

Adaptive Sampling Algorithms for Probabilistic Risk Assessment of Nuclear Simulations

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

Diego Mandelli

Idaho National Laboratory
diego.mandelli@inl.gov

ABSTRACT

Nuclear simulations are often computationally expensive, time-consuming, and high-dimensional with respect to the number of input parameters. Thus exploring the space of all possible simulation outcomes is infeasible using finite computing resources. During simulation-based probabilistic risk analysis, it is important to discover the relationship between a potentially large number of input parameters and the output of a simulation using as few simulation trials as possible. This is a typical context for performing adaptive sampling where a few observations are obtained from the simulation, a surrogate model is built to represent the simulation space, and new samples are selected based on the model constructed. The surrogate model is then updated based on the simulation results of the sampled points. In this way, we attempt to gain the most information possible with a small number of carefully selected sampled points, limiting the number of expensive trials needed to understand features of the simulation space.

We analyze the specific use case of identifying the limit surface, i.e., the boundaries in the simulation space between system failure and system success. In this study, we explore several techniques for adaptively sampling the parameter space in order to reconstruct the limit surface. We focus on several adaptive sampling schemes. First, we seek to learn a global model of the entire simulation space using prediction models or neighborhood graphs and extract the limit surface as an iso-surface of the global model. Second, we estimate the limit surface by sampling in the neighborhood of the current estimate based on topological segmentations obtained locally.

Our techniques draw inspirations from topological structure known as the Morse-Smale complex. We highlight the advantages and disadvantages of using a global prediction model versus local topological view of the simulation space, comparing several different strategies for adaptive sampling in both contexts. One of the most interesting models we propose attempt to marry the two by obtaining a coarse global representation using prediction models, and a detailed local representation based on topology. Our methods are validated on several analytical test functions as well as a small nuclear simulation dataset modeled after a simplified Pressurized Water Reactor.

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

1 Introduction

To study the reliability and safety of complex systems (e.g. nuclear power plants, airplanes, chemical plants), one typically performs a combination of event tree and fault tree analysis [1]. However, these methods are

characterized by the following disadvantages: (a) The timing of events is not explicitly modeled; (b) The ordering of events is preset by the analysts; (c) The modeling of complex accident scenarios is driven by expert judgment. For these reasons, there is currently an increasing interest in the development of dynamic Probabilistic Risk Assessment methodologies (DPRA) [2] since they can be used to address the deficiencies of the conventional methods listed above.

For DPRA, propagation of uncertainties in the analysis of complex systems employs sampling algorithms which perform a series of simulation runs given a set of uncertainty parameters. However, typically the set of uncertain parameters is very large and the computational cost of each run is very high. Consequently, the space of the possible solutions, the response surface, can be sampled only very sparsely and this precludes the ability to fully analyze the impact of uncertainties on the system. For safety analysis applications, the following points often emerge: (a) Many regions of the response surface are not of interest; (b) The set of parameters that are of safety concern is a small subset of the original set of uncertainty parameters.

Currently, the class of most used sampling algorithms include: classic stochastic sampling (e.g., Monte-Carlo [3,4], stratified sampling such as Latin hypercube sampling [5], importance sampling [6] and orthogonal arrays based [7] algorithms) and deterministic algorithms (e.g., Polynomial Chaos Expansions [8] and Quasi Monte-Carlo [9]). However, all of these sampling algorithms do not possess a sampling strategy that explicitly takes into account the results of previous simulations. Adaptive sampling algorithms, on the other hand, adapts a sampling strategy that chooses the next sample based on the results obtained by previous samples through a statistical learning models and, thus, focus sampling in risk sensitive regions such as boundaries between system safe and system failure: the limit surface.

Contributions. In this paper, we present three adaptive sampling models that could be used to infer the behavior of system by estimating the limit surface. All three models share a general adaptive sampling pipeline but differ from one another in terms of the candidate selection process. The first model relies on globally probing a prediction model to estimate the limit surface as an iso-surface of the global response surface. The second model is topology-based. It improves upon the first model by taking into account the topology of the prediction model based on its Morse-Smale complex structure, and samples points in local areas that are considered topologically relevant. The final model is designed to be completely data-driven. It computes a candidate set directly from the neighborhood graph constructed from the training data, without any dependencies on a particular prediction model. We discuss the advantages and limitations of our proposed models and demonstrate their fidelity on several synthetic datasets as well as a small nuclear simulation dataset modeled after a simplified Pressurized Water Reactor (PWR).

2 Background

One of our proposed adaptive sampling models draws inspiration from the topological structure known as the Morse-Smale complex. We review the relevant concepts here for non-specialists, including Morse-Smale complex, its approximation, and neighborhood graphs.

Morse-Smale Complex. Topological structures, such as contour trees [10], Reeb graphs [11,12] and Morse-Smale complexes [13–15] provide abstract representations for scalar functions defined on point clouds. These structures can be used to define a wide variety of features in various applications in medical [16], physics [17,18] and material science [19]. Other topological approaches have been proposed in the analysis and visualization of high-dimensional data, e.g. [20,21].




Figure 1: For a height function defined on a 2D domain (where maxima, minima and saddles are colored red, blue and green respectively): (a) For each descending manifold, the gradient flow (white arrow) ends at the same maxima; (b) For each ascending manifold, the gradient flow starts at the same minimum; (c) For each Morse-Smale crystal, the gradient flow begins and ends at the same maximum-minimum pair.

We propose a topology-driven adaptive sampling technique inspired by the Morse-Smale complex, which is based on the Morse complex and Morse theory [22]. Given a smooth function f defined on a smooth manifold embedded in the Euclidean space, a point in the domain is considered *regular* if its gradient is non-zero. At any regular point the gradient is well-defined and integrating it traces out an integral line which begins at local minimum and ends at local maximum of f . In a nutshell, based upon the line integral of the gradient flow at each regular point in the domain, the Morse complex of f describes the segmentation of the domain into subregions called *descending manifolds* around local maxima, and the Morse complex of $-f$ decomposes the domain into subregions called *ascending manifolds* surrounding local minima. The set of intersections of ascending and descending manifolds creates the *Morse-Smale complex* (MSC) of f . This is explained in Figure 1. Each *crystal* of the Morse-Smale complex consists of points in the domain whose gradient flow begins and ends at the same maxima-minima pair, that is, they share uniform gradient flow behavior. These crystals yield a decomposition of the domain into monotonic, non-overlapping regions.

Approximation of MSC. To approximate Morse-Smale complex in the discrete setting, especially in high dimension, we impose a neighborhood graph on the given set of points. At each point, we choose the steepest ascending/descending edge based on function values to represent the ascending/descending gradient flow directions. All points that have no ascending/descending gradient are considered local extrema. We label all vertices according to the local extrema pair its ascending/descending gradient terminates. Vertices with the same maximum-minimum pair then form an approximation of the Morse-Smale crystal [23].

Neighborhood Graphs. Many neighborhood graphs could be used to impose combinatorial structures on a given set of points as there are a number of ways to connect points to form a neighborhood. A k -nearest neighbor graph, for example, is based on distance between points; while alternatively, empty region graphs [24] take into account the inclusion of other points in regions defined for edges. We employ the

relaxed Gabriel graph [25] in the proposed work here, as it has certain desirable properties (such as relatively dense connectivities). The Gabriel graph [26] is a neighborhood graph on a set of points \mathbb{X} where two vertices p and q are joined by an edge if and only if the empty region test returns true, that is, the disk with diameter $\|p - q\|$ has no other points of \mathbb{X} in its interior. Such requirement is then relaxed by prescribing that only neighbors of p and q are considered in the empty region test. That is, if r lies in the circumscribing disk determined by p and q , and is not a neighbor of either p or q , then p and q are joined by an edge.

3 Adaptive Sampling Models

We introduce three different adaptive sampling models. They share a general pipeline but differ from one another in terms of the candidate selection process. In the first model M_1 , we rely on globally probing a prediction model (PM), e.g. a Gaussian process model (GPM), to estimate a candidate set that lies on the limit surface. The second model M_2 improves upon this technique by computing the topology of the PM and sampling in areas deemed interesting by the topology rather than querying the entire domain. The final model M_3 computes a candidate set directly from the training data without the use of a PM.

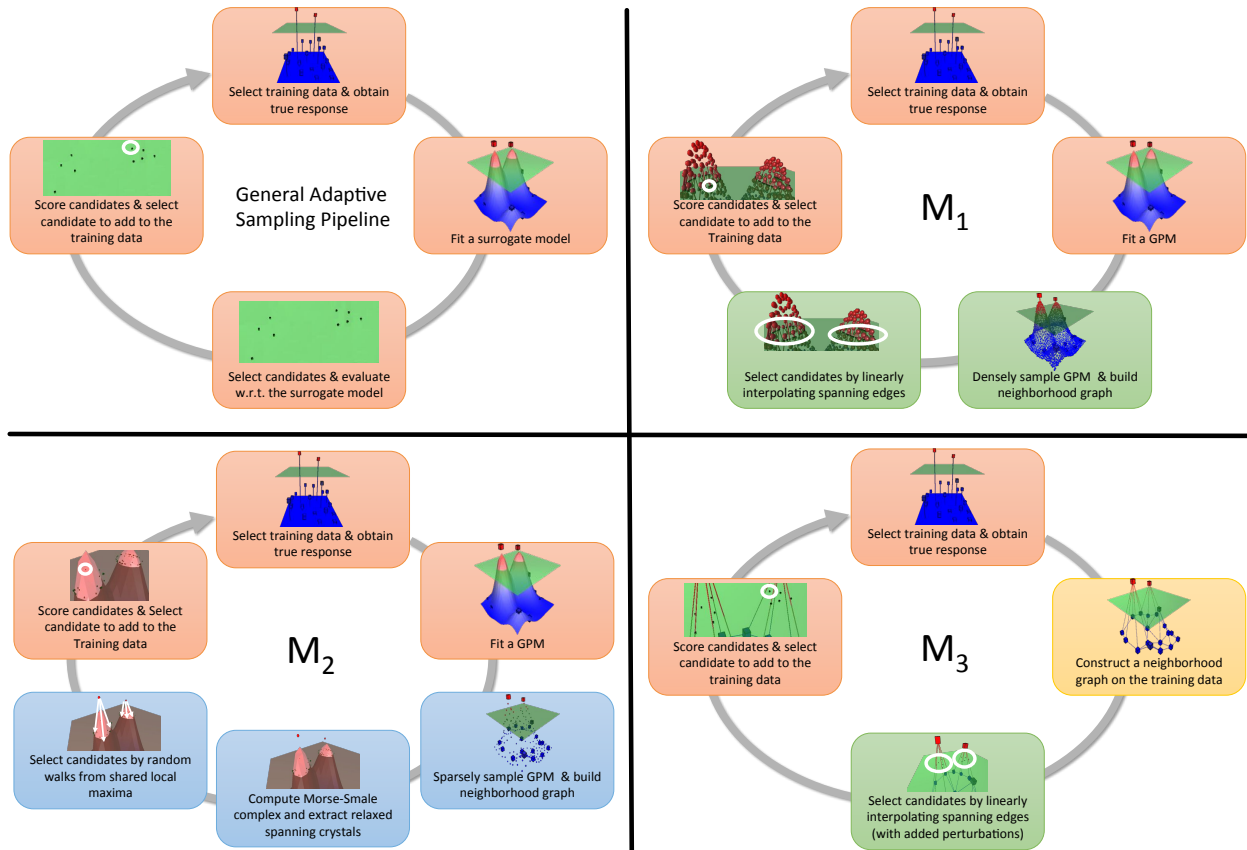


Figure 2: General adaptive sampling pipeline and pipelines for M_1 , M_2 and M_3 respectively.

The general adaptive sampling pipeline (Figure 2 top left) begins by selecting some initial training data, running the simulation and obtaining a collection of true responses at these data points. Second, it fits a response surface surrogate model from the initial set of training data. Third, a set of candidate points is chosen in the parameter space based on certain sampling technique, and the surrogate model is evaluated at these points, obtaining a set of approximated values. Fourth, each candidate point is assigned a score based on some adaptive sampling scoring function (usually derived from qualitative or quantitative relations

between the training points, their true and estimated response values). Finally, the candidate(s) with the highest score(s) are selected and added to the set of training data to begin a new cycle.

Before we describe each of the three models in detail, we describe their common ingredients. Two main families of PMs are used in the uncertainty quantification: regression models such as Multivariate Adaptive Regression Splines (MARS) and stochastic models such as Gaussian processes. Both M_1 and M_2 uses an (easily replaceable) PM, i.e. a Gaussian process model (GPM), as the surrogate model. We use a specific implementation known as the Sparse Online Gaussian Process [27]. To sample from a PM, we employ a central Voronoi tessellation (CVT) scheme [28]. CVT samples cover the domain relatively densely but are sparser than points obtained by Latin Hypercube Sampling (LHS). All three models select a set of candidate points by locating the current estimate of the limit surface. The limit surface is estimated by imposing a neighborhood graph on all points with known or estimated observations and linearly interpolating along edges that span the threshold value. In the case of points estimated with a GPM, their mean value predictions are used. While many neighborhood graphs are possible [21, 29], we use the *relaxed Gabriel graph* [25] which has been shown to have relatively dense connectivities. In terms of the chosen adaptive sampling scoring function, we employ one based on point density, referred to as the *point density criterion*, where candidate points which are further away from the existing training data get higher scores. Other scoring criteria could also be employed (e.g. Active Learning MacKay criterion [30] or expected improvement criterion [31–33]), arguably leading to different convergence behaviors. However, since our focus in this paper is not on scoring function design but model selection, we employ one of the simplest scoring functions.

Finally, we need a few concepts and notations. Suppose t is the function value threshold that corresponds to the limit surface and without loss of generality, suppose the failure region is considered part of the response surface whose function value is above t . An edge in a relaxed Gabriel graph is a *spanning edge* if it spans the threshold value t in the range space, in other words, if its endpoints have function values on the opposite sides of t . A Morse-Smale crystal is a *spanning crystal* (i.e. spanning the limit surface) if its local maximum has a function value above t and its local minimum has a function value below t . A Morse-Smale crystal is a *relaxed spanning crystal* if its local maxima has a function value above $t - \delta$ and its local minimum has a function value below $t + \delta$, for some limit surface tolerance parameter $\delta \geq 0$. In a nutshell, relaxed spanning crystals allow us to investigate regions from the response surface of a PM that are on or near the current estimate of the limit surface, thus are important for exploratory analysis. This would become clearer subsequently when M_2 is described in detail.

M_1 : Global PM Limit Surface Recovery. As illustrated in Figure 2 top right, we begin with a set of training points and build a PM model using the GPM. Second, we *densely* sample the GPM over the entire domain using a CVT scheme, and construct a relaxed Gabriel graph on the union of these CVT samples with the training points. Third, candidate points are generated by linearly interpolating at the threshold value t along each spanning edge (i.e. edge in the graph that spans the threshold value t in the range space). Fourth, the set of candidate points are ranked based on the point density criterion, where a candidate with the highest score is selected and added to the set of training points to begin a new cycle.

M_2 : Topology-Based PM Limit Surface Recovery. As illustrated in Figure 2 bottom left, M_2 improves the global PM query by extracting the topology of the PM and focusing samples in specific regions. First, we fit a GPM on the initial set of training points. Second, we *sparsely* sample the GPM over the entire domain using a CVT scheme. Sparsity means that we query just enough points to obtain a reasonable estimation of the topology of the GPM. Then we impose a relaxed Gabriel graph on the union of these CVT samples with the training points, and approximate a Morse-Smale complex based on the graph structure. Third, by examining the local extrema of each crystal in the Morse-Smale complex, we can quickly discard

those crystals that are far away from the limit surface, and keep only the crystals that are located in regions of interests. The remaining crystals are precisely the relaxed spanning crystals defined previously, which are on or near the limit surface of the current surrogate model. Fourth, we store the local maxima of all relaxed spanning crystals. From each such maxima (which are considered in the failure region with high probability), we employ some sampling techniques (i.e. random walks) in its local neighborhood to obtain a set of candidate points that are near the boundary of the failure region. Those points are determined based on the GPM such that their predicted function values are in the range $[t - \delta, t + \delta]$ for some limit surface tolerance parameter $\delta \geq 0$. Finally, the set of candidate points collected across all such local maxima are ranked and selected the same way as in \mathbf{M}_1 to begin a new cycle.

To sample the neighborhood of a local maxima shared among some relaxed spanning crystals, we perform random walks described below. Starting from the maxima, we perform walks in enough random directions (i.e. scale with dimension) to characterize the domain space. The path of a random walk terminates when it hits the domain boundary or it reaches a point whose predicted function value from the GPM lies within $[t - \delta, t + \delta]$. As the distance between the random walks increases as we move further away from the seeding maxima, we add new seed points and increase the number of paths at fixed intervals along each path to better cover the domain.

In comparison with \mathbf{M}_1 , we expect several advantages of \mathbf{M}_2 with added benefits coming from the topological knowledge of the surrogate model. First, \mathbf{M}_2 allows us to more densely sample the interesting regions and quickly reject uninteresting regions of the domain space. One could achieve a similar result by refining the sampling of \mathbf{M}_1 near the estimated limit surface, but this would require more samples to be generated from an already densely sampled PM. Second, \mathbf{M}_2 introduces a tolerance parameter δ in defining the relaxed spanning crystals, therefore allowing exploratory analysis of potential failure regions within a controlled level of uncertainty. Overall, as we demonstrate in Section 4, \mathbf{M}_2 behaves comparably with \mathbf{M}_1 under general settings, and for certain special scenarios it performs drastically better.

\mathbf{M}_3 : Data-Driven Limit Surface Recovery. As opposed to building a PM, \mathbf{M}_3 (Figure 2 bottom right) begins by directly building a neighborhood structure as the surrogate model (e.g. a relaxed Gabriel graph) on the initial training data. It then creates a candidate set by first obtaining linearly interpolated points at t along spanning edges of the graph, and introducing a random perturbation ϵ along all dimensions to these points. Finally, these candidates are ranked and selected in a similar fashion as in \mathbf{M}_1 and \mathbf{M}_2 to begin a new cycle.

The candidate set we obtained through \mathbf{M}_3 is arguably sparser, however introducing a certain amount of perturbation to the linearly interpolated points enables us to explore the region surrounding the limit surface further. Note that it is not necessary to introduce this type of perturbation to candidates obtained from linear interpolation in \mathbf{M}_1 since it relies on a PM which always queries a new set of sampled points during each round, resulting in drastically different graph structures. On the other hand, since \mathbf{M}_3 does not use a PM, the graph obtained during each round changes only slightly, such that without a random perturbation the candidate points are generally located linearly along the edge of the graph, which is less desirable.

4 Experiments

We demonstrate our three adaptive sampling models on six two-dimensional datasets with specially prescribed threshold values to exploit their limit surfaces, which are the boundaries of failure regions of varying sizes and shapes. We start with several synthetic datasets (with analytic closed forms) of various complexity,

and end with a nuclear dataset modeled after a simplified Pressurized Water Reactor (PWR). It is important to note that our proposed models could be extended and applied to high-dimensional datasets, although it becomes much harder to validate the results through visual inspection.

Our first three synthetic datasets are generated as mixtures of Gaussians: **MixtureA**, **MixtureB** and **MixtureC**. The fourth and fifth synthetic datasets are based upon a modified distance field for **DistanceD** and an inverted modified distance field for **DistanceE**. The last nuclear dataset **NuclearF** comes from a simplified PWR model that has been used for station blackout analysis. We now describe the parameter settings for the experiments performed. The domain of each dataset has been normalized to be $[0, 1] \times [0, 1]$. For the initial training points, we use the same 10 points for **MixtureA** and **MixtureB**, and the same 20 points for the remaining datasets. We add additional 100 adaptively sampled points (one at a time) to the training points through our experiments. To recover points that lie on the true limit surface, we sample 10K CVT samples from the ground truth model (either in closed form for the synthetic datasets or in terms of black-box simulation for the nuclear dataset), impose a relaxed Gabriel graph on these samples and linearly interpolate the spanning edges at the given threshold value. Similar operations are performed to recover points that lie on the estimated limit surface for models M_1 and M_2 . Since M_3 does not rely on a PM, thus not allowing CVT sampling of the domain, the set of points that lies on the estimated limit surface is obtained by linearly interpolating the spanning edges of the relaxed Gabriel graph constructed on the current training points. M_2 introduces a limit surface tolerance parameter δ , which is typically set to 5% of the range of the function (except for $\delta = 10\%$ for **MixtureC** and $\delta = 20\%$ for additional tests run on **DistanceD**). M_3 employs a perturbation parameter ϵ , which is set to be 5% of the length of each dimension in the domain space (in the case where each dimension is normalized to be $[0, 1]$, $\epsilon = 0.05$). The limit surface threshold values are 0.6 for **MixtureA**, **MixtureB**, **MixtureC** and **DistanceD**; 2.58 for **DistanceE**; 800 for **NuclearF**. For M_1 , we need to sample the GPM densely to recover the points on or near the estimated limit surface, therefore we use 2000 CVT samples. On the other hand, for M_2 , we only need to sample the GPM sparsely, using just enough samples to recover its topology, thus we query 200 CVT samples.

4.1 Mixture of Gaussians

We start with three functions generated as mixtures of Gaussians. As shown in Figure 3, in all three cases, each model is able to extract each component of the limit surface to varying degrees of fidelity.

The first dataset **MixtureA** is sampled from a single Gaussian function centered in the domain, given by the function, $z = e^{-\frac{(x-0.5)^2+(y-0.5)^2}{0.2}}$. Its limit surface is the boundary of a failure region that covers a large portion of the interior of the domain and contains a single connected component. The limit surface is smooth and circular with uniform curvature. Therefore it serves as a trivial validation dataset as most models should recover it with good accuracy and with few adaptively sampled points. As shown in Figure 3 **MixtureA**, all three models succeed in this regard.

The second dataset **MixtureB** adds complexity by adding a second component and moving the failure regions such that they intersect with the domain boundary. The function used for this is just the direct combination of two uniform Gaussian kernels: $z = e^{-\frac{x^2+y^2}{0.3}} + e^{-\frac{(x-1)^2+(y-1)^2}{0.3}}$. It is interesting to see in Figure 3 **MixtureB** that both M_1 and M_2 estimate similar limit surfaces with inaccuracies near the domain boundary, while M_3 recovers the limit surface with a more accurate geometry.

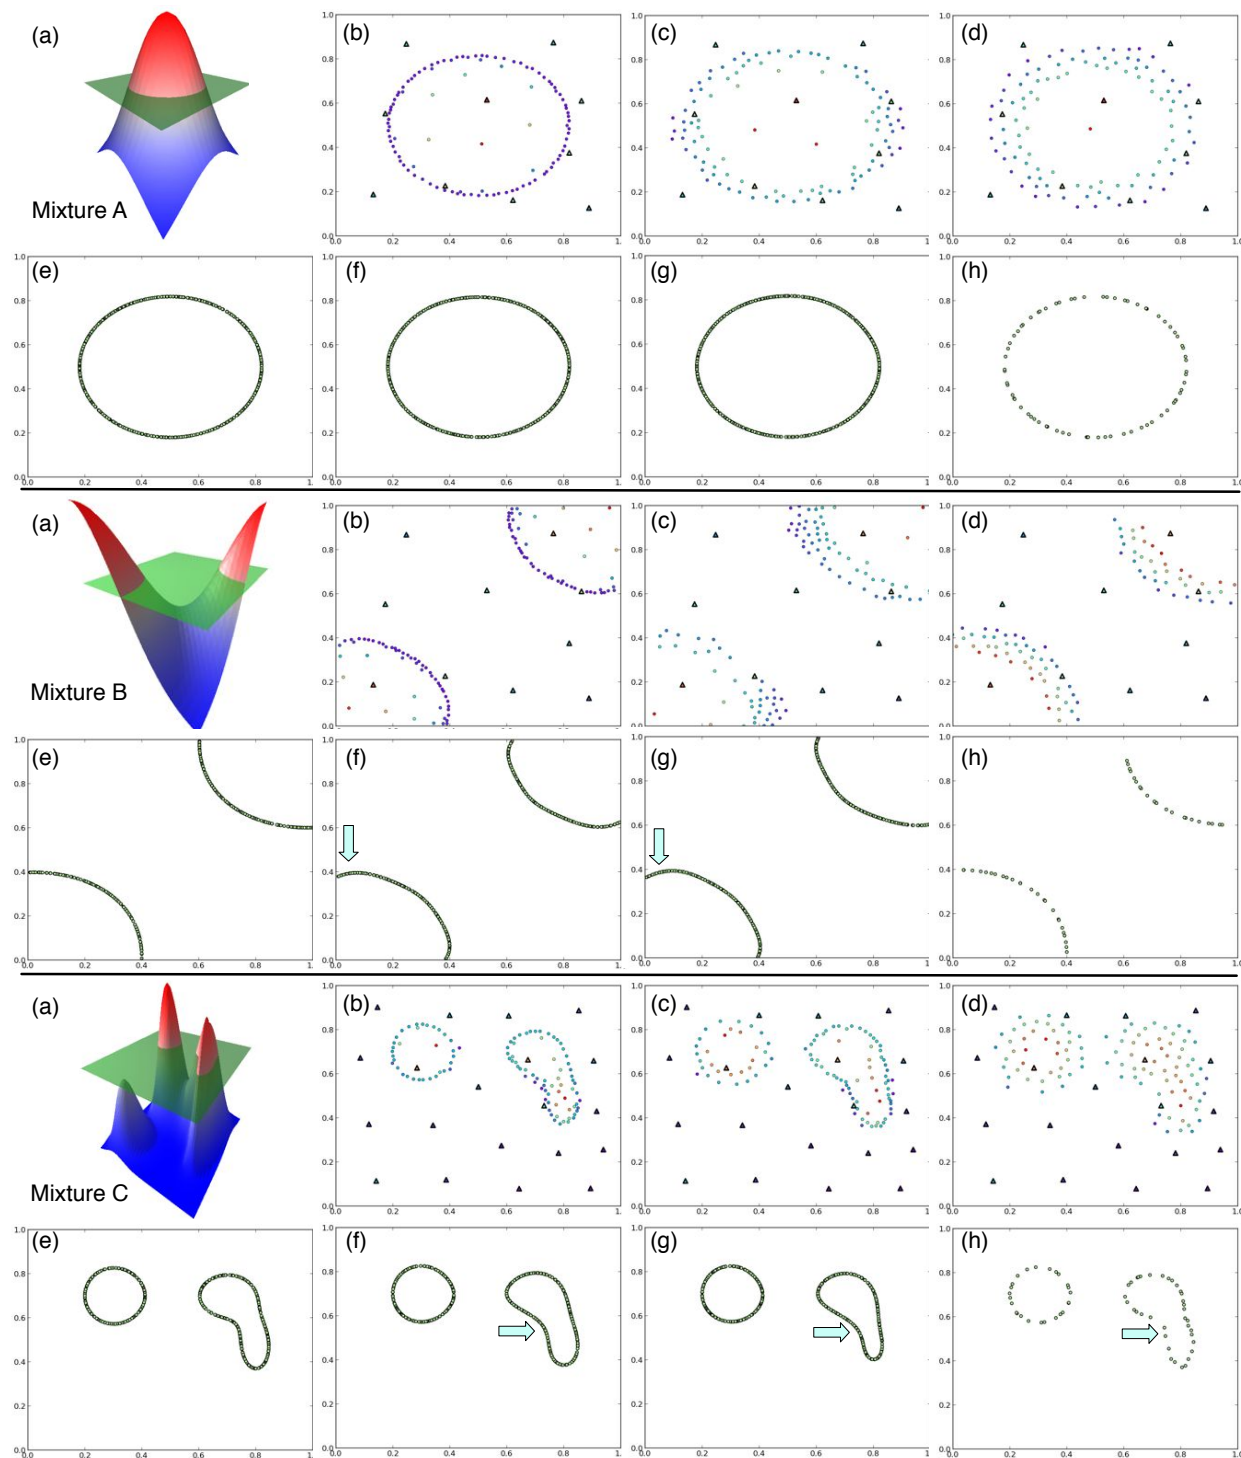


Figure 3: From top to bottom: Synthetic datasets **MixtureA**, **MixtureB** and **MixtureC**. For each dataset: (a) Surface rendering of the true response model and the true limit surface; Adaptively sampled points and training points at the end of the process under (b) M_1 , (c) M_2 and (d) M_3 ; Limit surfaces extracted from (e) the true response model, (f) M_1 , (g) M_2 and (h) M_3 . Initial training (adaptive sampled) points are solid triangles (circles), colored by true responses.

The third dataset **MixtureC** distorts the shape of the Gaussian peaks to be ellipsoidal and places two of the Gaussian centers close together to form a more complicated shape for one component of the failure region. In addition, another Gaussian function is centered nearby and has a saddle point near the limit surface. A fourth Gaussian function is created whose local maxima lies just below the limit surface. The closed form equation is given by: $z = 0.5e^{-\frac{(x-0.2)^2}{0.02} + \frac{(y-0.2)^2}{0.02}} + e^{-\frac{(x-0.3)^2}{0.02} + \frac{(y-0.7)^2}{0.03125}} + 0.8e^{-\frac{(x-0.7)^2}{0.03125} + \frac{(y-0.7)^2}{0.03125}} + 0.8e^{-\frac{(x-0.8)^2}{0.005} + \frac{(y-0.45)^2}{0.02}}$. In Figure 3 **MixtureC**, both \mathbf{M}_1 and \mathbf{M}_3 seem to capture the geometry of the right component of the limit surface better than \mathbf{M}_2 .

For **MixtureC**, since the current limit surface tolerance parameter δ is set to be 10% of the range space, neither the saddle or the lowest peak (see Figure 3 **Mixture C** (a)) are deemed interesting (that is, considered as part of the relaxed spanning crystals), therefore no candidate points are placed in their local neighborhood. As we will see in Section 4.2, in certain situations, varying δ allows exploration of the domain space that may lead to identifications of new failure regions that are otherwise undetected by the global PM.

4.2 Distance Field and Inverted Distance Field

To avoid performance bias of GPM on certain types of functions (e.g. mixture of Gaussians), it is important to construct test datasets that are geometrically distinct, as shown in Figure 4. We start with a dataset **DistanceD** constructed from a modified distance field. Such a function is composed of five centers and to each center we assign an anisotropic covariance matrix. The function value at any location is given as an L_p -norm over the vector of (anisotropic) distances to each of the centers. Therefore, each of the five center points defines a local minimum. However due to the interactions between the different centers, it produces geometrically different shapes than the ellipses, or combinations of ellipses, given by the Gaussian family of functions. The function in closed form is expressed as $f(\vec{x}) = (\sum_i \langle \vec{x} - \vec{c}_i, A_i(\vec{x} - \vec{c}_i) \rangle)^p)^{1/p}$, where \vec{c}_i is the center point, A_i is the covariance matrix and the parameter p allows tuning the smoothness of this approximation to a non-differentiable distance field ($p = \infty$ gives a true distance field). We set $p = 2$ here.

As shown in Figure 4 **DistanceD**, all three models fail to recover the third component of the limit surface (pointed by cyan arrow). This is due to the fact that the initial training points in or near that region have function values lower than the threshold, therefore the GPM from $\mathbf{M}_1/\mathbf{M}_2$ and the relaxed Gabriel graph from \mathbf{M}_3 fail to obtain a fit that predicts a potential interesting region surrounding that location. On the other hand, by increasing the tolerance level from 5% to 20% of the range space, we allow further exploration of the response surface. Subsequently our topology-based \mathbf{M}_2 is able to extract the third component successfully, as shown in Figure 5. This is our first example illustrating the potential exploratory power of a topology-based adaptive sampling model, where the parameter δ measures the amount of uncertainty.

Finally, our last synthetic dataset **DistanceE** is constructed by negating (i.e. inverting) the output values of the functional form defined in **DistanceD**, adjusting p to be 3, and deriving new centers with new covariance matrices. A limit surface threshold value that leads to three connected components in the failure region is prescribed. As shown in Figure 4 **DistanceE**, only the topology-based \mathbf{M}_2 manages to recover all three components of the limit surface using a tolerance threshold of 5% in the range space. This is our second example showcasing that the tolerance parameter δ allows \mathbf{M}_2 to detect areas near the limit surface and place candidate points in that region, potentially enhancing the exploratory power of the model.

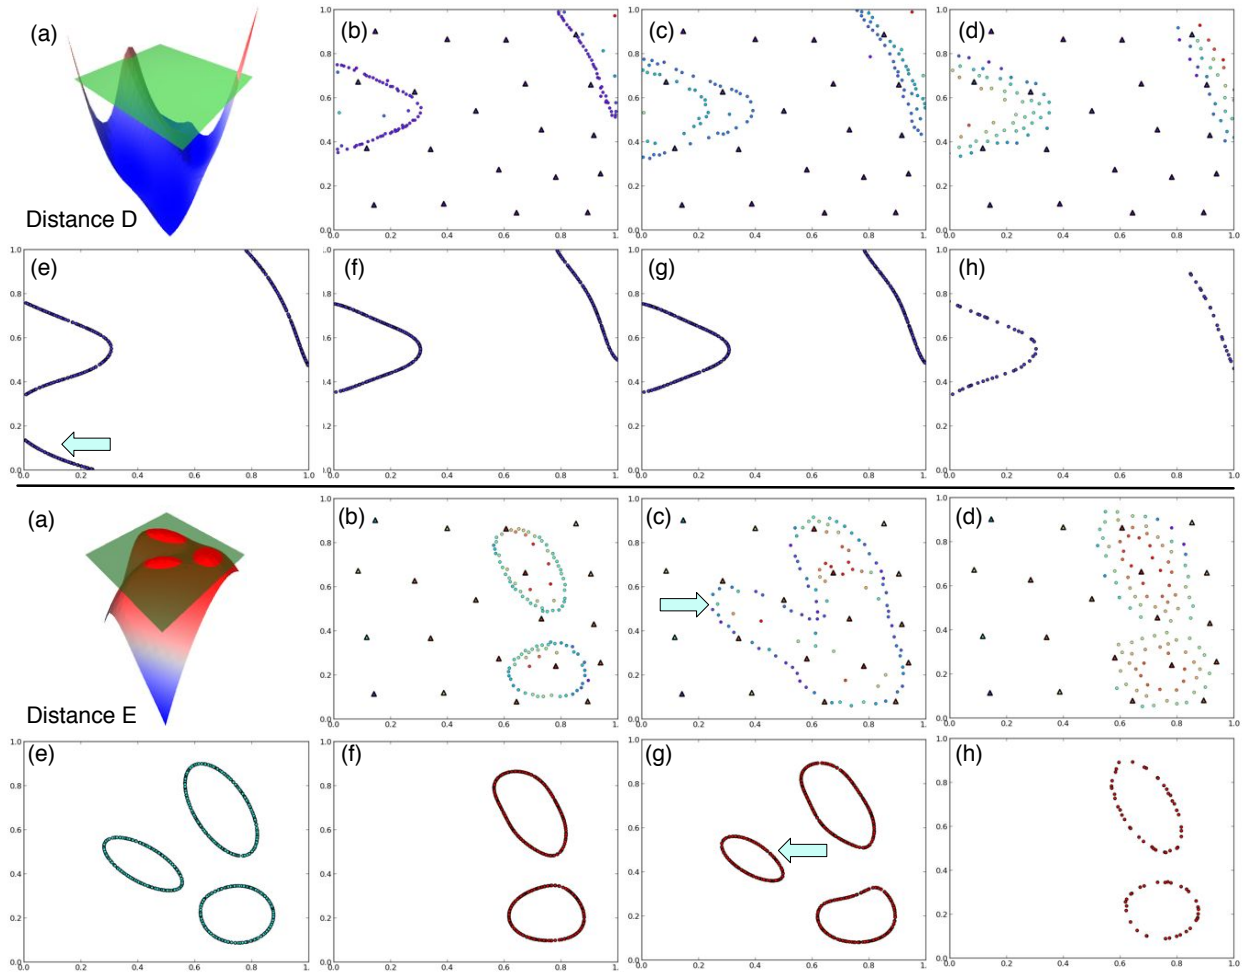


Figure 4: Synthetic datasets **DistanceD** (top) and **DistanceE** (bottom). For each dataset: (a) Surface rendering of the true response model and the true limit surface; Adaptively sampled points and training points at the end of the process under (b) M_1 , (c) M_2 and (d) M_3 ; Limit surfaces extracted from (e) the true response model, (f) M_1 , (g) M_2 and (h) M_3 . Initial training (adaptive sampled) points are solid triangles (circles), colored by true responses.

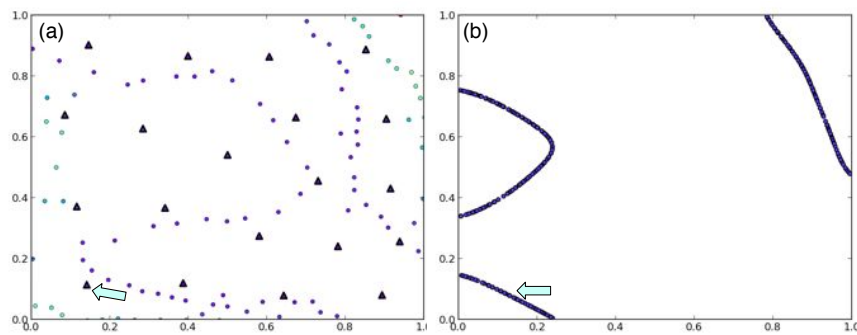


Figure 5: M_2 for **DistanceD**: Increasing limit surface tolerance parameter δ from 5% to 20% of the range space allows us to accurately capture the third component of the limit surface even though we have no initial training data whose response is above the threshold for this region.

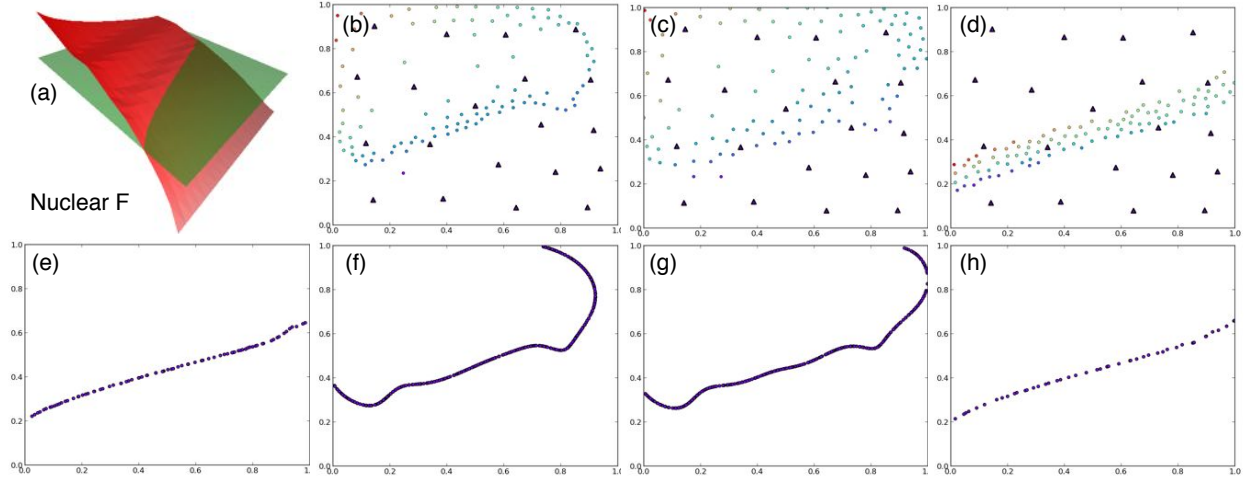


Figure 6: Nuclear dataset **NuclearF**. (a) Surface rendering of the true response model and the true limit surface; Adaptively sampled points and training points at the end of the process under (b) M_1 , (c) M_2 and (d) M_3 ; Limit surfaces extracted from (e) the true response model, (f) M_1 , (g) M_2 and (h) M_3 . Initial training (adaptive sampled) points are solid triangles (circles), colored by true responses.

4.3 Nuclear Simulator

To place our techniques in the appropriate setting, we test the validity of our proposed adaptive sampling models for a PRA application of nuclear systems, in which a two-dimensional simplified PWR simulation has been implemented using Simulink [34]. It considers a simplified PWR primary loop including the following systems: automatic depressurization, high pressure injection system (HPIS), low pressure injection systems (LPIS) and diesel generators (DG). The scope of this analysis is to evaluate the impact of Loss Of Offsite Power (LOOP) on the maximum fuel temperature. In particular, domain scientists aim to analyze uncertainties associated with the time at which the power grid is disconnected from the plant and the time associated with the diesel generators starting successfully. There are two input variables of interests in the simulation: T_{PG_SD} and T_{DG} , and the output variable is the core temperature. A typical scenario is as follows (see Figure 7):

- External events force the reactor trip at time 0;
- Offsite power is no longer available at T_{PG_SD} (LOOP condition);
- Diesel generators fail to start (station blackout condition) and batteries provide energy only for instrumentation;
- Temperature of the core starts to rise since heat removal is incapacitated. A failure condition is reached when temperature reach 800C;
- Diesel generators become available T_{DG} after LOOP condition and the ECCS is now able to remove decay heat.

For the above dataset named **NuclearF**, the results of the adaptive sampling models are shown in Figure 6. It is interesting to note that M_3 manages to recover the most accurate estimate of the limit surface. In the situation where the end users do not have enough prior information to know the *right* choice of PM to impose on the data, M_3 could become extremely valuable since it does not impose the bias from the PM (as the GPM has done in this case). However, M_3 ignores areas of the domain without any training data, therefore it is not a good exploration but a good exploitation model. Both M_1 and M_2 explore more of the

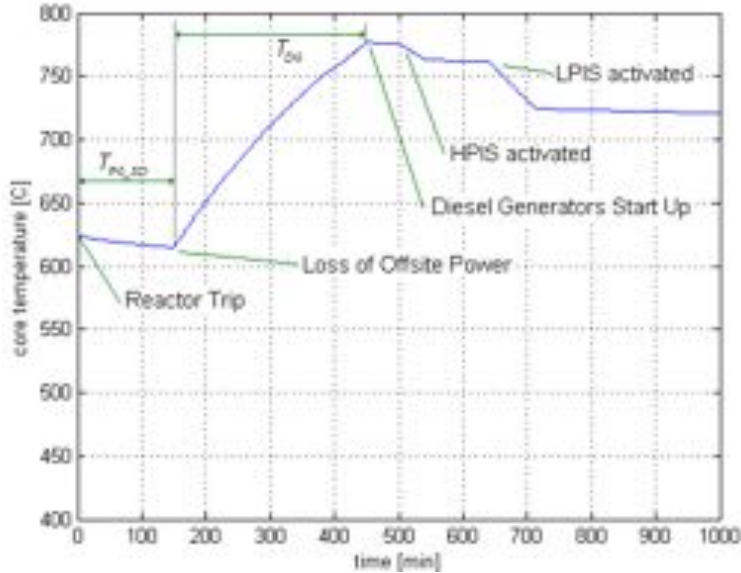


Figure 7: NuclearF: Example of a transient for a LOOP scenario.

domain space, at the cost of wasting sample points in potentially uninteresting regions. The catch here is that if we do not know what regions are interesting, it might be good to explore the domain for validation and verification purposes through M_1 or M_2 .

To further illustrate the bias GPM in M_1 and M_2 impose on the data which lead to less desirable limit surface estimates in Figure 6(f) and (g), we illustrate the evolution of the estimated response surfaces as well as estimated limit surfaces for all three models in Figure 8. These surfaces are rendered after every 10 adaptively sampled points are added to the training data up to 90 points. It is again evident that M_1 and M_2 based upon GPM introduce certain bias towards response surface fitting that more adaptively sampled points do not improve the fit drastically after 50 points or so. On the other hand, since M_3 is not PM-dependent, it converges to the true limit surface (through visual inspection) reasonably fast after 30 adaptively sampled points.

4.4 RMSE Convergence for M_1 and M_2

To further understand the convergence behavior of the limit surfaces estimated by M_1 and M_2 , we obtain convergence plots in terms of the root mean squared error (RMSE) in Figure 9. All plots show the RMSE between estimated limit surfaces and the true limit surface*, versus the number of samples added in the training data up to 100 points. The RMSE is computed using points located on the true limit surfaces.

It is important to note that M_2 behaves comparably with M_1 for general cases, e.g. for **MixtureA**, **MixtureB**, **MixtureC** and **DistanceD**; and for **DistanceD**, M_2 converges to the same level a lot faster than M_1 . On the other hand, M_2 appears to perform drastically better in certain situations, e.g. for **DistanceE** and **NuclearF**. The improved performance is due to the fact that, the topology-based model M_2 , with the tolerance parameter, allows exploration of potentially interesting regions of the domain that could contain failure regions. In the specific case of **DistanceE**, M_2 manages to find the third failure region while M_1 does not. In the case of **NuclearF**, M_2 converges faster than M_1 . This is because

*In the case where ground truth limit surface is not available, we could compute RMSE with respect to the previous round to estimate the convergence behavior.

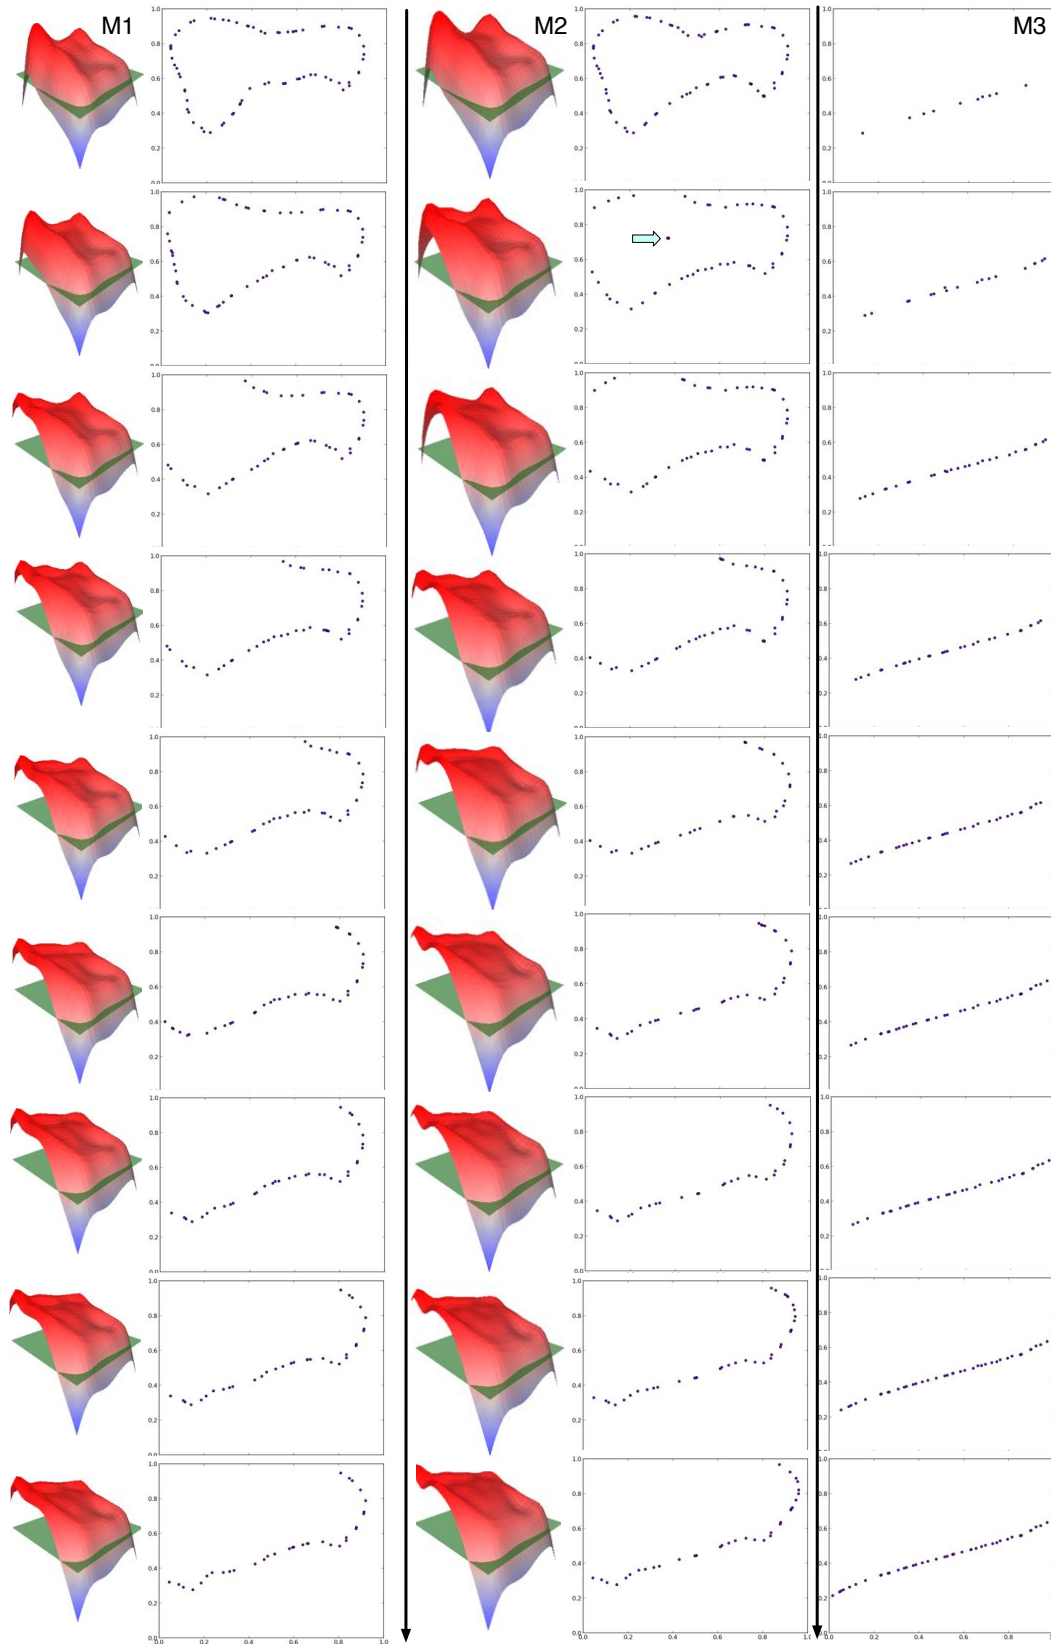


Figure 8: Convergence of estimated response surfaces and estimated limit surfaces for **NuclearF**, with all three models. Images are obtained after every additional 10 adaptively sampled points.

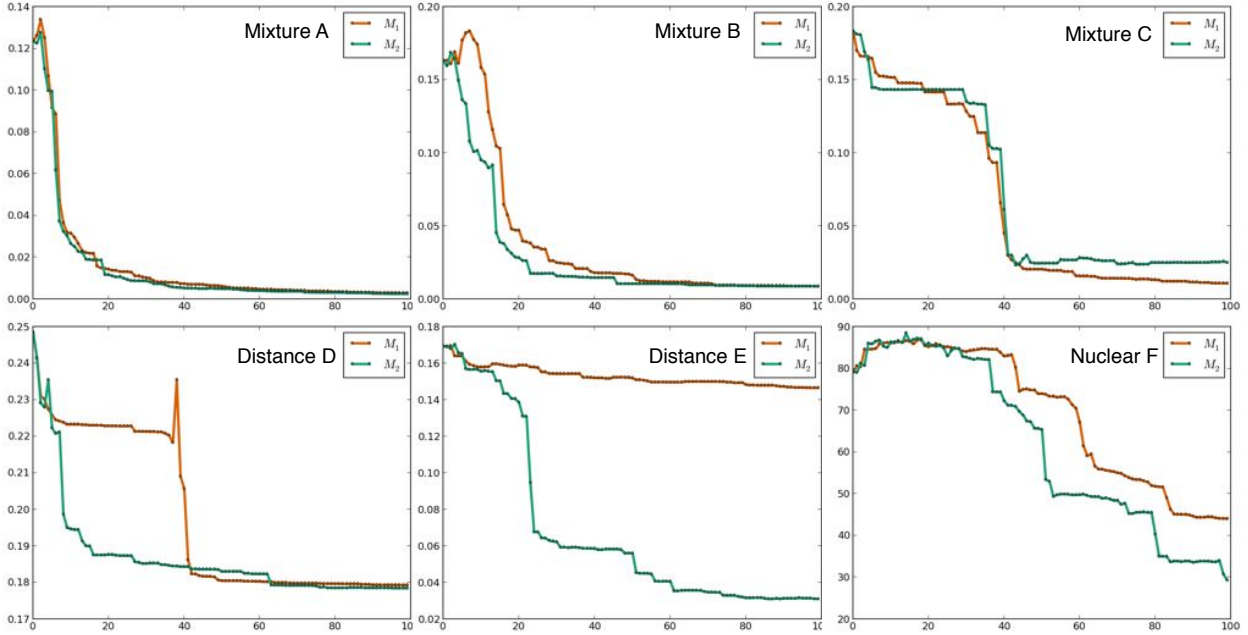


Figure 9: Convergence of RMSE of the estimated limit surfaces based on M_1 and M_2 for a single trial.

during the early stage of adaptive sampling, M_2 tries to explore the region topologically, placing a selected sampled point in the interior of the failure region, leading to faster convergence afterwards (as shown in Figure 8 second row).

To understand the average behavior of M_1 and M_2 , in addition to RMSE plots with respect to a single trial, we compute RMSE with different initial training sets and obtain the mean and median plots across 10 trials for each dataset. The mean RMSE plots are shown in Figure 10, while the median RMSE plots are displayed in Figure 11. Our conclusions do not change much with respect to mean or median convergence behavior. In general, M_1 and M_2 perform comparably for majority of the datasets. M_1 converges to a slightly lower RMSE than M_2 for **MixtureC**, while M_2 has a slight upper hand in the case of **DistanceD**. Both M_1 and M_2 converge to the same RMSE for **NuclearF**, however M_2 converges slightly faster. In the case of **DistanceE**, M_2 performs drastically better than M_1 for both mean and median plots, which is expected as M_2 recovers the third component in the limit surface which is missing from the other models.

Since M_3 does not rely on a PM, therefore it is difficult to evaluate the estimated function values of the set of points located on the true limit surface from M_3^\dagger . Thus we do not have RMSE convergence for M_3 . A more reasonable metric to measure the convergence of limit surfaces would be the dynamic time warping (DTW) distance, which is currently under consideration.

5 Conclusions and Discussions

In this paper, we explore three adaptive sampling models to recover the limit surface. From M_1 and M_3 , we learn a global model of the entire response surface using PMs such as GPMs or neighborhood graphs such as relaxed Gabriel graphs, and extract the limit surface as an iso-surface of the global model. For

[†]However this is not entirely impossible, in 2D, this involves linearly interpolating within triangles; In high dimensions, this becomes a lot harder and requires construction of high-dimensional simplexes.

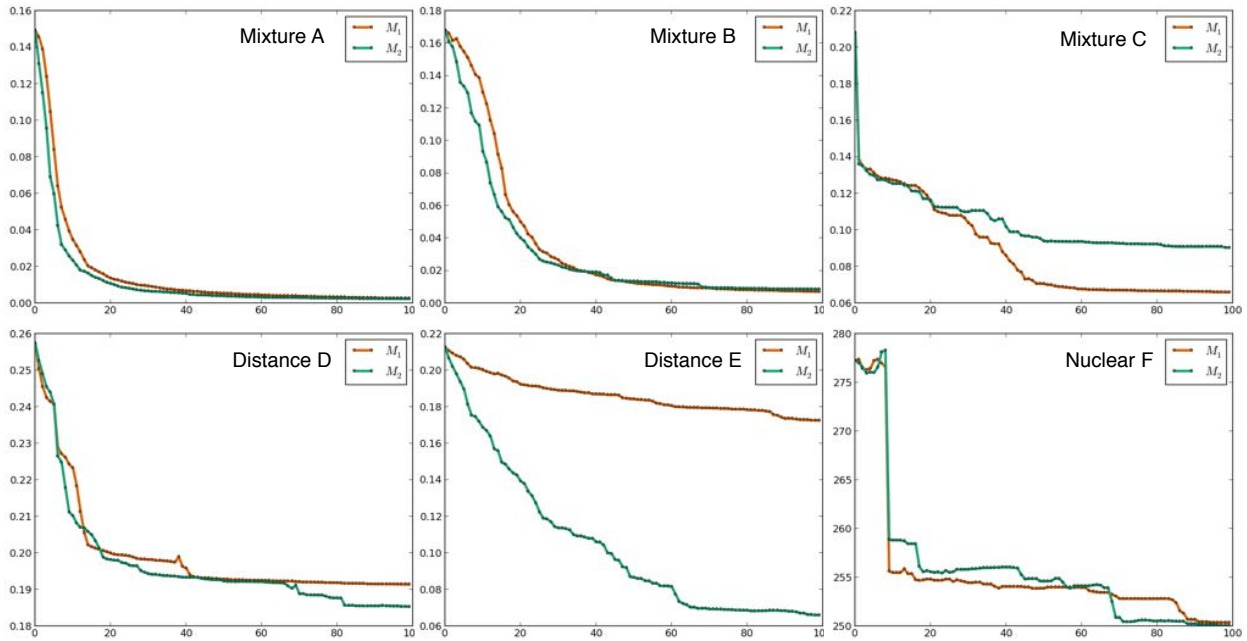


Figure 10: Convergence of mean RMSE of the estimated limit surfaces based on M_1 and M_2 across 10 trials.

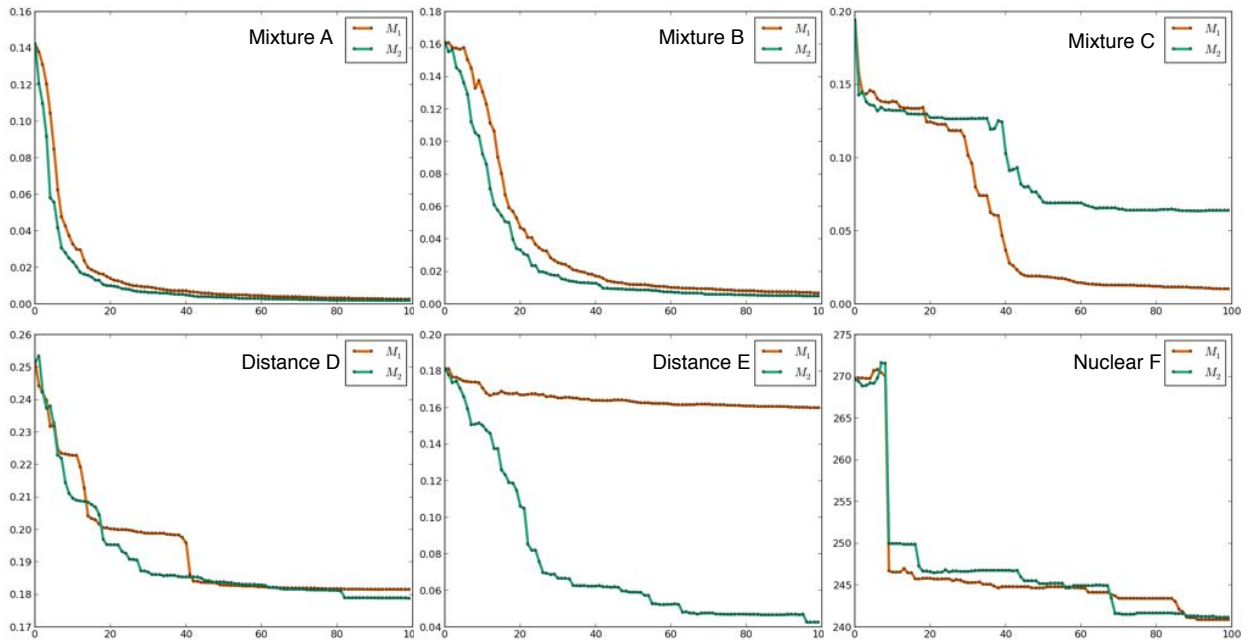


Figure 11: Convergence of median RMSE of the estimated limit surfaces based on M_1 and M_2 across 10 trials.

M_2 , we first study the topological segmentation of the global model and estimate the limit surface locally through the relaxed spanning crystals of the corresponding Morse-Smale complex of the global model. We demonstrate our models on several analytical test functions as well as a nuclear simulation dataset modeled after a simplified PWR. All models typically perform well on most of the datasets tested in recovering the limit surfaces with reasonable accuracy. M_2 performs comparably as M_1 for most cases and under certain

situations performs drastically better. M_1 and M_2 fit the training data with a PM while M_3 is completely data-driven and does not rely on a PM. M_3 performs exceptionally well on the nuclear dataset as it does not suffer the bias from a fixed PM.

Among all three models, the topology-driven model M_2 is most noticeable as it attempts to combine information from a coarse global representation using a PM with detailed local view based on topology. We expect M_2 to have several advantages over some other adaptive sampling frameworks. First, M_2 allows sampling densely around interesting regions of the response surface while avoiding uninteresting regions of the domain. In particular, it introduces a limit surface tolerance parameter δ that allows exploratory analysis of potential failure regions with controlled uncertainty. This phenomena could be investigated further in the near future where one could establish a multi-scale uncertainty characterization of the response surface by gradually varying δ . Second, we suspect that when dimension of the response surface is high and/or when the failure regions are small with respect to the entire domain area, M_2 could be an efficient model which queries much smaller number of points from the PM. It is an on-going work to validate our claim by applying the model to datasets whose response surface is high dimensional and/or with small failure regions. We may further extend topology-driven adaptive sampling schemes to include other topological structures such as contour trees or Reeb graphs, which may offer new insights to our understanding of the response surfaces.

In summary, one would need to consider the tradeoff between exploration and exploitation when choosing the appropriate adaptive sampling model. A data-driven model such as M_3 exploits the properties associated with existing training points with limited exploratory capabilities, while a PM-based model such as M_1 focuses on sampling the entire domain. A topology-driven PM-based model such as M_2 tries to strike a balance between the two through an uncertainty parameter δ . As δ increases, the model explores more uncertain neighborhoods while reduces samples in regions that strengthen the signal.

REFERENCES

- [1] E. Zio, "Reliability engineering: Old problems and new challenges," *Reliability Engineering and System Safety*, vol. 94, no. 2, pp. 125–141, 2009.
- [2] N. Siu, "Risk assessment for dynamic systems: an overview," *Reliability Engineering and System Safety*, vol. 43, no. 1, pp. 43–73, 1994.
- [3] W. Hastings, "Monte carlo sampling methods using markov chains and their applications," *Biometrika*, pp. 97–109, 1970.
- [4] J. S. Liu, *Monte Carlo strategies in scientific computing*. Springer, 2001.
- [5] J. C. Helton and F. J. Davis, "Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems," *Reliability Engineering & System Safety*, vol. 81, no. 1, 2003.
- [6] J. Geweke, "Bayesian inference in econometric models using monte carlo integration," *Econometrica*, vol. 57, no. 6, pp. 1317–39, 1989.
- [7] Owen, "Orthogonal arrays for computer experiments, integration and visualization," *Statistica Sinica*, vol. 2, pp. 439–452, 1992.
- [8] H. N. Najm, "Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics," *Annual Review of Fluid Mechanics*, vol. 41, pp. 35–52, 2009.

- [9] R. E. Caflisch, “Monte carlo and quasi-monte carlo methods,” *Acta Numerica*, vol. 7, pp. 1–49, 1998.
- [10] H. Carr, J. Snoeyink, and U. Axen, “Computing contour trees in all dimensions,” *Computational Geometry Theory and Applications*, vol. 24, no. 3, pp. 75–94, 2003.
- [11] G. Reeb, “Sur les points singuliers d’une forme de pfaff complément intégrable ou d’une fonction numérique,” *Comptes Rendus de L’Académie ses Séances Paris*, vol. 222, pp. 847–849, 1946.
- [12] V. Pascucci, G. Scorzelli, P.-T. Bremer, and A. Mascarenhas, “Robust on-line computation of reeb graphs: simplicity and speed,” *ACM Transactions on Graphics*, vol. 26, no. 3, p. 58, 2007.
- [13] H. Edelsbrunner, J. Harer, and A. J. Zomorodian, “Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds,” *Discrete and Computational Geometry*, vol. 30, pp. 87–107, 2003.
- [14] H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci, “Morse-Smale complexes for piecewise linear 3-manifolds,” *Proceedings 19th ACM Symposium on Computational Geometry*, pp. 361–370, 2003.
- [15] 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*, vol. 13, pp. 1440–1447, 2007.
- [16] H. Carr, J. Snoeyink, and M. van de Panne, “Simplifying flexible isosurfaces using local geometric measures,” in *Proceedings 15th IEEE Visualization*, pp. 497–504, 2004.
- [17] D. Laney, P.-T. Bremer, A. Mascarenhas, P. Miller, and V. Pascucci, “Understanding the structure of the turbulent mixing layer in hydrodynamic instabilities,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, pp. 1052–1060, 2006.
- [18] P.-T. Bremer, G. Weber, V. Pascucci, M. Day, and J. Bell, “Analyzing and tracking burning structures in lean premixed hydrogen flames,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 2, pp. 248–260, 2010.
- [19] A. Gyulassy, M. Duchaineau, V. Natarajan, V. Pascucci, E. Bringa, A. Higginbotham, and B. Hamann, “Topologically clean distance fields,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1432–1439, 2007.
- [20] P. Oesterling, C. Heine, H. Jaenicke, G. Scheuermann, and G. Heyer, “Visualization of high-dimensional point clouds using their density distribution’s topology,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 11, pp. 1547–1559, 2011.
- [21] C. D. Correa and P. Lindstrom, “Towards robust topology of sparsely sampled data,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, pp. 1852–1861, 2011.
- [22] J. Milnor, *Morse Theory*. New Jersey, NY, USA: Princeton University Press, 1963.
- [23] S. Gerber, P.-T. Bremer, V. Pascucci, and R. Whitaker, “Visual exploration of high dimensional scalar functions,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, pp. 1271–1280, 2010.
- [24] J. Jaromczyk and G. Toussaint, “Relative neighborhood graphs and their relatives,” *Proceedings of the IEEE*, vol. 80, no. 9, pp. 1502–1517, 1992.
- [25] P. Bose, J. Cardinal, S. Collette, E. D. Demaine, B. Palop, P. Taslakian, and N. Zeh, “Relaxed gabriel graphs,” *Proceedings Canadian Conference on Computational Geometry*, 2009.

- [26] K. Gabriel and R. Sokal, “A new statistical approach to geographic variation analysis,” *Systematic Zoology*, vol. 18, pp. 259–278, 1969.
- [27] L. Csató and M. Opper, “Sparse online gaussian processes,” *Neural Computation*, vol. 14, pp. 641–668, 2002.
- [28] Q. Du, V. Faber, and M. Gunzburger, “Centroidal voronoi tessellations: Applications and algorithms,” *SIAM Review*, vol. 41, no. 4, pp. 637–676, 1999.
- [29] 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.” Workshop on Topological Methods in Data Analysis and Visualization (accepted), 2013.
- [30] D. MacKay, “Information-based objective functions for active data selection,” *Neural Computation*, vol. 4, no. 4, pp. 589–603, 1992.
- [31] M. Schonlau, “Computer Experiments and Global Optimization.” PhD thesis, University of Waterloo, 1997.
- [32] D. Jones, M. Schonlau, and W. Welch, “Efficient global optimization of expensive black-box functions,” *Journal of Global Optimization*, vol. 13, pp. 455–492, 1998.
- [33] C. Q. Lam, “ Sequential Adaptive Designs In Computer Experiments For Response Surface Model Fit.” <http://etd.ohiolink.edu/>, 2008.
- [34] D. Mandelli and C. Smith, “Adaptive sampling using support vector machines,” in *Proceeding of American Nuclear Society*, vol. 107, pp. 736–738, 2012.