

Exploration of High-Dimensional Scalar Function for Nuclear Reactor Safety Analysis and Visualization

Dan Maljovec, Bei Wang, Valerio Pascucci

Scientific Computing and Imaging Institute, University of Utah
{maljovec,beiwang,pascucci}@sci.utah.edu

Peer-Timo Bremer

Lawrence Livermore National Laboratory
bremer5@llnl.gov

Michael Pernice, Diego Mandelli, Robert Nourgaliev

Idaho National Laboratory
{michael.pernice,diego.mandelli,robert.nourgaliev}@inl.gov

ABSTRACT

The next generation of methodologies for nuclear reactor Probabilistic Risk Assessment (PRA) explicitly accounts for the time element in modeling the probabilistic system evolution and uses numerical simulation tools to account for possible dependencies between failure events. The Monte-Carlo (MC) and the Dynamic Event Tree (DET) approaches belong to this new class of dynamic PRA methodologies. A challenge of dynamic PRA algorithms is the large amount of data they produce which may be difficult to visualize and analyze in order to extract useful information. We present a software tool that is designed to address these goals. We model a large-scale nuclear simulation dataset as a high-dimensional scalar function defined over a discrete sample of the domain. First, we provide structural analysis of such a function at multiple scales and provide insight into the relationship between the input parameters and the output. Second, we enable exploratory analysis for users, where we help the users to differentiate features from noise through multi-scale analysis on an interactive platform, based on domain knowledge and data characterization. Our analysis is performed by exploiting the topological and geometric properties of the domain, building statistical models based on its topological segmentations and providing interactive visual interfaces to facilitate such explorations. We provide a user's guide to our software tool by highlighting its analysis and visualization capabilities, along with a use case involving data from a nuclear reactor safety simulation.

Key Words: high-dimensional data analysis, computational topology, nuclear reactor safety analysis, visualization

1 INTRODUCTION

Dynamic Probabilistic Risk Assessment (PRA) methodologies [6] couple numerical simulation tools and time-dependent stochastic models (i.e., probabilistic failure models or parameter uncertainties), to perform system safety analysis. Widely used dynamic PRA methodologies are based on Monte-Carlo [16] or Dynamic Event Tree algorithms [11]. The common underlying idea is to run a large number of simulations (by employing system simulators) where values of system stochastic parameters (e.g., timing of failure of a specific component or an uncertain parameter) are sampled from their own distribution at each run. This type of PRA analysis can be very time-consuming when a large number of stochastic parameters are considered and when large and complex system simulators are used. Moreover, a large volume of data is typically

generated. Such large amounts of information can be difficult to organize for extracting useful information. Furthermore, it is often not sufficient to merely calculate a quantitative value for the risk and its associated uncertainties. The development of risk insights that can increase system safety and improve system performance requires the interpretation of scenario evolutions and the principal characteristics of the events that contribute to the risk.

The need for software tools able to both analyze and visualize large amount of data generated by Dynamic PRA methodologies has been emerging only in recent years. A first step has been shown in [18, 19] using clustering-based algorithms which focus more on the analysis part than the visualization side.

In this paper, we present a software tool that provides scientists and domain experts with an interactive analysis and visualization environment for understanding the structures of high-dimensional nuclear simulation data. Our tool adapts and extends the innovative framework called HDViz first proposed by Gerber et. al. [12] in exploring high-dimensional scalar functions, and applies the underlying techniques to nuclear reactor safety analysis and visualization. Our tool includes a host of various analysis and visualization capabilities. We describe each of these capabilities on a modular basis, by explaining the underpinning theories and presenting usage cases. The software segments the domain of a high-dimension function into regions of uniform gradient flow by decomposing the data based on its approximate Morse-Smale complex. Points belonging to a particular segment have similar geometric and topological properties, and from these we can create compact statistical summaries of each segment. Such summaries are then presented to the user in an intuitive manner that highlights features of the dataset which are otherwise hidden in a global view of the data. In addition, the visual interfaces provided by the system are highly interactive and tightly integrated, providing users with the ability to explore various aspects of the datasets for both analysis and visualization purposes.

2 PRELIMINARIES

We provide relevant background in Morse-Smale Complex, its approximation in high dimension and persistence simplification. We provide minimal technical definitions and illustrate the related concepts by examples whenever possible. We hope these intuitive examples are sufficient to convey a basic understanding of these concepts for non-specialists.

Morse-Smale Complex. The body of work presented in this paper relies heavily on the topological structure known as the Morse-Smale complex [3, 8, 9, 13, 14]. The Morse-Smale complex is based on Morse theory [20, 21]. Let \mathbb{M} be a smooth manifold embedded in \mathbb{R}^n without boundary and $f : \mathbb{M} \rightarrow \mathbb{R}$ be a smooth function with gradient ∇f . A point $x \in \mathbb{M}$ is called *critical* if $\nabla f(x) = 0$, otherwise it is *regular*. At any regular point x the gradient is well-defined and integrating it in both (ascending/descending) directions traces out an integral line, which is a maximal path whose tangent vectors agree with the gradient [9]. Each integral line begins and ends at critical points of f . Therefore, all regular points can trace their ascending integral line to a local maximum. Similarly, tracing the descending integral line of a regular point will associate a point with a local minimum. The *unstable/stable manifolds* (or *ascending/descending manifolds*) of a critical point p are defined as all the points whose integral lines start/end at p . The set of intersections of unstable and stable manifolds creates the *Morse-Smale complex* of f . Each cell (*crystal*) of the Morse-Smale complex is a union of integral lines that all share the same origin and the same destination. In other words, all the points inside a single crystal have uniform gradient flow behavior. These crystals yield a decomposition into monotonic, non-overlapping regions of the domain. Figure 1 illustrates these concepts associated with a height function defined on a 2-dimensional domain. For such a two-dimensional function, critical points




Figure 1. Top: (a) unstable manifolds, (b) stable manifolds and (c) Morse-Smale complex. Bottom: focused views of (d) an unstable manifold surrounding a local maximum x , (e) a stable manifold surrounding a local minimum y and (f) a Morse-Smale crystal associated with the maximum-minimum pair (x, y) .

include local maxima (red), local minima (blue) and saddles (green). (a) illustrates the unstable manifolds with a focused view surrounding a (local) maximum x in (d). (b) illustrates the stable manifolds with a focused view surrounding a (local) minimum y in (e). (c) illustrates the Morse-Smale complex created by the intersection of unstable and stable manifolds, and (f) focuses on a Morse-Smale crystal associated with the maximum-minimum pair (x, y) . Each partition is colored by its corresponding gradient behavior, where points share the same color if: (d) their ascending gradient flow (white arrow) end at the same local maximum; (e) their descending gradient flow (white arrow) end at the same local minimum and (f) their ascending and descending gradient flow (white bidirectional arrow) correspond to the same maximum-minimum pair.

Approximating the Morse-Smale Complex in High Dimension. Suppose our input domain is a finite set of points \mathbb{X} in \mathbb{R}^n sampled from some unknown n -manifold. To approximate the Morse-Smale complex of a high-dimensional scalar function f defined on \mathbb{X} , our first task is to estimate the gradient at the input points, \mathbb{X} . First, we compute a neighborhood graph such as the k -nearest neighbor graph of \mathbb{X} (that is, there is an undirected edge between points $x, y \in \mathbb{X}$ if x is among the k -nearest neighbors of y or y is among the k -nearest neighbors of x). At each point in \mathbb{X} , we choose the steepest ascending/descending edge to represent the gradient. With this gradient approximation, we can determine the local extrema by labeling all points with no neighbors of higher values as local maximum and all points with no neighbors of lower values as local minimum. We then label all points in \mathbb{X} according to the local extrema at which its ascending and descending gradients terminate. Subsequently, we collect all vertices with the same pair of labels into crystals and add the extrema to all crystals that share the corresponding label. Therefore the set of sampled points \mathbb{X} are then clustered into crystals \mathbb{X}_i (where $\bigcup \mathbb{X}_i = \mathbb{X}$), with function values observed as $\mathbb{Y}_i = f(\mathbb{X}_i)$. These crystals then serve as an approximation of the Morse-Smale complex, as first introduced in [12]. Some research is underway [5, 17] to understand how different neighborhood graphs (i.e. empty region graphs [4]) may impact our approximations, and what sampling conditions we should impose on

the data to guarantee the approximation quality. Figure 2 illustrates how our approximation works for a 2-dimensional height function where the function values are only given at a finite set of sampled points in the domain. The graph shown is part of a k -nearest neighbor graph where $k = 4$. Point x (y) is a local maximum (local minimum) since it has no neighbors in the graph with higher (lower) function values. By following the approximated gradients, all the cyan points belong to the same Morse-Smale crystal since they share the same pairs of maximum-minimum labels, e.g. (x, y) .

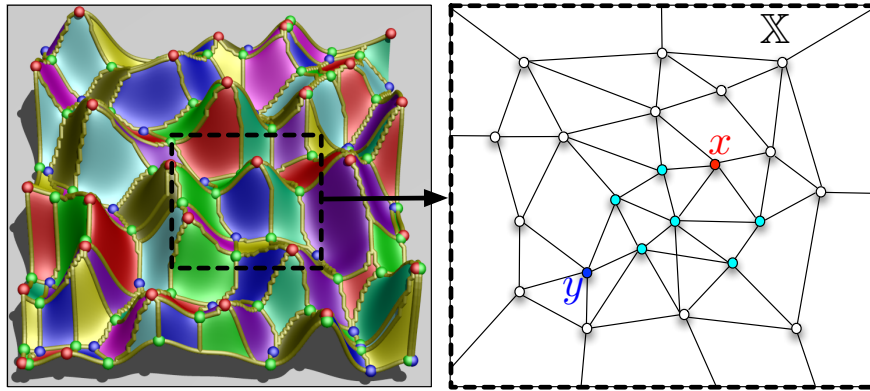


Figure 2. Left: the Morse-Smale complex of a 2-dimensional height function f . Right: the approximation of a Morse-Smale crystal of f , where the values of f are only given at a finite set of sampled points in the 2-dimensional domain.

Persistence Simplification. In real datasets, there is often noise which may manifest itself as small topological artifacts, either spurious extrema that may not truly exist in the data or small features the user does not deem relevant. To account for this and allow the user to select a *scale* (complexity level) appropriate for the specified dataset, we use the notion of *persistence simplification* whereby less salient features are merged with neighboring, more significant features. In the case of the Morse-Smale complex, we assign a persistence to each critical point in the complex which intuitively describes the scale at which a critical point would disappear through simplification. Here we omit the technical definition of persistence and refer the curious readers to [7, 10]. We illustrate the idea of persistence through the following example, as shown




Figure 3. A 2-dimensional Morse-Smale complex before (a) and after (b) persistence simplification. Maxima are red, minima are blue and saddles are green.

in Figure 3 for a 2-dimensional height function. As usual, red, blue and green points denote local maxima, local minima and saddles respectively. In (a), the left peak at the local maximum x is considered less important topologically than its nearby peak at local maximum z , since x is lower. Therefore, at a certain scale, we would like to represent this feature as a single peak instead of two separate peaks, as shown in (b), by redirecting ascending gradient flow (white arrow) that originally terminates at x to terminate at z . In this way, we simplify the function by removing (canceling) the local maximum x with its nearby saddle y . Such a simplification procedure is referred to as the persistence simplification.

A Complete Illustrative Example. Finally for a complete intuitive understanding of the above concepts, we give the following example, as shown in Figure 4. The top of the figure illustrates its corresponding unstable manifolds, stable manifolds and Morse-Smale complex, while the middle of the figure showcases the progressive simplification of its critical points: at the finest level there are four maxima in the data, but as we increase the scale, neighboring topological features are merged where the upward gradient flow is directed to its more salient neighboring maxima. Circled points correspond to the critical points that are removed at each simplification step. We thus build a filtration of segmentations where Morse-Smale crystals are merged based on the persistence value of their associated extrema.

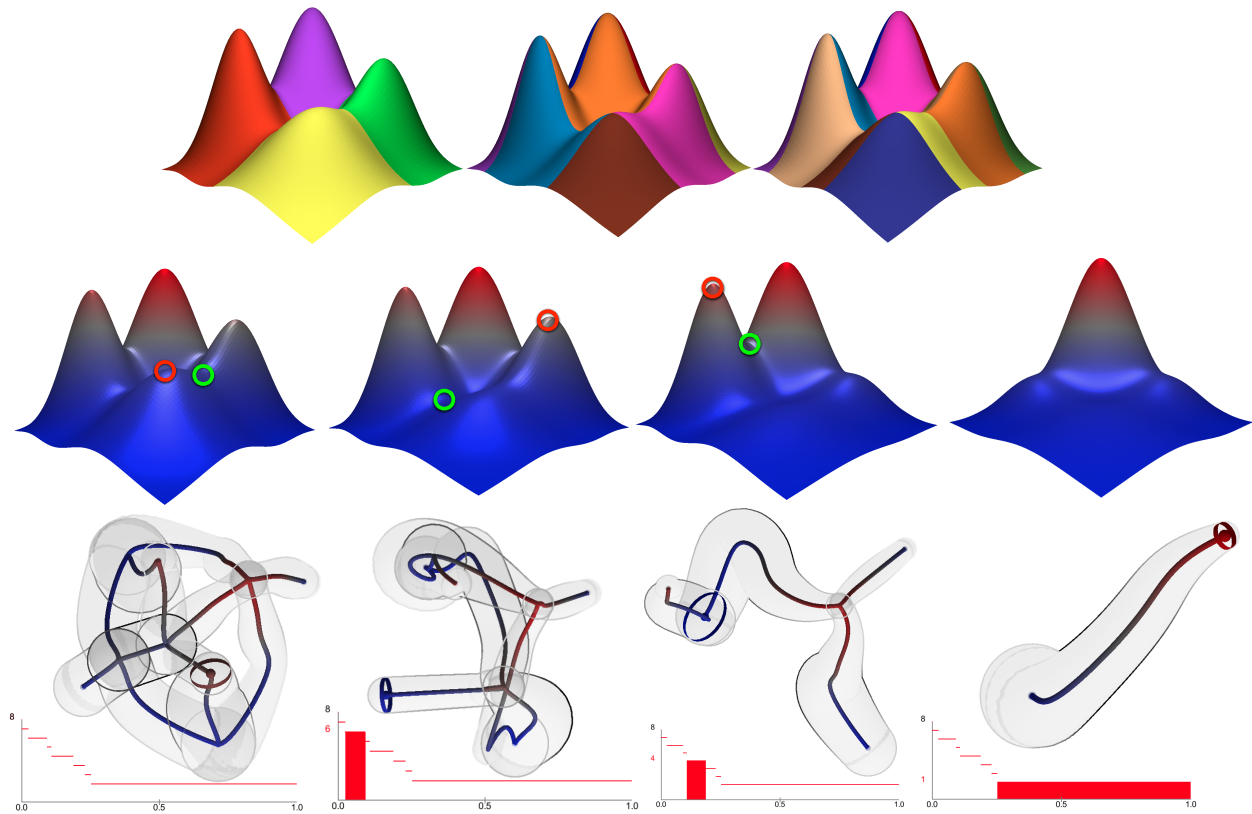


Figure 4. An illustrative example based on a 2-dimensional height function. **Top:** unstable manifolds, stable manifolds and Morse-Smale complex. **Middle:** progressive simplification of its critical points. **Bottom:** topological summaries shown at various scales (see Section 3).

3 ANALYSIS AND VISUALIZATION MODULES

In this section, we describe each analysis and visualization module within our integrated system that are either part of the original capabilities provided by HDViz [12] or part of our extension. We demonstrate our infrastructure with a 6-dimensional example dataset from nuclear plant safety analysis. The data is extracted from a VR_2^+ nuclear reactor simulator and represents an ensemble of 10000 simulation trials where a SCRAM is simulated due to a failure in the system. A SCRAM event is when the control rods of the reactor are inserted into the core in order to prevent overheating of the reactor core. The output variable is the peak coolant temperature (PCT*), measured in Kelvin. The domain scientists are interested in what

*Note that PCT used here does not stand for peak clad temperature.

combination of conditions (in the form of input parameters) can cause potential reactor failure (i.e. nuclear meltdown witnessed by PCT exceeding a threshold value). The input space is defined by six parameters:

- **PumpTripPre** - the minimum pressure (MPa) in the heat exchange pump causing the SCRAM to trip
- **PumpStopTime** - the relaxation time (sec) of pump’s phase-out
- **PumpPow** - end power of the pump
- **SCRAMtemp** - the maximum temperature in the system causing the SCRAM to trip
- **CRinject** - the control rod position at the end of SCRAM
- **CRtime** - the relaxation time (sec) of the control rod system.

3.1 Topological Summary

The visual interface designed for topological summary is inherited and extended from the capabilities provided by HDViz [12]. Considered as the main visual display of our software, this interface summarizes each Morse-Smale crystal into a 1D curve in high-dimensional space which is then projected onto a viewable 3D space. The interface encodes three steps (detailed in [12], with a high level description below), to arrive at a 3D representation for analysis and visualization of the d -dimensional scalar function f , defined on a set of sampled points \mathbb{X} . (1) *Morse-Smale approximation*: We approximate the Morse-Smale crystals, \mathbb{X}_i and $\mathbb{Y}_i = f(\mathbb{X}_i)$, in high dimension using an approximate k -nearest neighbor graph [1]. (2) *Geometric summaries*: For each crystal of the Morse-Smale complex, since each point has similar approximated monotonic gradient behavior, a geometric summary is constructed by an inverse regression, yielding a 1D curve r_i in the d -dimensional domain of f . Intuitively, $r_i(y)$ (where $y \in \mathbb{Y}_i$) yields a representation of the crystal as the average of the function values of level sets within the crystal, see [12] for its detailed derivation. (3) *Dimension reduction*: The set of regression curves can be represented by a graph embedded in \mathbb{R}^d with each edge corresponding to a curve and vertices corresponding to extremal points. For visualization, we embed this graph into 2D preserving the spatial relation among extrema and the geometry of the crystals that connect them. First, vertices are embedded using PCA or ISOMAP [22]. Second, edges are embedded individually through their first two principle components. And third, the resulting 2D curves are attached to the projected vertices through affine transformations. The third dimension is reserved for the output parameter.

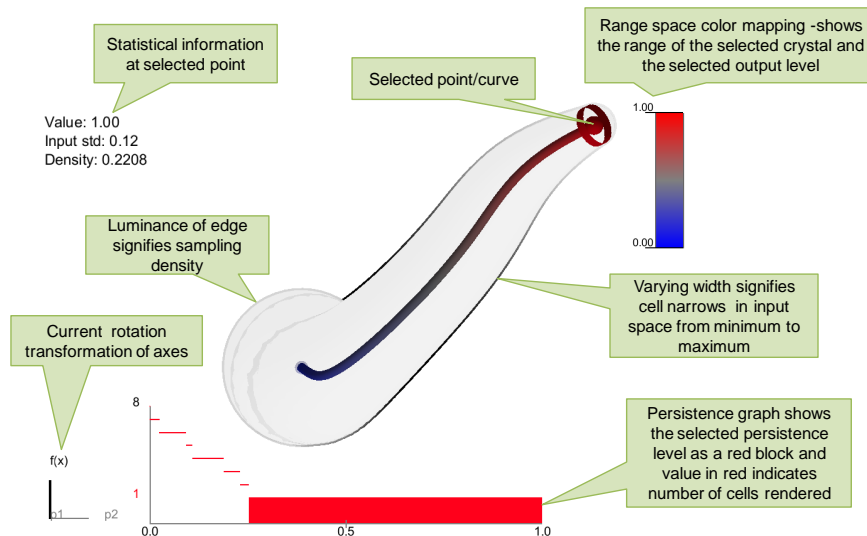


Figure 5. The topological summary visual interface of a simple 2D function.

Visual Components. The visual components of the above framework are shown in Figure 5 for a simple 2D function with one crystal (example revisited from Figure 4). Each curve is encased in a transparent tube

where the width of the tube represents the “spread” of the data at a particular scale, and the luminance of the tube encodes the density of data points within each crystal. Users are given the flexibility to view the topological summary of a high-dimensional function by switching between PCA and ISOMAP projections, and using affine transformations to manipulate the projection directly on the screen. In order to preserve the “width” of a crystal at a given scale, we compute the standard deviation with respect to each input parameter and also a single *average* standard deviation across all input parameters. The latter is direction-independent and can therefore be used as a generalized width of the data at a particular output level. The radius of the outer transparent tubes are defined by this direction-independent standard deviation. The last visual cue is the darkness of the edge of the transparent tube. Where the sampling density is high, the outline of the tube is drawn black and as the sampling density decreases, the luminance of this edge increases. To enable multi-scale analysis, we use a modified version of the persistence diagram [10], referred to as the *persistence graph*. This is shown as a visual component at the bottom of Figure 5. It shows the number of Morse-Smale crystals (y -axis) as a function of scale (i.e. x -axis, persistence threshold normalized by the range of the dataset). A selected scale is drawn with a red box and a corresponding number of Morse-Smale crystals is displayed along the y -axis in red. Stable features are considered as those that exist over a large range of scales (i.e. a sequence of persistence simplification with increasing scales), which correspond to long horizontal lines in the persistence graph.

We demonstrate how the topological summary visual interface allows the users to explore the data at multiple scales. We revisit our simple dataset under several levels of persistence simplification in Figure 4 bottom. Here the leftmost image shows the full resolution of data with four local maxima, while each subsequent image reduces the number of maxima by one until we are left with a single crystal describing the gradient flow from the global minimum to the global maximum. The numbers in red, shown in the persistence graphs, indicate the total number of crystals displayed, from left to right, as 8, 6, 4 and 1. Instead of giving the users a representation of the data at a fixed scale, we provide an interactive platform to help them differentiate features from noise through multi-scale analysis and to choose the appropriate scale based on domain knowledge and data characterization.

A Simple Example of Geometric Summary. Figure 6 shows several examples of simple 2D surfaces represented using the techniques described in this section. Each function has one local maximum and a single local minimum surrounded by a flat area. Using an evenly sampled grid, the density of points available at each level is mapped to the color on the edge of the transparent tubes. The widths of the tubes vary with respect to the spread of the data at each level set. Note how the sharp spike in (a) is represented in the topological display. The steep peak covers a small area and the small width of the transparent tube demonstrates this in the area near the maximum. (b) has a wider area surrounding the peak and while its representing curve remains a single straight line, the tube surrounding it changes behavior to account for the width of each level set. The plateau function in (c) even maps the discontinuity in the curve by making a sudden jump from a low value to a high value. Note how the edge of the tube has high luminance in the middle section denoting a lack of data used to compute this section, whereas the most densely sampled region is at the maximum value which has a black outline.

6D Demo Dataset. As a final example, we illustrate our interface with our 6D demo example in Figure 7 under multiple scales. This dataset contains 753 individual crystals at the finest level, though due to low persistence and visual clutter caused by these crystals, we restrict our analysis to only a handful of the most salient crystals. During the persistence simplification steps, the topological summaries consist of 6, 3 and 1 Morse-Smale crystal(s), respectively. At each of the shown persistence levels, the dataset is characterized by a single global minimum with high sampling density. The widths of all the crystals at the minimum is expansive compared to the widths at the maxima. We could infer from this analysis result that most of the




Figure 6. Three simple surfaces represented using geometric summary tubes.

data points (simulations) represent lower PCTs and the conditions to reach these lower PCTs varies widely in the domain space. On the other hand, the maximum PCTs are reached at more specific input ranges which is made clear through the narrow widths of the tubes (of one standard deviation) surrounding the red maxima.




Figure 7. The topological summary of our demo dataset shown at various scales.

3.2 Statistical Summary

The visual interface shown in Figure 8 demonstrates statistical and geometric (i.e. gradient) information associated with the selected point in the topological summary interface. Each input dimension, or coordinate, is viewed as an inverse function of the output parameter. In the left column of the statistical summary window, each horizontal axis describes the range of values of each input parameter and the coordinate mean and coordinate standard deviation associated with the selected point. The right column encodes the gradient information, that is, the change in the output with respect to the change in each input parameter. In Figure 8, we can see that three parameters have quite large standard deviations, PumpStopTime, SCRAMtemp, and CRtime, whereas the other parameters vary much less. The topological summary interface on the left of Figure 8 is a global summary of the data that presents the user with a transparent tube of moderate width that approximate such kind of detail. The statistical summary and the remaining display methods are integrated within the topological summary to provide deeper meanings to the analyst.

3.3 Inverse Coordinate Plots

In the inverse coordinate plots, each input parameter is considered as a one-dimensional function of the output variable. This visual interface is shown in Figure 9 for our 6D example. On the left, at the persistence level with three crystals, the interface displays the inverse coordinate plot for data points associated with the




Figure 8. A snapshot of the statistical summary visual interface with highlighted visual components.

selected Morse-Smale crystal in the topological summary interface. On the right, at the persistence level with six crystals, the interface shows combined inverse coordinate plots associated with all the crystals. For the selected crystal(s), the regression curve is drawn in the inverse coordinate plots with a grey tube representing the parameter-specific standard deviation.

From the right image in Figure 9, we see the top set of axes where each of the six regression curves is readily distinguishable with respect to different ranges of values for PumpTripPre, whereas in the lower plots they vary little toward the left and only slightly at the right. On the other hand, parameters such as PumpPow and CRinject result in high temperatures only at specific levels. This conclusion is based upon the tight configuration of points resulting in high PCT and it is also supported by the consistent locations of the mean values across all crystals and the low standard deviations of these parameters.

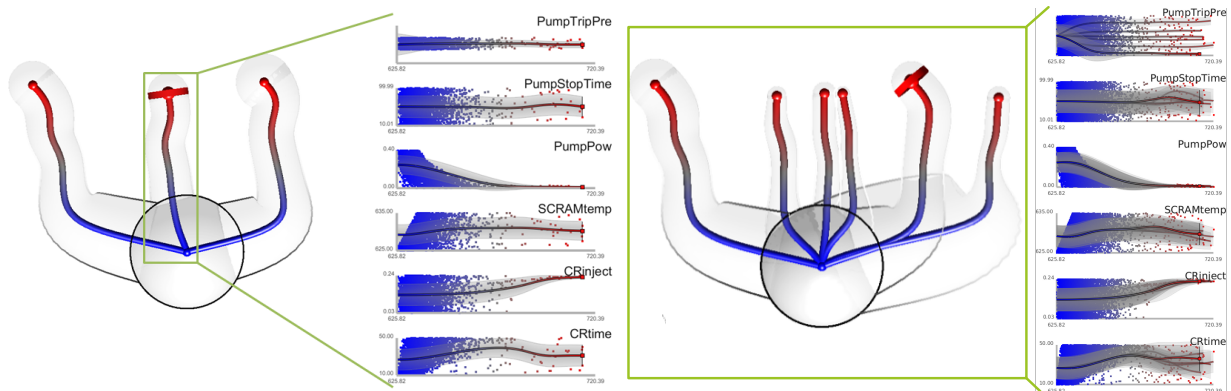


Figure 9. Left: Inverse coordinate plots shown for the highlighted crystal for the three crystal case. Right: Inverse coordinate plots shown for the six crystal case of the same 6D data.

3.4 Interactive Projection, Parallel Coordinate Plots and Pairwise Scatter Plots

As the number of dimensions increases for a high-dimensional function defined on the sampled point set, selecting the interesting projection dimensions while interacting with the high-dimensional function could become counter-intuitive. We design an interactive visualization interface that provides simple and fully explanatory pictures that give comprehensive insights into the global structure of the high-dimensional function. Based on hypervolume visualization techniques developed in [2], our basic idea is a generalization of direct parallel projection methods. We first create an independent viewing system that scales with the number of dimensions where the user is allowed to manipulate how each axis is projected. We then apply these manipulations to project the geometric summaries from the functional space to the screen space [2].

Users are provided with a wheel of labeled axes which they can manipulate by stretching, contracting, and rotating. The geometry of the inverse regression curves therefore is meaningfully preserved. With a single manipulated projection displayed at a time, the interactive nature of such a tool provides the user with intuition that is otherwise lost in a PCA or ISOMAP projection. An example for the 6D demo dataset is shown in Figure 10. Currently, this system only supports a view of the geometric summaries, but a possible extension is to create a full hypervolume visualization of the raw high-dimensional data points as proposed in [2].




Figure 10. Interactive Projection Visual Interface for the 6D demo example (3-crystal case). Left: projection of the high-dimensional inverse regression curves and their associated standard deviation tubes, based on the mapping shown in the right. Right: each input dimension is mapped to a single green line segment, which the user can manipulate by stretching and rotation, to emphasize or diminish the effect of a particular dimension on the projection.

In addition, we also use parallel coordinate [15] plots and pairwise scatter plots to illustrate the correlations among input parameters in the high dimensional datasets. These plots (figures omitted here) offer more exploratory values for the end-users.

4 CONCLUSIONS

In this paper, we have presented a software tool suitable to analyze and visualize datasets generated by safety analysis codes. In particular, our software could be interfaced with Dynamic PRA algorithms. In such a configuration, a large number of runs of the safety analysis code are performed, where during each run, initial conditions and timing/sequencing of events are changed accordingly to their statistical distribution. For large complex systems, this type of code generates a large amount of data, modeled as high-dimensional

scalar functions. Our software is designed to assist the user in the process of extracting useful information from such functions.

This extraction is performed by exploring the correlations between uncertain input parameters and simulation outcome. These correlations are modeled as a Morse-Smale complex and then summarized in various visual display modules, where the reconstructed simulation outcome is directly linked to the input parameters. From an uncertainty quantification point of view, such an analysis tool allows the user to identify and consequently rank variables based on their correlations with simulation outcome (e.g., maximum fuel temperature). Such software becomes even more effective for the analysis of complex systems such as nuclear power plants where the number of variables is very large.

We have described each analysis and visualization module in the system, by explaining the theories and techniques based on a specific application of the software in a 6-dimensional demo example involving nuclear safety simulation. Our example is based on code performed on 10000 simulations of a simplified pressurized water reactor (PWR) system for a SCRAM scenario. Our software allows us to identify and visualize the correlations among six parameters and the maximum coolant temperature. We first perform topological analysis to identify the underlying topological structure of the dataset. Statistical information is then summarized and linked to the topological structures extracted from the data, allowing the users to identify correlations between timing/sequencing of events and simulation outcome such as core temperature. We have obtained some intuitive understanding of the structures of such a high-dimensional function, although a more comprehensive interpretation of our results depends on extensive testing and explorations from our collaborators. We expect to work closely with the domain scientists and obtain feedback from the end users regarding (a) the interpretation of the testing datasets using our software, (b) the limitations and potential improvements of current analysis techniques, and (c) the potential improvements of the visualization interfaces in terms of usability and interactivity.

ACKNOWLEDGEMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and was funded by the Uncertainty Quantification Strategic Initiative LDRD Project at LLNL, Tracking code 10-SI-023.LLNL-CONF-625957. This work was also supported by Idaho National Laboratory LDRD award number 00115847A1DEAC0705ID14517 and Enabling Transformational Science grant UTA09000731.

REFERENCES

- [1] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM*, 45:891–923, 1998.
- [2] C. L. Bajaj, V. Pascucci, G. Rabbio, and D. Schikorc. Hypervolume visualization: a challenge in simplicity. *Proceedings IEEE Symposium on Volume Visualization*, pages 95–102, 1998.
- [3] P.-T. Bremer, H. Edelsbrunner, B. Hamann, and V. Pascucci. A topological hierarchy for functions on triangulated surfaces. *IEEE Transactions of Visualization and Computer Graphics*, 10:385–396, 2004.
- [4] J. Cardinal, S. Collette, and S. Langerman. Empty region graphs. *Computational Geometry: Theory and Applications*, 42:183–195, 2009.

- [5] C. D. Correa and P. Lindstrom. Towards robust topology of sparsely sampled data. *IEEE Transactions on Visualization and Computer Graphics*, 17:1852–1861, 2011.
- [6] J. Devooght. Dynamic reliability. *Advances in Nuclear Science and Technology*, 25:215–278, 1997.
- [7] H. Edelsbrunner and J. Harer. Persistent homology - a survey. *Contemporary Mathematics*, 453:257–282, 2008.
- [8] H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci. Morse complexes for piecewise linear 3-manifolds. In *Proceedings 19th ACM Symposium on Computational Geometry*, pages 361–370, 2003.
- [9] H. Edelsbrunner, J. Harer, and A. J. Zomorodian. Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete and Computational Geometry*, 30(87-107), 2003.
- [10] H. Edelsbrunner, D. Letscher, and A. J. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.
- [11] E. M. G. Cojazzi, J. M. Izquierdo and M. S. Perea. The reliability and safety assessment of protection systems by the use of dynamic event trees: the dylam-treta package. In *In Proc. XVIII Annual Meeting Spanish Nucl. Soc.*, 1992.
- [12] S. Gerber, P.-T. Bremer, V. Pascucci, and R. Whitaker. Visual exploration of high dimensional scalar functions. *IEEE Transactions on Visualization and Computer Graphics*, 16:1271–1280, 2010.
- [13] A. Gyulassy and V. Natarajan. Topology-based simplification for feature extraction from 3d scalar fields. In *Proceedings IEEE Visualization*, pages 535–542, 2005.
- [14] A. Gyulassy, V. Natarajan, V. Pascucci, and B. Hamann. Efficient computation of morse-smale complexes for three-dimensional scalar functions. *IEEE Transactions on Visualization and Computer Graphics*, 13:1440–1447, 2007.
- [15] A. Inselberg. *Parallel Coordinates: VISUAL Multidimensional Geometry and its Applications*. Springer, 2009.
- [16] E. Z. M. Marseguerra, J. Devooght, and P. Labeau. A concept paper on dynamic reliability via monte carlo simulation. *Mathematics and Computers in Simulation*, 47:371–382, 1998.
- [17] D. Maljovec, A. Saha, P. Lindstrom, P.-T. Bremer, B. Wang, C. Correa, and V. Pascucci. A comparative study of morse complex approximation using different neighborhood graphs. *TopoInVis*, accepted, 2013.
- [18] D. Mandelli, A. Yilmaz, and T. Aldemir. Data processing methodologies applied to dynamic PRA: an overview. In *Proceeding of Probabilistic Safety Analysis (PSA)*, 2011.
- [19] D. Mandelli, A. Yilmaz, and T. Aldemir. Scenario analysis and pra: Overview and lessons learned. In *Proceedings of European Safety and Reliability Conference (ESREL 2011), Troyes (France)*, 2011.
- [20] J. Milnor. *Morse Theory*. Princeton University Press, New Jersey, NY, USA, 1963.
- [21] M. Morse. Relations between the critical points of a real function of n independent variables. *Transactions of the American Mathematical Society*, 27:345–396, 1925.
- [22] J. B. Tenenbaum, V. D. Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000.