 4 for almost all of the possible weights. Note that stockpile A4 is on the second PF layer, and would never be considered as a solution if the focus has only been on the top PF layer.

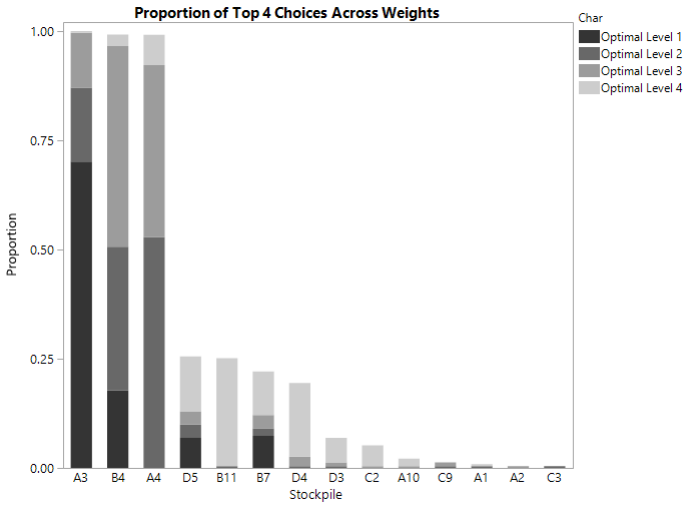

Figure 12. Proportion plot for stockpile prioritization example

Parallel Plot

The parallel plot (Lu and Anderson-Cook, 2014) allows the analyst to visually examine the tradeoffs of one solution over another in terms of the scaled criteria values. The best values for each criterion are on the top of the plot, and the worst values are on the bottom. Each line in the plot corresponds to a particular solution. The lines are on a gray scale where black indicates the highest scaled value for the first listed criterion and the lightest gray indicates the lowest scaled value for the same criterion. If many lines cross, then there are more tradeoffs between the solutions. Relatively steep lines indicate extreme tradeoffs between criteria for a given solution. Flatter lines indicate less tradeoff between the criteria. Figure 13 shows the parallel plot for the stockpile example.

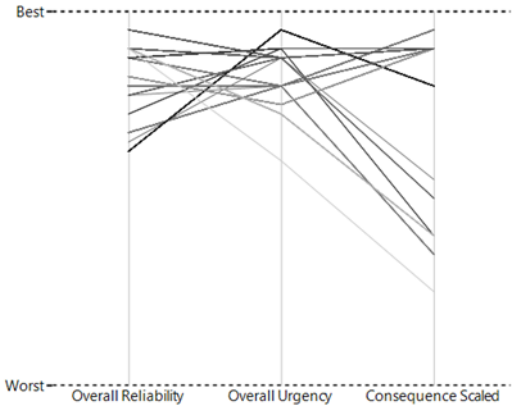

Figure 13. Parallel plot for stockpile prioritization example

Synthesized Efficiency Plot

The synthesized efficiency plot, first introduced by Lu and Anderson-Cook (2014) helps explain the implications of each choice for a given weight combination. For solutions that may or may not be a top N solution for certain weight combinations, their DF scores are compared to the best available at any particular weight combination. The synthesized efficiency value at a given weight combination $\boldsymbol{w}$ is defined as:

$$Syn\ Eff_{\boldsymbol{w}} = \frac{DF\ Score_{\boldsymbol{w}}}{\max(DF\ Score_{\boldsymbol{w}})}$$

By definition, the best available choice at a given weight combination has a synthesized efficiency value of 1. The plot has the same form as the Mixture Plot with the weight of one criterion on the horizontal axis and the ID of each choice along the vertical axis. When there are three criteria, one must be fixed (chosen in the settings window) to a given value. The drop down menu on the left in the local data filter allows you to change the weight value of this criterion. The values of synthesized efficiency are color-coded on a blue to white scale, where dark blue indicates a high efficiency value and white indicates poor performance. A high value for synthesized efficiency means that a given solution has a similar DF score to the best available at that weight combination. Figure 14 shows the synthesized efficiency plot for the stockpile example. Note that stockpile A3 has dark blue consistently over the range of all shown weights while A1 consistently has light blue for these weight combinations. This indicates that A3 is scored as more critical for all of the shown weight combinations, and A1 is relatively non-critical.

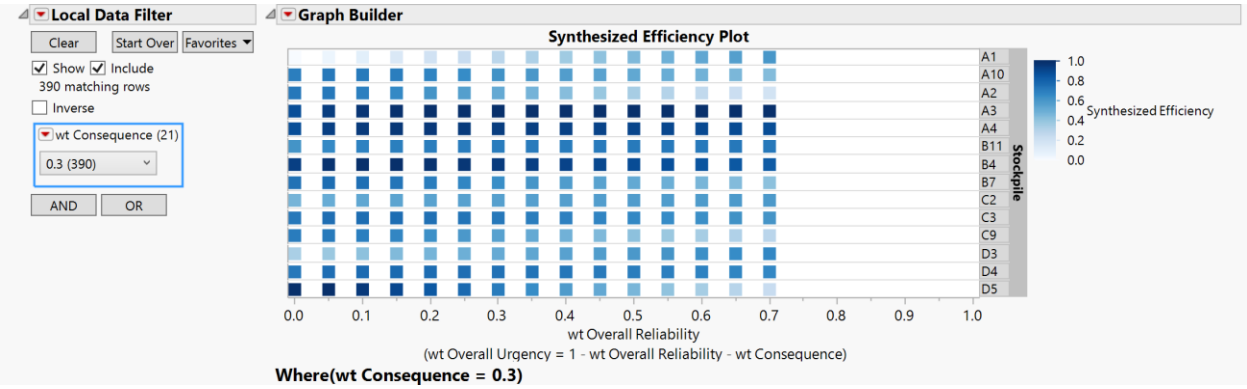

Figure 14. Synthesized efficiency plot for stockpile prioritization example

N Comparison Plot

The TopN-PFS algorithm is dependent on the choice of N selected by the user. In the stockpile example, N=4 since we are looking to identify the top 4 choices to receive additional funding. There may be some concern that the final decision may be strongly impacted by the choice of N. The N comparison plot helps the user decide whether they should consider a larger value of N in the algorithm. For each weight combination, the ratio of DF score for the (N+1)th best solution over the DF score for the Nth best solution is calculated:

$$NComp_{\boldsymbol{w}} = \frac{DFScore_{\boldsymbol{w},(N+1)}}{DFScore_{\boldsymbol{w},(N)}}$$

This ratio is bounded between 0 and 1. Low values indicate that the next best solution after the Nth choice is not very close to the Nth best solution. A high value (close or equal to 1) indicates that the (N+1)th solution has a DF score very close to the Nth best solution at that weight combination. High values suggest that the decision-maker may want to consider a larger value of N in the search, since there is no clear distinguishing feature between the top N and the next best choice. Note that a value of 1 indicates a tie in the N and (N+1)th solutions. By default, all tied solutions are automatically included as potential solutions in the TopN-PFS algorithm. Figure 15 shows the N comparison plot for the stockpile example. For many of the weights considered in the plot, the fifth best solution has very close performance compared to the fourth best choice.

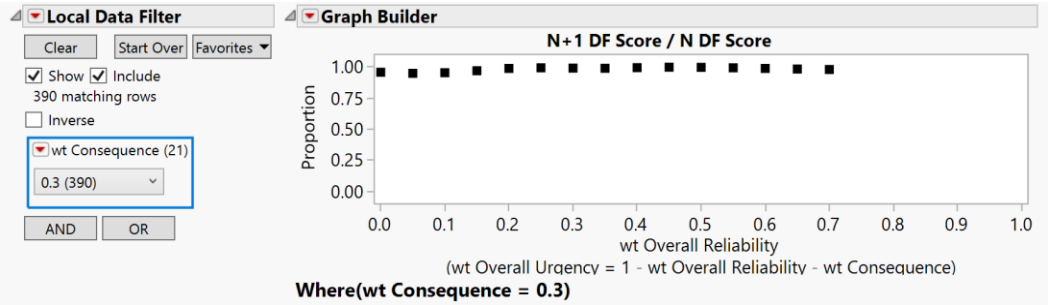

Figure 15. N comparison plot for stockpile prioritization example

Additional Output Tables

Results from the algorithm (used to create the five graphics described above) are available in three tables, initially hidden from view. Access to these tables is available at the bottom of the second tab ("Graphs") of the results window. By clicking their respective buttons, these tables can be opened.

- **Mixwgt**: A summary of the top N solutions for each weight combination. The table shows the weight values, ID, scaled criteria values, criteria values in their original units, the desirability function score, the Pareto Front Layer, the Synthesized Efficiency, and the optimality level (1, 2, …, N)
- **Prop Best**: A table showing the proportion a given solution was the best, second best, …, Nth best choice.
- **DF Summary Table**: A summary table showing the desirability function score and synthesized efficiency value for each point on the layered Pareto Front for each weight combination.

Miscellaneous Notes

- As of this version, if you would like to re-run this function on a data table, you must first close any opened results windows. Closing the results window will close the associated data tables that are produced from this tool.
- To uninstall this Add-In, go to the view menu, and scroll to "Add-Ins". Select the Add-In in the menu and then click "Unregister."

References

Anderson-Cook, C.M. and Lu, L. “Much-Needed Structure: A New 5-Step Decision-Making Process Helps You Evaluate, Balance Competing Objectives.” *Quality Progress*, vol. 48, no. 10, 2015, pp. 42–50.

Burke, S.E., Anderson-Cook, C.M., Lu, L., Borror, C.M., Montgomery, D.C. “Design of Experiment Selection using Layered Fronts” *Los Alamos National Laboratory Technical Report*, LAUR-16-25643, 2016.

Derringer, G. and Suich, R. "Simultaneous Optimization of Several Response Variables" *Journal of Quality Technology*, vol. 12, 1980, pp. 214-219.

Lu, L., and Anderson-Cook, C. M. “Balancing Multiple Criteria Incorporating Cost Using Pareto Front Optimization for Split-Plot Designed Experiments.” *Quality and Reliability Engineering International,* vol. 30, no. 1, 2014, pp. 37-55.

Lu, L., Anderson-Cook, C.M., and Robinson, T.J. "Optimization of Designed Experiments Based on Multiple Criteria Utilizing a Pareto Frontier" *Technometrics,* vol. 53, 2011, pp. 353-365.