SAND2013-10501P

# Ensemble Analysis of Electrical Circuit Simulations

Patricia J. Crossno, Timothy M. Shead, Milosz A. Sielicki, Warren L. Hunt, Shawn Martin, and Ming-Yu Hsieh

## 1 Ensembles and Sensitivity Analysis

With recent advances in computational power, scientists can now run thousands of related simulations to explore a single problem. We refer to such a group of related simulations as an *ensemble*. More generally, an ensemble can be thought of as a set of samples, each consisting of the same set of variables, in a shared high-dimensional space describing a particular problem domain. Practically, an ensemble is a collection of data sets with common attributes that we wish to analyze as a whole. Thus *ensemble analysis* is a form of meta-analysis that looks at the combined behaviors and features of the collection in an effort to understand and describe the underlying problem space. By looking at the ensemble as a whole, higher level patterns emerge beyond what can be seen by examining individual simulation runs.

As an example, a form of ensemble analysis called *sensitivity analysis* evaluates how changes in simulation input parameters correlate with changes in simulation results. In addition to revealing the types and strengths of relationships between inputs and outputs, sensitivity analysis can be used to verify that simulation results are within expected ranges and to validate that the underlying model is behaving correctly. Unexpected results can identify simulation errors and reveal flaws in simulation codes. Input parameters form the set of independent variables, and outputs the set of dependent variables. Commonly, sensitivity analyses are performed using either simple regression, correlating a single input to a single output at a time, or multiple regression, correlating a group of inputs to a single output. However, neither of these approaches provides a means for evaluating the collective relationships among multiple inputs and multiple outputs.

---

Patricia J. Crossno, Timothy M. Shead, Milosz A. Sielicki, Warren L. Hunt, Shawn Martin, and Ming-Yu Hsieh
Sandia National Laboratories, PO Box 5800, Albuquerque, NM 87185 e-mail: pjcross@sandia.gov, tshead@sandia.gov, masieli@sandia.gov, wlhunt@sandia.gov, smartin@sandia.gov, myhsieh@sandia.gov

Our introduction to this work began during an evaluation of the impacts on analysis and visualization workflows that were likely to result from architectural changes being proposed for exascale computing. As part of the evaluation, we interviewed analysts working at Sandia National Laboratories in a variety of simulation domains, including thermal, solid mechanics, and electrical circuit analysis. Sensitivity analysis is a common component within each domain's work flow, although the types of result data vary widely, ranging from simple tables of metrics to time series and finite element models. The analysts typically use Dakota (Adams et al, 2013) to manage ensemble creation, using custom scripts to extract scalar metrics from finite element outputs or time series. The metrics are merged with tables of original input parameters and analyzed using Dakota, JMP, Matlab, Minitab, Excel, and other tools. Existing visualization tools such as ParaView (Kitware, 2013) and EnSight (Computational Engineering International, 2013) are used for remote visualization of large simulation results, but are fundamentally designed to visualize individual simulations, or handfuls of simulations that are loaded into memory simultaneously and visually superimposed. Ensembles containing hundreds or thousands of simulations require a different type of analysis, a different visual abstraction, and a different system architecture to effectively manage integrating so many results.

This investigation led to the creation of Slycat, a system designed to meet the needs of ensemble analysis. For sensitivity analysis and parameter studies in particular, Slycat provides a visual interface to answer the following questions about a given ensemble:

- Which input parameters are most strongly correlated with specific outputs?
- Which input parameters exhibit little or no impact on output responses?
- Which simulations are anomalous, and in what ways do they differ?

We use canonical correlation analysis (CCA) to model the relationships between input and output parameters because it maps well to the structure of our problem, especially in its ability to correlate multiple inputs against multiple outputs. Although powerful and ideally suited to the problem, CCA results can be difficult-to-interpret; thus, the central contribution of this work has been making CCA accessible to domain experts through the tight integration of useful visualizations with iterative exploratory analysis.

## 2 Related Work

Previous work exploring ensemble visualization often focuses on the design of experiments, and is often less about understanding the simulations than with directing them to a particular outcome. For example, in (Bruckner and Moller, 2010) and (Coffey et al, 2013), the work is focused on providing a visual evaluation of the simulations, allowing a user to steer toward an optimal aesthetic result. This results-driven approach is directed at finding a "magic" spot within the input parameter space, based on an individual's aesthetic sense. Our work focuses on contributing confidence and understanding to physically valid modeling and simulation processes.

The work of (Matkovic et al, 2010) recognizes the need for "advanced tools" to support engineers in visualizing and understanding ensembles, and incorporates multivariate visualization techniques including parallel plotting, multiple linked views, and scatterplots to display one-to-one correlations.

Other work wrestles with the multivariate nature of the ensemble, but the data is often inherently spatial and the visualization techniques rely on this. (Wilson and Potter, 2009) explore geospatial weather data, and the authors discuss how ensembles mitigate uncertainty in simulations, a common thread throughout this research. Other work on weather simulations, (Potter et al, 2009b) and (Potter et al, 2009a), employs isocontours over spatial domains. (Sanyal et al, 2010) adds statistical metrics represented as ribbons and glyphs to communicate the inter-simulation uncertainty present in an ensemble of weather predictions.

Feature extraction can be a useful tool for ensemble visualization, as demonstrated in (Smith et al, 2006). The authors work with time-varying, spatial data to cluster based on feature identification. (Hummel et al, 2013) is another excellent example of feature extraction, used to visualize fluid flow variance across an ensemble. Linked views provide selection in the feature space to produce a visualization over the physical domain. (Steed et al, 2013) incorporates several multivariate feature detection techniques in a single interface, while (Piringer et al, 2012) visualize multivariate data using downsampling, 3D surface plots, extracted scalar features, and glyph-based visualizations to explore an ensemble of 2D functions. In addition to comparing ensemble members against each other, the latter work attempts to illustrate the distribution of features across the ensemble.

(Sukharev et al, 2009) use feature detection and CCA to reveal structure in multivariate data, demonstrating their analysis on time-varying climate data sets. Once their data has been clustered, segmented, and correlations computed, the results are geo-spatially overlaid on the weather prediction region for visualization and interpretation.

To more clearly identify the relationships in functional magnetic resonance imaging data sets, (Karhunen et al, 2013) exploit CCA prior to applying blind source separation techniques and achieve marked performance improvements. Other work (Ge et al, 2009) demonstrates CCA correlations that reflect the spatial correlations in multiple sensor arrays, even in the presence of noise. Another application of CCA (Degani et al, 2006) analyzes the correlations between the operating environment of a Boeing aircraft and the actions and responses of the pilots. In (Marzban, 2013) CCA is shown to capture complex weather relationships between model parameters and forecast quantities.

Sensitivity analysis can be defined as *the determination of the contributions of individual uncertain analysis inputs to the uncertainty in analysis results* (Helton, 2008). Our research is focused on understanding the behavior of an ensemble with the intent of exposing hidden relationships between the simulation input parameters and the results, which is similar to sensitivity analysis without the emphasis on numerical quantification of uncertainty. Sampling tools, (Adams et al, 2013) and

(Abdellatif et al, 2010), are typically relied on to provide coverage of the simulation parameter space. Even with a sampling method in place, much system behavior is unknown and there is more work to do to uncover those input-output relationships. (Song and Zhao, 2012) employ a variance-based method to identify the first-order model sensitivities when applied to forest growth simulations. Other work (Shearer et al, 2010) applies classic statistical methods, such as ANOVA, to U.S. immigration model results, and statistical aggregation to models of large distributed computing systems (Mills et al, 2011).

## 3 System Architecture

To support answering the questions outlined in Sect. 1, and to support additional analysis types in the future, we designed Slycat around the following general requirements:

- Remote ensemble analysis, in which large data is analyzed in-place to minimize data movement.
- Ubiquitous access to analysis results regardless of the availability of the original data sets or source platforms.
- Desktop delivery providing interactive exploration of ensemble analysis results, and collaborative sharing with appropriate access controls.

The need for remote ensemble analysis is driven by the ever widening gap between high performance computing (HPC) compute performance and I/O performance. Practically speaking, we have reached a point where computation is effectively "free" while data movement has become very expensive, and moving raw ensemble data from the HPC platform where it was generated to the host running Slycat would take significantly more time than the analysis computations to follow! Better instead to perform those computations on the machine where the ensemble data is located, so that only the model – typically orders of magnitude smaller than the original data – must be moved across the network to the Slycat host. This led Slycat to the design of Fig. 1.

An important practical consideration for users of HPC platforms is that ensemble results may often become temporarily or permanently unavailable – login nodes come and go due to resource contention, users often must archive or delete their data as scratch filesystems near capacity, and so on. Because Slycat stores its own greatly-reduced models of the underlying raw data, and only those models are necessary to produce a visualization, users can continue to access their Slycat analysis results even when their HPC resources are unavailable.

Finally, we wanted a system architecture that could support easy desktop delivery and collaboration: for example, we wanted Slycat users to be able to share results seamlessly with a colleague, across the network, without any software downloads or installation. That meant using existing web standards and clients, and dictated

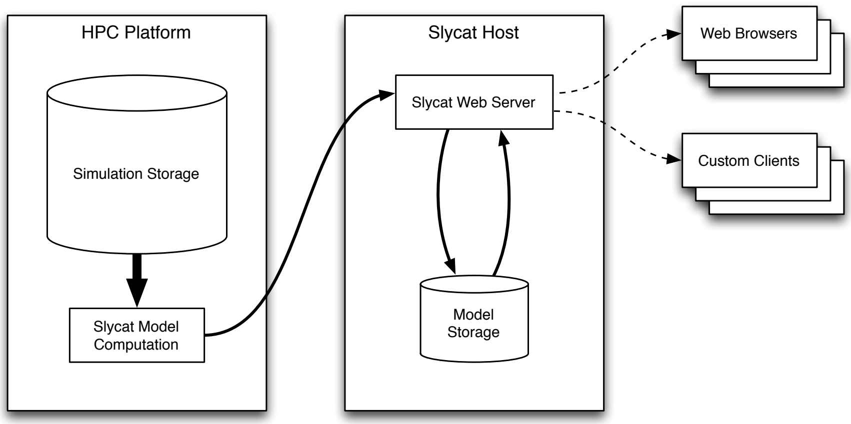


**Fig. 1** Slycat system diagram depicting how large data on an HPC platform is analyzed in-place to produce greatly reduced model artifacts that are stored by the Slycat web server. Later, these artifacts are delivered – incrementally and on-demand – to interactive clients.

much of the subsequent design and derived requirements for Slycat: straightforwardly, it meant adopting a web server as the front-end for the system, and standard web browsers as clients (or custom clients using standard web protocols to communicate). In-turn, this meant that we had to design interactions and visualizations that could work using the set of technologies widely available within web browsers, such as HTML5, JavaScript, AJAX, SVG, and Canvas. Unlike dedicated visualization tools such as ParaView or Ensight, we could not rely on the client to perform serious calculations, necessitating pre-computation of visualization artifacts, organized for rapid, incremental retrieval from the server on-demand. As an example, we allow users to interact with data tables that can contain thousands of columns and millions of rows - data that would cause unacceptable delays if it had to be transferred from server to client in its entirety before the visualization could be viewed. Instead, only the subset of data that is needed to display the currently-visible table rows is transferred “just-in-time”, minimizing total bandwidth consumption and keeping the interface responsive. Although working around the constraints of web browsers has been a challenge, the rewards have been significant, enabling Slycat users to “bookmark” the state of a visualization and share it with colleagues simply by sharing a hyperlink.

The Slycat source code and documentation are freely available under an open source license, at https://github.com/sandialabs/slycat.

## 4 Canonical Correlation Analysis

Canonical Correlation Analysis (CCA) was proposed by H. Hoteling (Hotelling, 1936) in 1936 (he was also instrumental in developing Principal Component Analysis in 1933). CCA is a method that can be used to understand relationships between two sets of multivariate data. One set $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ is presumed to be independent, and the other set $Y = \{\mathbf{y}_1, \ldots, \mathbf{y}_n\}$ is dependent, where we have $n$ samples, with $X \subset \mathbb{R}^{p_1}$ and $Y \subset \mathbb{R}^{p_2}$ (i.e. each vector $\mathbf{x}_i$ has $p_1$ components and each vector $\mathbf{y}_j$ has $p_2$ components). CCA attempts to find projections $\mathbf{a}$ and $\mathbf{b}$ such that $R^2 = \mathrm{corr}(\mathbf{a}^T X, \mathbf{b}^T Y)$ is maximized. The vectors $\mathbf{a}$ and $\mathbf{b}$ are known as the first pair of canonical variables, and are computed by solving an eigenvalue problem (Anderson, 2003). Further pairs of canonical variables are then identified such that they are all orthogonal and ordered by decreasing importance. For each $R^2$ computed, various statistics can be computed to determine the significance of the correlation. A common statistic used in this context is the $p$-value associated with Wilks' $\lambda$ (Krzanowski, 1988).

Once the canonical variables are determined, they can be used to understand how the variables in $X$ are related to the variables in $Y$, although this should be done with some caution. The components of the vectors $\mathbf{a}$ and $\mathbf{b}$ can be used to determine the relative importance of the corresponding variables in $X$ and $Y$. These components are known as canonical coefficients. However, the canonical coefficients are considered difficult to interpret and may hide certain redundancies in the data. For this reason, it is more typical to analyze the canonical loadings, also known as the structure coefficients. The structure coefficients are given by the correlations between the canonical variates and the original variables (e.g. $\mathrm{corr}(\mathbf{a}^T X, X)$). The structure coefficients are generally preferred to the canonical coefficients due to the fact that they are more closely related to the original variables.

CCA is well-known in the field of statistics and is included in most statistical software packages. However, it has not received as much use in applications as the methods it generalizes, such as multivariate regression and Principal Component Analysis.

## 5 Visualization

In Slycat, sensitivity analysis is performed through an iterative cycle of variable selection, CCA analysis, and visual exploration of the resulting CCA model. Analysis typically starts with an all-to-all evaluation to get an initial sense of the data, revealing the most strongly correlated combinations of variables. Some cases require iterative refinement to tease apart disparate groups of inputs and outputs.

We combine three levels of representation in a single web page using multiple linked views, as shown in Fig. 2. In the upper left, the *Correlation View* represents the relationships found in the ensemble as a whole, displaying CCA coefficients for each variable in tabular form, grouped by correlation component into columns. All

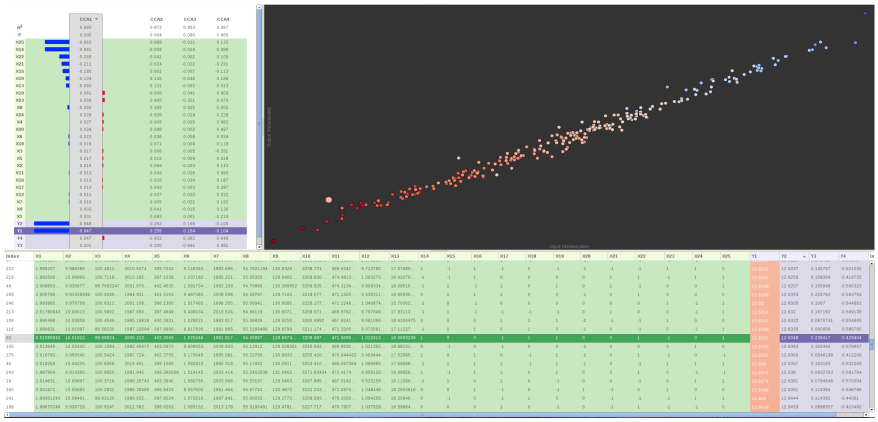

**Fig. 2** Slycat visualization showing an all-to-all CCA analysis of a 250 simulation electrical circuit ensemble. As seen through the length and shared color/direction of the bars in the *Correlation View* in the upper left, the first CCA component exhibits a strong positive correlation predominantly between the combined inputs X25 and X14 and both of the outputs Y1 and Y2, an example of a many-to-many relationship.

but the top two table rows correspond to variables, which are labeled along the left edge of the view. The first two rows provide each component's $R^2$ and $p$-value. Input variables/rows are colored green, while output variables/rows are lavender. This green/purple color-coding is used consistently throughout the interface to designate inputs and outputs. Each column can have its rows sorted by correlation weights, in ascending or descending order. Sorting is performed within each column independently.

Users select a component by clicking its column header, expanding its coefficients into an inline bar chart. The bars visually encode the signed correlation weights for each variable, with left-facing blue bars representing negative weights and right-facing red bars representing positive weights. Color-coding the bars visually reinforces the relationship types. Variables with matching colors are positively correlated, variables with mismatched colors are inversely correlated. Bar length indicates a variable's relative importance. Sorting a component in descending order forms a tornado plot, which makes it simple to evaluate which variables exhibit strong correlations, what types of relationships exist between variables, and which inputs have the most impact on the results.

The scatterplot in the upper right of Fig. 2 is the *Simulation View*. It visualizes how well individual simulations are described by the correlations of the ensemble as a whole. The coordinate space for this plot is abstract, and the axes can be considered to be input and output metavariables, with the x-axis representing all the inputs and the y-axis all the outputs. The axes are the projections of the data onto the canonical variables, **a** and **b**. Each simulation is rendered as a point in the scatterplot, with coordinates computed as sums of variable values, weighted by the canonical variable values for the selected CCA component. Consequently, the scatterplot changes when selecting a new component. A perfect correlation between inputs and outputs would form a diagonal line, while anomalous simulations frequently appear as positional outliers.

Additionally, points can be color-coded by the values of any of the input or output variables, providing another way to identify outliers. Clicking on a row in the *Correlation View* or a column header in the *Variable Table* (see below) selects that variable's values for display in the scatterplot. We use Moreland's blue/white/red diverging color map (Moreland, 2009) for encoding the values, where blue is at the low end of the scale and red at the high end. While we did not assign any particular meaning to the central values, we found in testing that the diverging color map made it easier to interpret values of nearby points.

Across the bottom of the display, the *Variable Table* displays the raw variable values from each simulation. Each column corresponds to a variable and each row to an individual simulation. There is a bi-directional link between selecting table columns and bar chart rows. Selecting a column not only changes the color-coding of the scatterplot points, but also correspondingly colors the column element backgrounds to visually correlate the two views. Using the same interface as the bar chart, the table columns can be sorted. There is a bi-directional link between selecting rows in the table and points in the scatterplot. Darker green/purple backgrounds highlight the table rows, while selected points in the scatterplot are enlarged.

## 6 Electrical Simulation Sensitivity Analysis

Our users model circuits using Xyce, an open source, high-performance circuit simulator developed by Sandia National Laboratories as part of the Advanced Simulation and Computing (ASC) Program (Thornquist et al, 2013). We will look at two circuit ensembles of differing scales, first a small ensemble of 250 runs, followed by a large ensemble of 2641 runs. In both cases, some of the input variables take a restricted set of values (-1, 0, or 1). These values correspond to selecting different models whose responses are *low*, *nominal*, and *high*, respectively. The models act to encapsulate groups of input variables, thereby reducing the number of variables and the number of runs needed in the ensemble.

### *6.1 Small Ensemble*

This ensemble has 250 runs, each with 25 input parameters and 4 output metrics. Outputs Y1 and Y2 measure voltage responses, while Y3 and Y4 measure current. The goal with this analysis was to answer the first two questions from the list in Section 1: Which input parameters are most strongly correlated with specific outputs? Which input parameters exhibit little or no impact on output responses?

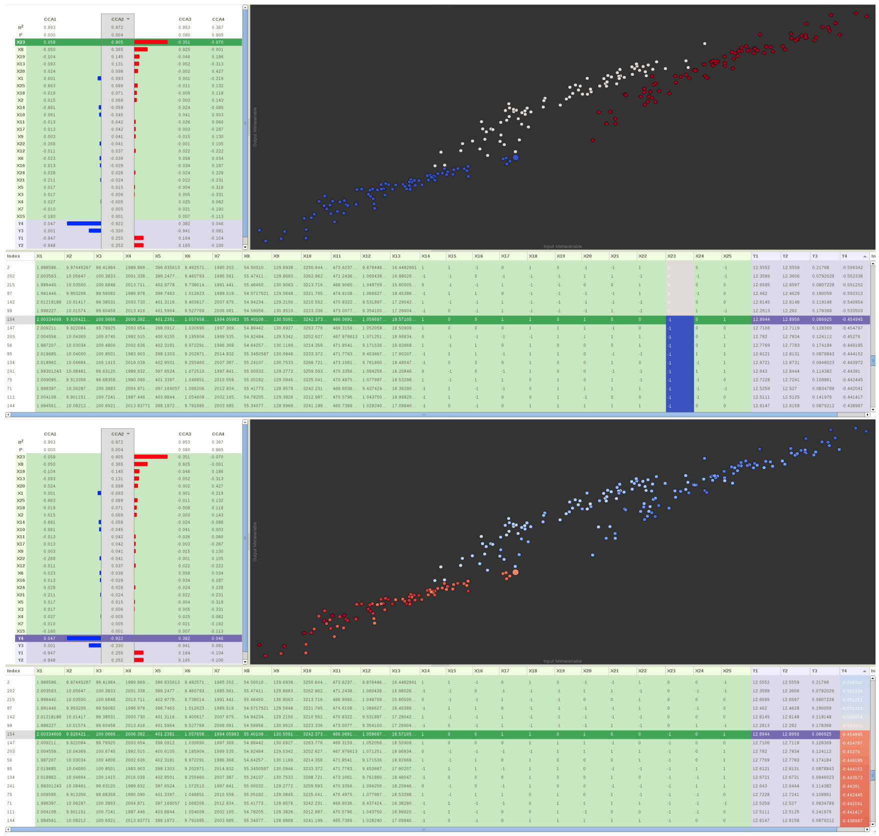

**Fig. 3** In the second CCA component, the input parameter X23 is highly correlated with the current response, Y4. Color-coding the scatterplot by the values of X23 (top image) and Y4 (bottom image), we can see that there is a one-to-one correspondence between the *low*, *nominal*, and *high* values of X23 and 3 distinct groups in the Y4 response values. In both images, the table rows are sorted by Y4 value and we have selected a simulation on the boundary between two of the Y4 groups. Note that this run is also on the edge of the transition between the *low* and the *nominal* value groups in X23.

As seen in the *Correlation View* bar chart in Fig. 2, the first CCA component shows a positively correlated relationship that is mostly between the input parameters X25 and X14 and both of the voltage outputs. The inputs are listed in decreasing order of importance, so the parameters at the bottom of the green region in the first column exhibit little or no impact on voltage responses.

In Fig. 3, the sorted variables for CCA2 reveal a strong inverse correlation predominantly between the input X23 and the current response, Y4. Color-coding the runs alternatively by the first the input values, then the outputs, we can see a one-to-one correspondence between the *low*, *nominal*, and *high* values of X23 and 3 distinct groups in the Y4 response values. Although the other current response, Y3, is also to a lesser degree present in CCA2, it is more strongly described by CCA3,

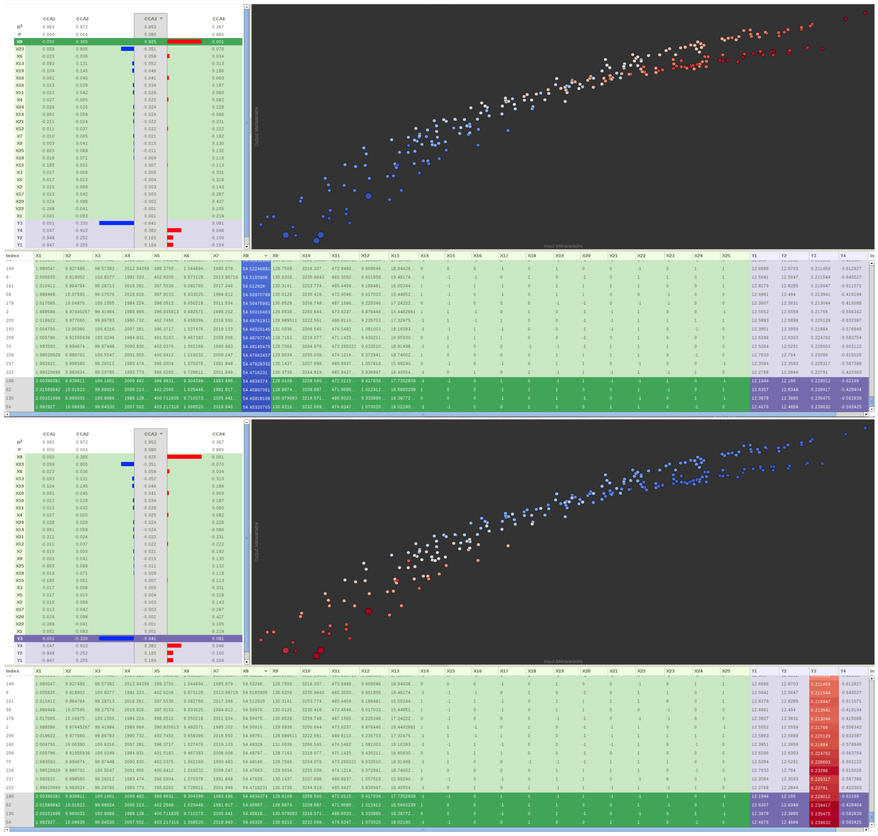

**Fig. 4** In the third CCA component, the input parameter X8 is inversely correlated with the other current response Y3. Color-coding the scatterplot by the values of X8 (top image) and Y3 (bottom image), we can see the inverse relationship between low values (dark blue) in the input parameter and high values (dark red) in the output. In both images, the table rows are sorted by decreasing X8 values and we have selected the four runs with the lowest values in X8.

as shown in Fig. 4. The central relationship is an inverse correlation between the input X8 and the output Y3.

## *6.2 Large Ensemble*

This ensemble from a different circuit simulation is an order of magnitude larger with 2641 runs, including 266 input variables and 9 outputs. The outputs for this circuit are more varied than the previous circuit, so they do not fall into the two simple categories of voltage and current. Rather they capture events and features, where some of the outputs are naturally grouped together and others are not.

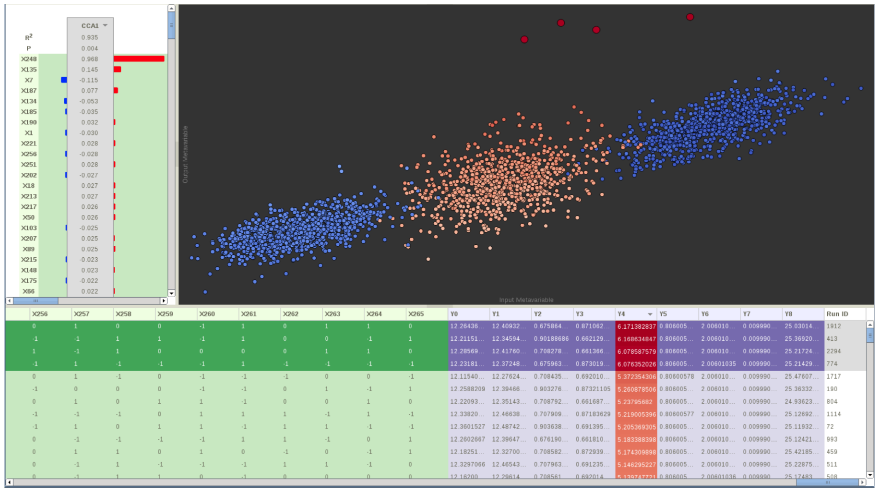

**Fig. 5** In the first CCA component of the all-to-all analysis of the large ensemble, anomalous runs (in red) are highlighted near the top of the scatterplot. We initially noticed them based on their position. The vertical position indicates that the difference between these simulations and the others is based on one of the outputs. Color-coding, combined with sorting the table, shows that these four simulations have much higher values in Y4 than any of the other simulations.

Given the large number of input variables, our initial analysis goal is to reduce the number of variables that we are using to drive the simulations. Our analysis uses a process like the one described for the small ensemble to evaluate which inputs are driving which outputs. Slycat easily scales to handle the increased variables and additional simulation runs (our largest ensemble to date has been about 500,000 runs). We are able to reduce the number of input variables from 266 down to 21. Unfortunately, due to space considerations, we are unable to provide the details of this analysis result. Instead, we return to the list in Sect. 1, and demonstrate how Slycat can be used to answer the third question : Which simulations are anomalous, and in what ways do they differ?

In the all-to-all analysis done as part of the variable reduction, we notice four anomalous points in the upper part of the scatterplot, shown in red in Fig. 5. Distance off of the diagonal is a metric for how well the linear correlation found by CCA describes a particular simulation. These four runs visually stand out immediately. What sets them apart? Since the Y-axis in the scatterplot is a metavariable based on the simulation outputs, vertical placement is a function of output variable values. Color-coding by the various outputs, we discover that Y4 values for these four simulations are at the high end of the scale. Sorting the Y4 values in the table, we find that these four simulations have Y4 values that are distinctly higher than any of the others (notice that they are in red, while the next largest values are in orange).

Next we investigate what is common amongst the four simulations' inputs, because presumably there is a common factor that is leading to these higher responses. We perform a many-to-one CCA analysis between all of the inputs and Y4. The top

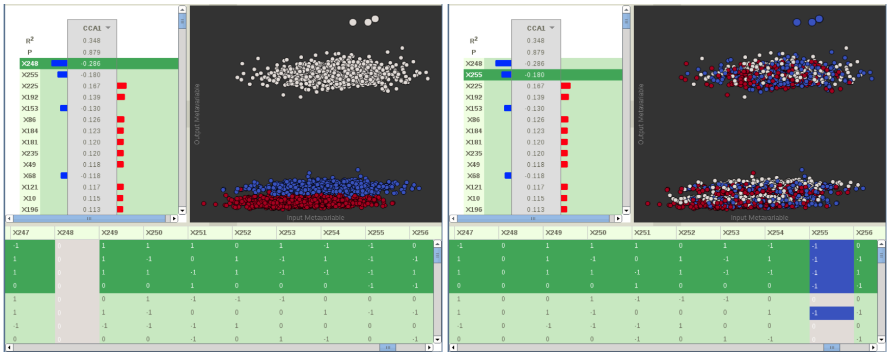

**Fig. 6** We perform an all-to-one analysis for Y4 to discover which input variables are causing the four anomalous output values. The top two variables, X248 (left image) and X255 (right image), have identical input values for the four anomalous simulation runs, but neither is the sole driver, since both demonstrate the same inputs driving a variety of outputs. The source of the anomalous outputs must stem from a combination of inputs acting in concert.

two input variables, X248 and X255, both have identical values for all four simulations. However, each of these variables provides a range of other Y4 responses for the same input values, as seen in Fig. 6, so neither variable in isolation is the cause. Using the table, we find nine variables that share identical values for all four simulations (X248, X255, X224, X175, X176, X187, X196, X213, and X229). Given X248's strong correlation with Y4, it is definitely involved. We will need to do further testing to isolate which additional input variables are involved. However, Slycat has enabled us to narrow down the possibilities to just a handful of variables.

## 7 Conclusions and Future Work

We have met our design goals for Slycat and demonstrated through two electrical circuit analysis examples of varying scale how the system can be used to answer the required three analysis questions in Sect. 1. Because our users also produce time-varying plots of voltage and current waveforms as part of these analyses, our next task will be to create an analysis type to handle time-varying data.

**Acknowledgements** The authors wish to thank Ben Long at Sandia National Laboratories for sharing his work processes, expertise, and data. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

## References


Abdellatif AS, El-Rouby A, Abdelhalim MB, Khalil A (2010) Hybrid latin hypercube designs. In: Informatics and Systems (INFOS), 2010 The 7th International Conference on, pp 1–5

Adams BM, Ebeida MS, Eldred MS, Jakeman JD, Swiler LP, Bohnhoff WJ, Dalbey KR, Eddy JP, Hu KT, Vigil DM, Bauman LE, Hough PD (2013) Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 5.3.1 users manual. Tech. Rep. SAND2010-2183, Sandia National Laboratories

Anderson TW (2003) An Introduction to Multivariate Statistical Analysis,3rd Edition. Wiley, New York

Bruckner S, Moller T (2010) Result-driven exploration of simulation parameter spaces for visual effects design. Visualization and Computer Graphics, IEEE Transactions on 16(6):1468–1476

Coffey D, Lin CL, Erdman A, Keefe D (2013) Design by dragging: An interface for creative forward and inverse design with simulation ensembles. Visualization and Computer Graphics, IEEE Transactions on 19(12):2783–2791

Computational Engineering International I (2013) URL http://www.ceisoftware.com/

Degani A, Shafto M, Olson L (2006) Diagram 2006, Lecture Notes in Artificial Intelligence 4045, Springer-Verlag, chap Canonical Correlation Analysis: Use of composite heliographs for representing multiple patterns., pp 93–97

Ge H, Kirsteins I, Wang X (2009) Does canonical correlation analysis provide reliable information on data correlation in array processing? In: Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on

Helton J (2008) Uncertainty and sensitivity analysis for models of complex systems. In: Graziani F (ed) Computational Methods in Transport: Verification and Validation, Lecture Notes in Computational Science and Engineering, vol 62, Springer Berlin Heidelberg, pp 207–228

Hotelling H (1936) Relations between two sets of variates. Biometrika 28(3/4)

Hummel M, Obermaier H, Garth C, Joy K (2013) Comparative visual analysis of lagrangian transport in cfd ensembles. Visualization and Computer Graphics, IEEE Transactions on 19(12):2743–2752

Karhunen J, Hao T, Ylipaavalniemi J (2013) Finding dependent and independent components from related data sets: A generalized canonical correlation analysis based method. Neurocomputing 113(0):153 – 167

Kitware I (2013) URL http://www.paraview.org/

Krzanowski WJ (1988) Principles of Multivariate Analysis. A User's Perspective. Oxford University Press, London

Marzban C (2013) Model tuning with canonical correlation analysis. Mon Wea Rev (Conditionally Accepted) URL http://faculty.washington.edu/marzban/cca.pdf

Matkovic K, Gracanin D, Jelovic M, Ammer A, Lez A, Hauser H (2010) Interactive visual analysis of multiple simulation runs using the simulation model view: Understanding and tuning of an electronic unit injector. Visualization and Computer Graphics, IEEE Transactions on 16(6):1449–1457

Mills K, Filliben J, Dabrowski C (2011) An efficient sensitivity analysis method for large cloud simulations. In: Cloud Computing (CLOUD), 2011 IEEE International Conference on, pp 724–731

Moreland K (2009) Diverging color maps for scientific visualization. In: Advances in Visual Computing, Springer Berlin Heidelberg, vol 5876, pp 92–103

Piringer H, Pajer S, Berger W, Teichmann H (2012) Comparative visual analysis of 2d function ensembles. Comp Graph Forum 31(3pt3):1195–1204

Potter K, Wilson A, Bremer PT, Williams D, Doutriaux C, Pascucci V, Johhson C (2009a) Visualization of uncertainty and ensemble data: Exploration of climate modeling and weather forecast data with integrated visus-cdat systems. Journal of Physics: Conference Series 180(1):012,089

Potter K, Wilson A, Bremer PT, Williams D, Doutriaux C, Pascucci V, Johnson C (2009b) Ensemble-vis: A framework for the statistical visualization of ensemble data. In: Data Mining Workshops, 2009. ICDMW '09. IEEE International Conference on, pp 233–240

Sanyal J, Zhang S, Dyer J, Mercer A, Amburn P, Moorhead R (2010) Noodles: A tool for visualization of numerical weather model ensemble uncertainty. Visualization and Computer Graphics, IEEE Transactions on 16(6):1421–1430

Shearer N, Khasawneh M, Zhang J, Bowling S, Rabadi G (2010) Sensitivity analysis of a large-scale system dynamics immigration model. In: Systems and Information Engineering Design Symposium (SIEDS), 2010 IEEE, pp 78–81

Smith KM, Banks DC, Druckman N, Beason K, Hussaini MY (2006) Clustered ensemble averaging: A technique for visualizing qualitative features of stochastic simulations. J of Computational and Theoretical Nanoscience 3(5):752–760

Song X, Zhao G (2012) Sensitivity analysis for a forest growth model: A statistical and time-dependent point of view. In: Plant Growth Modeling, Simulation, Visualization and Applications (PMA), 2012 IEEE 4th International Symposium

Steed CA, Ricciuto DM, Shipman G, Smith B, Thornton PE, Wang D, Shi X, Williams DN (2013) Big data visual analytics for exploratory earth system simulation analysis. Comput Geosci 61:71–82

Sukharev J, Wang C, Ma KL, Wittenberg A (2009) Correlation study of time-varying multivariate climate data sets. In: Visualization Symposium, 2009. PacificVis '09. IEEE Pacific, pp 161–168

Thornquist HK, Keiter ER, Rajamanickam S (2013) Electrical modeling and simulation for stockpile stewardship. XRDS 19(3):18–22

Wilson AT, Potter KC (2009) Toward visual analysis of ensemble data sets. In: Proceedings of the 2009 Workshop on Ultrascale Visualization, ACM, New York, NY, USA, UltraVis '09, pp 48–53