Visualization of Global Sensitivity Analysis Results Based on a Combination of Linearly Dependent and Independent Directions

Misty D. Davies* and Karen Gundy-Burlet†
NASA Ames Research Center, Moffett Field, CA, 94035

A useful technique for the validation and verification of complex flight systems is Monte Carlo Filtering—a global sensitivity analysis that tries to find the inputs and ranges that are most likely to lead to a subset of the outputs. A thorough exploration of the parameter space for complex integrated systems may require thousands of experiments and hundreds of controlled and measured variables. Tools for analyzing this space often have limitations caused by the numerical problems associated with high-dimensionality and caused by the assumption of independence of all of the dimensions. To combat both of these limitations, we propose a technique that uses a combination of the original variables with the derived variables obtained during a principal component analysis.

Nomenclature

α = The angle between the plotting plane and the planes of the projected hyperrectangle

I. Introduction

Finding bugs in modern flight software is an arduous process. Interdependencies between the systems and subsystems of the flight vehicle and the environment are usually explored in the design phase using high-fidelity simulations. At this level, the interplay between the interfaces can be fully explored. Advances in computing power make it possible to try many thousands of combinations of input variables to exercise the behavior of the system. However, a thorough exploration of the entire flight envelope results in gigabytes of data that requires expert domain knowledge to interpret; and because of the high dimensionality of the data, it is difficult to pick out patterns with the human eye.

The Robust Software Engineering (RSE) group within the Intelligent Systems Division at NASA Ames Research Center has developed a multi-step Monte Carlo Filtering technique[1,2] called the Margins Analysis[3-5] that automates much of the process necessary to identify unusual behaviors within the system and to isolate the critical ranges for suspicious inputs. We develop input test vectors using an n-factor combinatorial process[9]—the assumption is that most bugs are triggered by a maximum combination of two or three input variables; by guaranteeing that each combination of n input variables (where n is 2 or 3) appears in the test suite at least once we get a coverage guarantee while limiting the number of total test vectors within the suite. The output

* Research Computer Engineer, Intelligent Systems Division, Mail Stop 269-3, AIAA Member.
† Research Scientist, Intelligent Systems Division, Mail Stop 269-3, AIAA Associate Fellow.

behavior is analyzed for non-compliance with requirements and for unexpected structure using a combination of unsupervised[6] and supervised[7,8] machine learning techniques. Failures or unexpected results are collected and, when enough examples of a particular undesired behavior have been gathered, the data is automatically examined to determine which inputs (and ranges) demonstrate the most sensitivity with respect to the undesired behavior. This information is exploited in order to cause the undesired behavior to occur more often in future runs, and the aggregated data over all of the iterations can be used to identify bugs within flight software or within the simulation, or to determine the margins to failure as the systems pass from nominal to off-nominal behavior.

In current practice at NASA, the analysis can involve several hundred parameters over tens of thousands of trials. The unsupervised learning technique, a clustering technique that utilizes an expectation-maximization algorithm, has numerical difficulties at this level of dimensionality; we use a principal component analysis to reduce the dataset. The supervised learning technique, a treatment learning algorithm developed at West Virginia University, does not suffer from the same numerical difficulties as the clustering technique and can handle large numbers of parameters effortlessly. However, the treatment learning algorithm assumes that all of the parameters are independent; this assumption, when false, causes parameter ranges to be much larger than necessary and can cause the treatment learner to miss important sensitivities. A treatment learning analysis performed solely in the principle components space benefits from the fact that each transformed variable is linearly independent, but suffers from the difficulty of understanding the results in the transformed space. Each new direction in the principal component analysis is a linear combination of the original variables, so the answer that the treatment returns is a hyperrectangle in the space of the original variables.

We utilize the ability of the treatment learner to handle many dimensions, and ask it to perform an analysis over the data simultaneously in the original space and in the rotated space of the principle component analysis. As a post-processing step, we reduce the hyperrectangles in the principle component space to two- and three-dimensional projections in the original variable space. We then combine all of the information gained from the non-rotated and rotated spaces into two- and three-dimensional plots in the original variable space, so that they can be understood and interpreted by domain experts.

II. Methodology

We will discuss at length how we perform the principal component analysis, how we reduce dimensionality for the clustering algorithm, what data is given to the treatment learner and what data the treatment learner returns, how we choose the number of plots for each treatment, how we determine the dimensionality and axes for each plot, and how we plot the boundaries for the treatments.

III. Results

In standard practice, we ask the treatment learner to return a maximum of 10 treatments. Each treatment can consist of up to 4 variables. We place no restriction on whether these variables are in the original space or the principle component space. Ranges in the principle component space are first transformed into hyperrectangles in the original space, and then projected into two- or three-dimensional plots depending on the value of α, the angle between the primary two-dimensional plotting plane and the parallel planes of the projected hyperrectangle. The primary two-dimensional plotting plane is defined by the maximum components of the unit vector in the principle component direction when written as a linear combination of the variables in the original directions.

Figure 1 shows an example treatment for a dataset obtained during a bicycle ride. There were 10 variables for this dataset, and each of the 4435 time steps in the series was treated as a different experiment. The power calculation from the measured variables in the dataset was particularly noisy. Each data point in the set was rated according to the noise in the power calculation for some small time period surrounding the data point; the colormap goes from blue to red with the red points in the plots corresponding to the noisiest data points. The treatment returned for Figure 1 had 3 components (chosen by the algorithm from 20 components—10 components in the original variable space and the 10 rotated directions): cadence, wind speed, and a principle component direction in 10-dimensional space where the two largest components when written in the original space were hill slope and wind speed. The two variables in the original space are plotted together as a two-dimensional plot on the left, and the red box shows that the noisiest data occurs at the highest wind speeds and cadence. The angle α for the

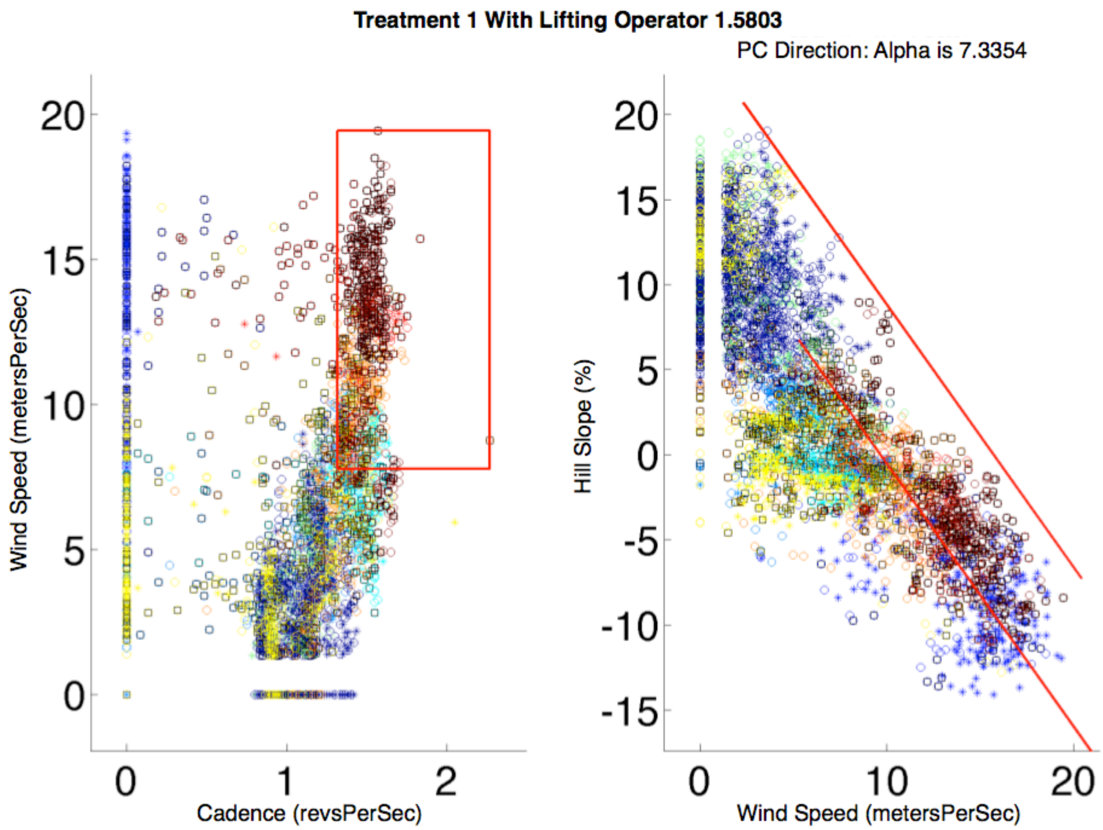

Figure 1. An example treatment for data obtained during the operation of a bicycle. *The data are color-coded from blue to red with red corresponding to the noisiest power data and blue corresponding to the least noisy. The red lines outline the regions identified by the treatment learner as the most likely to contain the undesirable noisy data. In this treatment, it is clear to see from the first plot that the undesired data (in this case, noisy power data, plotted in red and outlined with a black box) occurs most often at high wind speeds when the rider is continuing to pedal the bicycle. This second plot in the treatment is a projection of a principal component analysis direction result into the original space. This plot shows us that the hill slope and the wind speed have a strong linear correlation and that the least desirable data occurs for a linear combination of high wind speed and high hill slope. The two plots together constitute a visual representation of the highest scoring treatment.*

principal direction isolated in this treatment is less than 10 degrees, so the principal component information is plotted as a two-dimensional plot on the right. The red lines show the intersection of the almost vertical treatment hyperplanes with the plane defined by the hill slope and wind speed. What immediately becomes apparent from the plot on the right is that the hill slope and wind speed are negatively linearly correlated, so that the highest wind speeds occur at the steepest downhill slopes. Interpreting the two plots together leads to the conclusion that the calculated power measurement is most likely to be faulty when the cyclist is free pedaling on the downhill slopes.

IV. Conclusion

Acknowledgments

This research was conducted at NASA Ames Research Center. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government.

References

[1]Rose, K., Smith, E., Gardner, R., Brenkert, A. and Bartell, S. "Parameter Sensitivities, Monte Carlo Filtering, and Model Forecasting Under Uncertainty," *Journal of Forecasting*, Vol. 10, 1991, pp. 117-133.

[2]Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., and Tarantola, S., *Global Sensitivity Analysis: The Primer*, Wiley, Chichester, 2008, Chaps. 1, 5.

[3]Gundy-Burlet, K., Schumann, J., Barrett, T., and Menzies, T., "Parametric Analysis of ANTARES Re-entry Guidance Algorithms Using Advanced Test Generation and Data Analysis," *9th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2007.

[4]Schumann, J., Gundy-Burlet, K., Pasareanu, C., Menzies, T., and Barrett, T. "Tool Support for Parametric Analysis of Large Software Systems", *Proceedings of Automated Software Engineering, 23rd IEEE/ACM International Conference,* 2008.

[5]Gundy-Burlet, K., Schumann, J., Barrett, T., and Menzies, T., "Parametric Analysis of a Hover Test Vehicle Using Advanced Test Generation and Data Analysis," *AIAA Aerospace*, 2009.

[6]Fischer, B., and Schumann, J. "Autobayes: A System for Generating Data Analysis Programs From Statistical Models," *Journal of Functional Programming*, Vol. 13, 2003, pp. 483-508.

[7]Hu, Y., "Treatment Learning: Implementation and Application," Masters Thesis, Department of Electrical Engineering, University of British Columbia, 2003.

[8]Hu, Y., and Menzies, T. "Data Mining for Very Busy People," *IEEE Computer*, Vol. 36, No. 11, 2003, pp. 22-29.

[9]Barrett, A., "A Combinatorial Test Suite Generator for Gray-Box Testing", *Third IEEE International Conference on Space Mission Challenges for Information Technology*, 2009.