International Journal for Uncertainty Quantification, 4 (6): 535–554 (2014)

EFFECTIVE SAMPLING SCHEMES FOR BEHAVIOR DISCRIMINATION IN NONLINEAR SYSTEMS

Vu Dinh,[1,] Ann E. Rundell,[2] & Gregery T. Buzzard[1]*

[1] *Department of Mathematics, Purdue University, 150 North University Street, West Lafayette, Indiana 47907, USA*

[2] *Weldon School of Biomedical Engineering, Purdue University, 206 S. Martin Jischke Drive, West Lafayette, Indiana 47907, USA*

Original Manuscript Submitted: 01/13/2014; Final Draft Received: 08/13/2014

Behavior discrimination is the problem of identifying sets of parameters for which the system does (or does not) reach a given set of states. While there are a variety of methods to address this problem for linear systems, few successful techniques have been developed for nonlinear models. Existing methods often rely on numerical simulations without rigorous bounds on the numerical errors and usually require a large number of model evaluations, rendering those methods impractical for studies of high-dimensional and expensive systems. In this work, we describe a probabilistic framework to estimate the boundary that separates contrasting behaviors and to quantify the uncertainty in this estimation. In our approach, we directly parameterize the, yet unknown, boundary by the zero level-set of a polynomial function, then use statistical inference on available data to identify the coefficients of the polynomial. Building upon this framework, we consider the problem of choosing effective data sampling schemes for behavior discrimination of nonlinear systems in two different settings: the low-discrepancy sampling scheme, and the uncertainty-based sequential sampling scheme. In both cases, we successfully derive theoretical results about the convergence of the expected boundary to the true boundary of interest. We then demonstrate the efficacy of the method in several application contexts with a focus on biological models. Our method outperforms previous approaches to this problem in several ways and proves to be effective to study high-dimensional and expensive systems.

KEY WORDS: *representation of uncertainty, variance reduction methods, high-dimensional methods, classification, sequential data, probabilistic inference, biological modeling*

1. INTRODUCTION

Behavior discrimination, or parameter synthesis, is the problem of identifying sets of parameters for which the system does (or does not) reach a given set of states. This problem appears in science and engineering contexts in various forms. For examples, in studies of biological systems, regions of the parameter/input space with different qualitative behaviors need to be treated differently and should be identified before other tools, such as sensitivity analysis, identifiability analysis or model order reduction can be performed. Similarly, in optimal control theory, the constrained optimization problem is well-defined only on the feasible region of the parameter space and the task of computing that region is crucial in designing the control [1, 2]. In uncertainty quantification of discontinuous model response with limited model runs, a preliminary step to identify the structure of the discontinuity also needs to be done before the reconstruction of the model response can be established [3].

While there are a variety of methods to address this problem for linear systems, few techniques have been developed for nonlinear models. One of the recent successful methods for behavior discrimination of enzymatic reaction networks is proposed by Donze et al. in [4] and later [5], based on a system's reachability and robustness analysis. This algorithm employs Monte Carlo sampling and sensitivity analysis to explore the space P of possible parameter

*Correspond to Vu Dinh, E-mail: vdinh@purdue.edu

2152–5080/14/$35.00

values and identify subsets of P which robustly satisfy a property ϕ. If a subset of P does not satisfy/violate that property, the algorithm iteratively subdivides it until each subdivision entirely satisfies or entirely violates ϕ, or which are of insignificant size. The expected result of the overall procedure is a partition of P into small subsets around the boundary between satisfaction and violation of ϕ and larger regions where the satisfaction or violation is robust.

As discussed in [4] and [5], this approach has a few limitations. First, the refinement process implies that the number of partitions increases exponentially with the number of unknown parameters. Thus, in practice, some variables must be held fixed while analyzing the behavior of the model. Second, a common issue in studying complex systems of ordinary differential equations (ODEs) is that the cost for evaluation of system output for a particular set of parameter values can be very expensive. For systems with a large number of equations and unknown parameters, it is usually not computationally feasible to sample enough points by Monte Carlo sampling to explore the parameter space. Third, the method relies on numerical simulations, without rigorous bounds on the numerical errors. Finally, although the method generates a map that partitions the parameter space into regions of the same behavior, no analytical formula to describe the boundary is derived.

In this paper we introduce an alternative way to identify the boundary between the satisfaction and violation regions of a property ϕ. In our approach, we directly parameterize the, yet unknown, boundary by the zero level-set of a polynomial function, then use statistical inference on available data to identify the coefficients of the polynomial.

This idea of using statistical inference to locate the surface of discontinuity was proposed by Sargsyan et al. [3] for uncertainty quantification of mathematical models with discontinuity and limited model runs. Note that [3] focused on finding a probabilistic description of the boundary between the ”good” and ”bad” regions based on given data, and did not specify an effective way to choose the points where data should be collected for inference. In most examples, Monte Carlo sampling (which can be expensive in complex systems) was employed to choose the points of inference, with an implicit assumption that the selected data contain enough information to constrain the distribution on the set of possible boundaries effectively. Since no estimator of the true boundary was specified, no theoretical result about the convergence or rigorous bound on the numerical errors was provided.

To resolve the issue of computational cost, as well as to provide a theoretical foundation for the usage of statistical inference to locate the surface of discontinuity, in this paper we investigate two classes of sampling methods to choose the points that need to be evaluated for inference: the low-discrepancy sampling method and the sequential sampling method. The former is a well-known and effective class of sampling schemes to study high dimensional structure, while the latter selects points where current understanding about the location of the boundary is most uncertain. As we will show later in our computational results, the low-discrepancy sampling method can quickly identify the general structure of the boundary, while after a burn-in period, the sequential sampling only samples at points near the boundary of interest. With these effective sampling schemes, the number of model evaluations needed to produce a fair approximation of the boundary is significantly smaller than that required by the methods discussed above.

Our method has several advantages over previous approaches. First, the number of points to be sampled is relatively insensitive to the number of unknown parameters. Thus, the algorithm is a practical method to study high dimensional/computationally expensive systems. Second, by parameterizing the boundary as a zero level set of a polynomial function, we successfully explore a large class of boundary curves, including the case when the boundary has multiple components, which has not been addressed and analyzed in the literature. Third, for both classes of sampling schemes, we are able to provide theoretical results about the convergence of the estimated surface to the boundary of interest. Finally, by employing a probabilistic framework, our method provides a feasible way to quantify uncertainty in the discriminations. This enables further uncertainty analysis of the system of interest.

The paper is organized as follows. Section 2 provides the mathematical framework and describes our algorithm of behavior discrimination, as well as compares the method with other approaches to the problem. Section 3 establishes the convergence of the estimated surface to the boundary of interest for the two mentioned classes of sampling schemes. In Section 4, we investigate various mathematical models of biological systems to illustrate the efficacy of the method in applications. Further properties about the performance of the algorithm are provided by simulation in Section 5. Finally, in Section 6, we conclude the paper with discussions about the method and some descriptions of future work.

2. METHODOLOGY

2.1 Description of the Algorithm

In this work, we consider a continuous nonlinear system that can be described by a functional relation between a parameter vector ω and the system output y

$$y = F(\omega)$$

where $\omega = (\omega_1, \omega_2, ..., \omega_n) \in \Omega \subset \boldsymbol{R}^n$ is the parameter vector; F is a continuously differentiable function of its arguments, and $y = (y_1, y_2, ..., y_{n_y})$ is the vector of system outputs.

For a given set of parameter values, our algorithms first solve the system, then decide whether or not the corresponding trajectory satisfies the property of interest. Hereafter, we will use $G(\omega)$ to denote the response of the system with parameter ω, in which $G(\omega) = 1$ if the system satisfies the property of interest, and $G(\omega) = -1$ otherwise. Since the closed form of G is unknown and evaluating G might be expensive, we seek to approximate G by a simple rule of discrimination with small error using as few evaluations of G as possible.

2.1.1 Probabilistic representation and uncertainty quantification of the discrimination

To identify the boundary that separates the "good" and "bad" regions, we assume without loss of generality that the boundary of interest $\partial\Gamma$ is the zero level set of a smooth function $\Gamma(\omega)$ that can be well approximated by polynomials. Moreover, the polynomial approximations are expressed in the Legendre basis, instead of the monomial basis:

$$f(c, \omega) = \sum_{i=0}^{N} c_i \eta_i(\omega)$$

with the vector c contained in some coefficient space $\mathcal{C}$. In other words, the boundary of interest will be modeled as the set $\{\omega \in \Omega : f(c, \omega) = 0\}$, where the coefficient c needs to be inferred from available data. We note that the choice of N affects the performance of the method and will be investigated experimentally further in later sections.

In our method, based on the collected data $(\omega_i, G(\omega_i))$, $1 \leq i \leq m$, a probability distribution π_m is generated on $\mathcal{C}$, where $\pi_m(c)$ corresponds to the likelihood that the zero level set of the polynomial $f(c, \cdot)$ is the boundary that separates the two regions

$$\pi_m(c) \propto \exp\left(-\sum_{i=0}^{m} |G(\omega_i) - \phi(c, \omega_i)|\right) \tag{1}$$

where $\phi(c, \omega) = \text{sign}(f(c, \omega))$. Note that this distribution may be generalized (see Remark 2.1).

This distribution, when propagated to the space of all possible boundary curves, induces a probabilistic representation of the boundary. The expected prediction function with respect to the distribution π_m is defined as

$$\bar{\phi}_m(\omega) = E_{\pi_m}[\phi(c, \omega)] \tag{2}$$

while the uncertainty in the discrimination at a point ω can partially be represented by the variance in prediction

$$\text{Var}_{\pi_m}[\phi(c, \omega)] \tag{3}$$

By using formula (2), we aim to compute the prediction function associated with the true boundary, rather than the boundary itself. Formula (3) gives an estimate of the uncertainty in the prediction function. As we illustrate later in this paper, we use (1) as the basis for an Monte Carlo Markov chain (MCMC) sampling algorithm to approximate (2) and (3).

2.1.2 Effective Sampling Schemes for Behavior Discrimination

We will illustrate in this paper that when the collected data ω_i, $1 \leq i \leq m$ satisfy certain patterns, the zero level set of $\bar{\phi}_m(\omega)$ will converge to the true boundary. Specifically, we investigate the convergence of the algorithm in two separate settings:

1. Low-discrepancy sampling method: ω_i, $1 \leq i \leq m$ is a sequence with low discrepancy.
2. Sequential sampling method: data are collected sequentially, where the next data point ω_{m+1} is taken at the point where the maximum of variance in prediction with respect to π_m is achieved.

$$\omega_{m+1} = \arg\max_{\omega \in \Omega} \text{Var}_{\pi_m}[\phi(c, \omega)] \tag{4}$$

2.2 Main Results

The intuition behind the low-discrepancy sampling schemes is simple: if the data we sample on the parameter space are dense enough, and assuming that the boundary between the contrasting behaviors is smooth, we will have enough information to recover the true boundary that separates the two regions. Since low-discrepancy sampling is a well-known and effective scheme to unravel high-dimensional structure (see, for example, [6]), it is natural to use such sampling schemes for discrimination. This intuitive idea is supported by the following result:

Theorem 2.1.
Assume that the approximate model is correct, i.e.,

$$\exists c_0 \in \mathcal{C} : \qquad G(\omega) = \phi(c_0, \omega) \qquad \forall \omega \in \Omega$$

and $\{\omega_i\}$ has discrepancy D_m tending to 0 when $m \to \infty$.

Then with π_m and $\bar{\phi}_m$ defined as in (1) and (2), we have

$$\lim_{m \to \infty} \bar{\phi}_m(\omega) = G(\omega) \qquad \forall \omega \in \Omega$$

That is, the predicted classification converges pointwise to the true classification.

Moreover, for all $\epsilon > 0$, if we denote

$$D^{-1}(\epsilon) = \sup\{m + 1 : D_m \geq \epsilon\}$$

Then for $m = \Theta\left(D^{-1}(\epsilon) + N(1/\epsilon)\log(1/\epsilon)\right)$, where N is the number of terms in the polynomial expansion, we have

$$\int_\Omega |\bar{\phi}_m(\omega) - G(\omega)| d\omega \leq \epsilon$$

This theorem guarantees that when data are collected with a low-discrepancy scheme, the predicted curves converge to the true boundary that separates two regions. The theorem also provides a numerical bound for the number of points to be sampled to achieve a given level of accuracy. Notice that the number of evaluations needed to approximate the boundary within a given accuracy ϵ does not directly depend on the dimension d of the parameter space, but on the dimension of the coefficient space N. For a fixed degree of smoothness of the true boundary surface, N will increase as a polynomial in d when d becomes larger.

However, by exploring the whole parameter space by a low-discrepancy sequence, we also collect data at insignificant points that do not give much information about the location of the boundary. For complicated systems where the cost for each evaluation of data is high, such a strategy may not be practical. The sequential sampling scheme is proposed to address the issue. By choosing to observe the response at the point with highest uncertainty, the method introduces an effective way to refine the structure of the boundary. While at first sight this method may appear to be heuristic, the convergence of the method is also guaranteed by the following result:

Theorem 2.2.
Assume that the approximate model is correct, $\mathcal{C}$ has finite cardinality, and data are collected sequentially as in (4).

$$\omega_{m+1} = \arg\max_{\omega\in\Omega} \mathrm{Var}_{\pi_m}[\phi(c,\omega)]$$

Then

$$\lim_{m\to\infty} \bar{\phi}_m(\omega) = G(\omega) \qquad \forall \omega \in \Omega$$

That is, the predicted classification converges pointwise to the true classification.

In the rest of the paper, the algorithm will be analyzed under the setting specified above. However, we make the following remarks about how the results can be generalized and varied to adapt to different applications. The needed modifications to the proofs are straightforward but are omitted for clarity.

First, we notice that the form of the distribution π_m can be generalized to weighted forms. This accounts for the fact that in some applications, some samples are given more weight than others. For example, discriminations between oscillatory and nonoscillatory dynamics in the face of noise can be made only with certain degree of confidence, and those with higher confidence should be given more weight in the analysis. In feasible analysis of optimal control, one would prefer an under-approximation of the feasible set over an over-approximation, and two types of misclassification should be weighted differently.

Remark 2.1. (Generality of the distribution form)
Theorem 2.1 and Theorem 2.2 are still valid when the probability distribution in (1) is of the form

$$\pi_m(c) \propto h(c)\exp\left(-\sum_{i=0}^{m} k(\omega_i) D(G(\omega_i), \phi(c,\omega_i))\right)$$

where h, k are arbitrary positive weight functions on (C) and Ω, D is a metric on the set of real numbers. Note that the choice of D does not have much influence on the distribution since the values of G and ϕ are restricted to ± 1.

Second, criteria (4) corresponds to an optimization problem, which may be difficult in some situations. In Remark 2.2, we relax this condition to a sub-optimization problem that can be solved easily in most cases. Not only does this mean that we do not need to find the optimal sample with high accuracy, it also implies that as long as we take into account the uncertainty in prediction and make an effort to reduce it, a quick convergence toward the true boundary is to be expected.

Remark 2.2. (Generality of the sequential sampling criteria)
The convergence results from Theorem 2.2 are still valid when criteria (4) is replaced by the condition that the variance at the next evaluation point is within a fixed constant of the maximum variance. That is, there exists $C > 1$ so that for all $\omega \in \Omega$ and $m \geq 1$

$$\mathit{Var}_{\pi_{m(c)}}[\phi(c,\omega)] \leq C\, \mathit{Var}_{\pi_{m(\omega)}}[\phi(c,\omega_{m+1})] \tag{5}$$

2.3 Comparison to Other Approaches

One main difference between our method and other algorithms of boundary detection and behavior discrimination comes from our choice of estimator: instead of trying to maximize the likelihood function in order to estimate the boundary, we average the possible curves of discontinuity, weighted by the likelihood function. Our expected estimator has several advantages over the maximum likelihood estimator. First, a prediction using the averaging method is more stable than the prediction of an algorithm based on the maximum likelihood estimator [7]. Second, the problem of identifying the maximum likelihood estimator is a global optimization problem, which is more difficult in both theoretical and computational aspects, while the expected estimator can be computed easily by employing the MCMC method (whose convergence is relatively insensitive to dimension [8]). Finally, the probabilistic framework we propose

in this paper provides a feasible way to quantify uncertainty in the discriminations, which is not available in other techniques.

Another key difference is the choice of design points where data are collected. Designed to address the problem of behavior discrimination in high-dimensional and computationally expensive systems, the sequential design takes into account the information that one gains (or alternatively, the uncertainty one reduces) by collecting data. Not only does this provide a theoretical foundation for the convergence of the methods but also the number of model evaluations needed to produce a fair approximation of the boundary using our method is smaller than other methods in the literature.

Although our method shares the same idea of using statistical inference for boundary detection employed in [3], the method of inference here is quite different from the Bayesian framework suggested in that paper: instead of using the posterior distribution for inference, our computations are based on a prior distribution, which represents current understanding about the system. By this, we simplify the computing process and relax further assumption on the achieved values and the expected error in the next evaluation. The choice of basis functions (tensored Legendre polynomials) also helps produce more stable results than the monomial basis employed in [3]: a small perturbation on the leading coefficient of the monomial basis may lead to a very different function, which in turn gives rise to a very different boundary; the orthonormality of the Legendre basis guarantees that the L^2 difference between the original boundary and the perturbed one is small under small perturbation to the coefficients.

The mathematical formulation of the problem of behavior discrimination also shares a great deal of similarity with the problem of classification from machine learning. In fact, to some extent, our algorithm can be framed as an active Bayesian learning algorithm for classification. The distinguishing feature is the choice of data collected for learning: data for statistical learning (either in passive or active learning settings) are generally sampled from some underlying natural unknown distribution and learning this distribution is also a part of the process [9]. Our problem setting, however, allows us to get the response at any point on the parameter space, which provides us more freedom in data sampling. To the best of our knowledge, no analysis of a similar algorithm with either low-discrepancy or sequential data exists in the machine learning literature. From an application point of view, the ability to quantify and reduce uncertainty in inferences also distinguishes the method from other machine learning based classification techniques.

3. CONVERGENCE RESULTS

In this section, we establish results about the convergence of the expected prediction function, that is, in the limit when the number of samples m approaches infinity (following the algorithms described in the previous section), the approximation zero level set converges to the boundary between the two regions.

For simplicity, we define for every $c \in \mathcal{C}$, the binary classifications

$$\phi(c, \omega) = \text{sign}(f(c, \omega))$$

and

$$G(\omega) = \text{sign}(\Gamma(\omega))$$

(We recall that Γ is the smooth function whose zero level set is assumed to be the boundary that separates regions with different behaviors. Hence, G is the classification function of the parameter space by behaviors.)

3.1 Low-Discrepancy Sampling

We first consider the case when the sequence of samples $\{\omega_m\}$ has discrepancy approaching 0 when $m \to \infty$. More precisely, for a sequence of points $\{\omega_i\} \subset \Omega$ and a subset B of Ω, let $\#B_m$ be the number of points of $\omega_1, \ldots, \omega_m$ contained in B. Then the discrepancy of the first m points is

$$D_m(\{\omega_1, ..., \omega_m\}) = \sup_{B \subset \Omega} \left| \frac{\#B_m}{m} - \text{Vol}(B) \right|, \tag{6}$$

and the low-discrepancy condition means that D_m tends to 0 as m tends to infinity.

This condition covers a large class of sampling schemes, for example: (1) when the set of $\{\omega_i\}$ is grid points of a multilevel sparse grid in parameter space (2) $\{\omega_i\}$ is quasi-random, such as those that are collected by Latin hypercube sampling (3) when $\{\omega_i\}$ is sampled independently from an absolutely continuous distribution, or a Markov chain whose invariant measure is such a distribution.

Theorem
Assume that the approximate model is correct, i.e.,

$$\exists c_0 \in \mathcal{C}: \qquad G(\omega) = \phi(c_0, \omega) \qquad \forall \omega \in \Omega$$

and $\{\omega_i\}$ has discrepancy D_m tending to 0 when $m \to \infty$.

Then with π_m and $\bar{\phi}_m$ defined as in (1) and (2), we have

$$\lim_{m\to\infty} \bar{\phi}_m(\omega) = G(\omega) \qquad \forall \omega \in \Omega$$

That is, the predicted classification converges pointwise to the true classification.

Moreover, for all $\epsilon > 0$, if we denote

$$D^{-1}(\epsilon) = \sup\{m+1 : D_m \geq \epsilon\}$$

then for $m = \Theta\left(D^{-1}(\epsilon) + N(1/\epsilon)\log(1/\epsilon)\right)$, where N is the number of terms in the polynomial expansion, we have

$$\int_\Omega |\bar{\phi}_m(\omega) - G(\omega)| d\omega \leq \epsilon$$

Theorem 2.1 guarantees that when data are collected by one of the schemes discussed above, the predicted curves converge to the true boundary that separates two regions. The theorem also provides a numerical bound for the number of points to be sampled to achieve a given level of accuracy.

3.1.1 Preliminary lemmas

Before moving forward to provide the proof for Theorem 2.1, we first establish the following two lemmas:

Lemma 3.1.
Let g be a nonconstant polynomial on a compact subset $\mathcal{P} \subset \mathcal{R}^n$. Denote $V = \{p \in \mathcal{P} : g(p) = 0\}$.

Then there exists k and C_1 depending only on g such that

$$\text{Vol}(\{p \in \mathcal{P} : |g(p)| \leq \epsilon\}) \leq \text{Vol}(\{p \in \mathcal{P} : \text{dist}(p, V) \leq (C_1\epsilon)^{1/k}\})$$

for all $\epsilon > 0$.

Proof. By Lojasiewicz inequality (see, for example, [10]), there exists k and C_1 depends only on g such that

$$\text{dist}(p, V)^k \leq C_1 |g(p)|$$

Hence, for every p that satisfies $|g(p)| \leq \epsilon$, we have $\text{dist}(p, V) \leq (C_1\epsilon)^{1/k}$. This deduces

$$\{p \in \mathcal{P} : |g(p)| \leq \epsilon\} \subset \{p \in \mathcal{P} : \text{dist}(p, V) \leq (C_1\epsilon)^{1/k}\}$$

□

Lemma 3.2.
Let V be an algebraic surface with Hausdorff dimension $n-1$ on a compact subset $\mathcal{P} \subset \mathcal{R}^n$. Then

$$\text{Vol}(\{p \in \mathcal{P} : \text{dist}(\omega, V) \leq \epsilon\}) \leq C_2\epsilon \qquad \forall \epsilon > 0$$

where C_2 depends only on V and the volume constant in n dimensions.

Proof. Since V has Hausdorff dimension $n-1$, there exist C depending only on V such that for all $\epsilon > 0$, the number of balls with radius ϵ needed to cover V ($K(\epsilon)$) satisfies

$$K(\epsilon) \leq \frac{C}{\epsilon^{n-1}}$$

Such cover of V will also contain the set

$$\{p \in \mathcal{P} : \text{dist}(\omega, V) < \epsilon\}$$

as a subset. Since the volume of an n-dimensional ball with radius r is equal to $C_n r^n$, we deduce that

$$\text{Vol}(\{p \in \mathcal{P} : \text{dist}(p, V) \leq \epsilon\}) \leq \frac{C}{\epsilon^{n-1}} C_n \epsilon^n = C C_n \epsilon$$

which completes the proof. □

3.1.2 Proof of Theorem 2.1

Throughout this section, we denote

$$e_m(c) = \frac{1}{m} \sum_{i=1}^{m} |\phi(c, \omega_i) - G(\omega_i)| \tag{7}$$

$$e(c) = \int_\Omega |\phi(c, \omega) - G(\omega)| \, d\omega \tag{8}$$

and

$$\mathcal{C}_\epsilon = \{c \in \mathcal{C} : e(c) \leq \epsilon\} \tag{9}$$

The proof for Theorem 2.1 can be summarized as follows: In Lemma 3.3, we prove that outside the set of "good" candidates $\mathcal{C}_\epsilon$ (on which the error of prediction is less than ϵ), the distribution π_m converges to zero exponentially at a rate depending on $\text{Vol}(\mathcal{C}_\epsilon)$. Lemmas 3.5 and 3.6 establish a lower bound on $\text{Vol}(\mathcal{C}_\epsilon)$ in term of ϵ and N (the number of terms in the polynomial expansion). The combination of those results completes the proof of Theorem 2.1.

Lemma 3.3.
Let $\epsilon > 0$ and $m \geq D^{-1}(\epsilon/8)$. For all $c \in \mathcal{C} \setminus \mathcal{C}_\epsilon$, we have

$$\pi_m(c) \leq \frac{\exp(-m\epsilon/4)}{\text{Vol}(\mathcal{C}_{\epsilon/4})}$$

Proof. Denote $r_m(c) = \exp(-e_m(c))$ and

$$q_m(c) = \exp\left(-\sum_{i=1}^{m} |\phi(c, \omega_i) - G(\omega_i)|\right)$$

Consider $\epsilon > 0$ and $m \geq D^{-1}(\epsilon/8)$, if we denote $B = \{\omega : \phi(c, \omega) \neq G(\omega)\}$, then

$$|e(c) - e_m(c)| = 2 \left| \frac{\#\{\omega_i \in B\}}{m} - \text{Vol}(B) \right| \leq 2D^{-1}(\epsilon/8) = \epsilon/4$$

for all $c \in C$.

Hence, for $c \in \mathcal{C}_{\epsilon/4}$, we deduce that $e_m(c) \leq |e(c) - e_m(c)| + e(c) \leq \epsilon/2$.

Therefore

$$\|r_m\|_m^m = \int_{\mathcal{C}} |r_m(c)|^m \geq \int_{\mathcal{C}_{\epsilon/4}} |r_m(c)|^m \geq \exp(-m\epsilon/2)\,\mathrm{Vol}\,(\mathcal{C}_{\epsilon/4})$$

Now for $c \in \mathcal{C} \setminus \mathcal{C}_\epsilon$, since $|e(c) - e_m(c)| \leq \epsilon/4$, we deduce that $e_m(c) > 3\epsilon/4$. Hence $r_m(c) \leq \exp(-3\epsilon/4)$ and

$$\pi_m(c) = \frac{q_m(c)}{\int_{\mathcal{C}} q_m(c)dc} = \left(\frac{r_m}{\|r_m\|_m}\right)^m$$

$$\leq \frac{\exp(-m\epsilon/4)}{\mathrm{Vol}\,(\mathcal{C}_{\epsilon/4})}$$

□

Lemma 3.4.
For $c_1, c_2 \in \mathcal{C}$, we have

$$|e(c_1) - e(c_2)| \leq 2\,\mathrm{Vol}(\{\omega : |f(c_1,\omega)| \leq |f(c_2,\omega) - f(c_1,\omega)|\}).$$

Proof. For any $c \in \mathcal{C}$, denote

$$\Omega_c = \{\omega \in \Omega : \phi(c,\omega) \neq G(\omega)\}$$

We have

$$e(c_2) - e(c_1) = 2\left(\int_{\Omega_1} d\omega - \int_{\Omega_2} d\omega\right)$$

$$= 2\left(\int_{\Omega_1 \setminus \Omega_2} d\omega - \int_{\Omega_2 \setminus \Omega_1} d\omega\right)$$

$$\leq 2\,\mathrm{Vol}(\{\omega : \mathrm{sign}(|f(c_1,\omega)|) \neq \mathrm{sign}(|f(c_2,\omega)|)\})$$

Note that for $a, b \in \mathbb{R}$ and $b \neq 0$, $\mathrm{sign}(a) \neq \mathrm{sign}(b)$ implies $|a| < |a-b|$. Then

$$|e(c_1) - e(c_2)| \leq 2\,\mathrm{Vol}(\{\omega : |f(c_1,\omega)| \leq |f(c_2,\omega) - f(c_1,\omega)|\})$$

□

Lemma 3.5.
There exists $C > 0, k \geq 1$ depending only on $\mathcal{C}$ and Ω such that for all $c_1, c_2 \in \mathcal{C}$ and $\epsilon > 0$

$$|e(c_1) - e(c_2)| \leq C|c_1 - c_2|^{1/k}$$

Proof. Let $c_1, c_2 \in \mathcal{C}$. Since $\mathcal{C}$ and Ω are compact and f is smooth, there exist C_3 depending only on $\mathcal{C}$ and Ω such that

$$|f(c_2,\omega) - f(c_1,\omega)| \leq C_3|c_2 - c_1|$$

We deduce

$$\begin{aligned} &\mathrm{Vol}(\{\omega \in \Omega : |f(c_1,\omega)| \leq |f(c_2,\omega) - f(c_1,\omega)|\}) \\ \leq\ &\mathrm{Vol}(\{\omega \in \Omega : |f(c_1,\omega)| \leq C_3|c_2 - c_1|\}) \end{aligned} \tag{10}$$

Using this and Lemma 3.4, we have

$$|e(c_1) - e(c_2)| \leq 2\,\mathrm{Vol}(\{\omega \in \Omega : |f(c_1,\omega)| \leq C_3|c_2 - c_1|\}) \tag{11}$$

Applying Lemma 3.1 with $\epsilon = C_3|c_2 - c_1|$ and $V = \{\omega \in \Omega : f(c_1, \omega) = 0\}$, we deduce that there exist C_1 and k depending only on c_1 such that

$$\begin{aligned} &\text{Vol}(\{\omega : |f(c_1, \omega)| \le C_3|c_2 - c_1|\} \\ &\le \text{Vol}(\{\omega \in \Omega : \text{dist}(\omega, V) \le (C_1C_3|c_2 - c_1|)^{1/k}\}) \end{aligned} \tag{12}$$

By Lemma 3.2 with $\epsilon = (C_1C_3|c_2 - c_1|)^{1/k}$, we have

$$\begin{aligned} \text{Vol}(\{\omega \in \Omega : \text{dist}(\omega, V) \le &(C_1C_3|c_2 - c_1|)^{1/k}\}) \\ &\le C_2(C_1C_3|c_2 - c_1|)^{1/k} \end{aligned} \tag{13}$$

Combining (10)–(13), we deduce

$$|e(c_1) - e(c_2)| \le 2C_2(C_1C_3)^{1/k}|c_2 - c_1|^{1/k}$$

Since $\mathcal{C}$ is compact, the constants can also be chosen independent of c_1.

□

Lemma 3.6.
For $\mathcal{C}_\epsilon$ defined as in (7), k the constant defined in Lemma 3.5 and N the dimension of the coefficient space, there exists C depending only on $\mathcal{C}$ and Ω such that for all $\epsilon > 0$

$$\text{Vol}(\mathcal{C}_\epsilon) \ge C\epsilon^{Nk}$$

Proof. Recall that c_0 is the true vector of coefficients [hence $e(c_0) = 0$]. Then for $c \in \mathcal{C}$ such that $|c - c_0| \le (\epsilon/C)^k$, where k and C are defined as in Lemma 3.5, we have

$$|e(c) - e(c_0)| \le C|c - c_0|^{1/k} \le \epsilon$$

which implies that $c \in \mathcal{C}_\epsilon$.

Therefore

$$B = \left\{ c \in \mathcal{C} : |c - c_0| \le \left(\frac{\epsilon}{C}\right)^k \right\} \subset \mathcal{C}_\epsilon$$

and

$$\text{Vol}(\mathcal{C}_\epsilon) \ge \text{Vol}(B) = C_1 \left(\frac{\epsilon}{C}\right)^{Nk}$$

where C_1 is the volume constant in N-dimensional space. □

Proof of Theorem 2.1

Proof. We have

$$\int_\Omega \left|G(\omega) - \bar{\phi}_m(\omega)\right| d\omega = \int_\Omega \left|E_{\pi_m}[G(\omega) - \phi(c, \omega)]\right| d\omega$$

$$\le \int_\Omega E_{\pi_m}[|G(\omega) - \phi(c, \omega)|]d\omega = E_{\pi_m}\left[\int_\Omega |G(\omega) - \phi(c, \omega)|\, d\omega\right]$$

$$= E_{\pi_m}[e(c)] = \int_{e(c) > \epsilon/2} e(c)\pi_m(c)dc + \int_{e(c) \le \epsilon/2} e(c)\pi_m(c)dc \tag{14}$$

$$\le \text{Vol}(\mathcal{C}) \exp(-m\epsilon/8) \frac{1}{\text{Vol}(\mathcal{C}_{\epsilon/8})} + \epsilon/2 \tag{15}$$

with the inequality obtained by using Lemma 3.3 for $m \geq D^{-1}(\epsilon/16)$. Hence if we further choose m that satisfies

$$\text{Vol}(\mathcal{C}) \exp(-m\epsilon/8) \frac{1}{\text{Vol}(\mathcal{C}_{\epsilon/8})} \leq \epsilon/2$$

or equivalently, in Θ notation (using Lemma 3.6)

$$m = \Theta\left(D^{-1}(\epsilon/16) + Nk \frac{1}{\epsilon} \log \frac{1}{\epsilon} \right)$$

then

$$\int_{\Omega} \left| G(\omega) - \bar{\phi}_m(\omega) \right| d\omega \leq \epsilon$$

which completes the proof of Theorem 2.1. □

3.2 Sequential Sampling

The sequential sampling scheme is proposed to study high-dimensional and computationally expensive systems. By choosing to observe the response at the point with highest uncertainty, the method introduces an effective way to reduce the uncertainty and refine the structure of the boundary. While at first sight this method may appear to be heuristic, the convergence of the method is also guaranteed by the following result:

Theorem 3.1.
Assume that the approximate model is correct, $\mathcal{C}$ has finite cardinality, and data are collected sequentially as in (4).

$$\omega_{m+1} = \arg\max_{\omega \in \Omega} \text{Var}_{\pi_m}[\phi(c, \omega)]$$

Then

$$\lim_{m \to \infty} \bar{\phi}_m(\omega) = G(\omega) \qquad \forall \omega \in \Omega$$

That is, the predicted classification converges pointwise to the true classification.

Before moving forward to provide the proof of the theorem, note that in the sequential setting, an additional condition is imposed on the coefficient space: the space is supposed to have finite cardinality. This condition comes from the fact that in a continuous framework, a coefficient vector has measure zero and good performance of the true coefficient vector does not always guarantee that it will stay in the support of the limit distribution. Though this situation is perhaps unlikely to happen in practice, we cannot exclude such a possibility for a convergence result. We also note that this assumption was also commonly used in the context of parameter identification [11, 12]. In practice this condition may be achieved without affecting the model's ability to approximate the true boundary, for example, by subdividing each coordinate axis using a fixed step size and taking the set of points that lie on the resulting grid.

Proof. Denote

$$q_m(c) = \exp\left(-\sum_{i=1}^{m} |\phi(c, \omega_i) - G(\omega_i)| \right)$$

Then $\pi_m = C_m q_m$. Also, let A be the set of cluster points of $\{\omega_m\}$: points $\omega \in \Omega$ such that there exists a subsequence $\{\omega_{m_k}\}$ of with $\omega_{m_k} \to \omega$.

Step 1: We claim first that if $\pi_m(c)$ does not tend to 0 with m (so that c has probability above some fixed $\rho > 0$ for infinitely many m), then $f(c, \omega) = G(\omega)$ for all $\omega \in A$.

Proof:

Consider any $c \in \mathcal{C}$, $\omega \in A$ such that $|\phi(c,\omega) - G(\omega)| > 0$. Then there exists a subsequence $\{\omega_{m_k}\}$ of $\{\omega_m\}$ such that $\omega_{m_k} \to \omega$ and $|\phi(c,\omega_{m_k}) - G(\omega_{m_k})| \geq 1$. Hence

$$\sum_{i=1}^{m} |\phi(c,\omega_i) - G(\omega_i)| \to \infty$$

when $m \to \infty$, and so $q_m(c) \to 0$.

The assumption that there exists c_0 such that $\phi(c_0,\omega) = G(\omega)$ for all ω implies that $q_m(c_0) = 1$. Therefore, $\pi_m(c_0)/\pi_m(c) \to \infty$. Since $\mathcal{C}$ is a finite space, $\pi_m(c_0) \leq 1$, and hence $\pi_m(c) \to 0$. Hence $\pi_m(c) \not\to 0$ implies $\phi(c,\omega) = G(\omega)$ for all $\omega \in A$.

We deduce that

$$\lim_{m\to\infty} E_{\pi_m}[\phi(c,\omega)] = G(\omega) \qquad \forall \omega \in A$$

and

$$\mathrm{Var}_{\pi_m}[\phi(c,\omega)] \to 0 \qquad \forall \omega \in A$$

Step 2: Now we claim that

$$\mathrm{Var}_{\pi_m}\left[\phi(c,\omega_{m+1})\right] \to 0 \tag{16}$$

Proof:

We have

$$E_{\pi_m}\left[|\phi(c,\omega_{m+1}) - G(\omega_{m+1})|^2\right]$$
$$= \sum_{c:\pi_m(c)\to 0} |\phi(c,\omega_{m+1}) - G(\omega_{m+1})|^2 \pi_m(c) + \sum_{c:\pi_m(c)\not\to 0} |\phi(c,\omega_{m+1}) - G(\omega_{m+1})|^2 \pi_m(c)$$

However, by the same argument as above, $\pi_m(c) \not\to 0$ implies that c makes only a finite number of mistakes in prediction, which implies that $\phi(c,\omega_{m+1}) = G(\omega_{m+1})$ for m large enough. Therefore

$$E_{\pi_m}\left[|\phi(c,\omega_{m+1}) - G(\omega_{m+1})|^2\right] \to 0$$

and

$$\mathrm{Var}_{\pi_m}\left[\phi(c,\omega_{m+1})\right] = E_{\pi_m}\left[|\phi(c,\omega_{m+1}) - G(\omega_{m+1})|^2\right] - \left(G(\omega_{m+1}) - E_{\pi_m}\phi(c,\omega_{m+1})\right)^2$$
$$\leq E_{\pi_m}\left[|\phi(c,\omega_{m+1}) - G(\omega_{m+1})|^2\right] \to 0$$

Step 3: The choice of ω_{m+1} gives

$$\mathrm{Var}_{\pi_m}\left[\phi(c,\omega)\right] \leq \mathrm{Var}_{\pi_m}\left[\phi(c,\omega_{m+1})\right] \qquad \forall \omega \in \Omega \tag{17}$$

and the right-hand side tends to 0 as $m \to \infty$ by step 2.

Then for all $\omega \in \Omega$

$$\lim_{m\to\infty} \sum_{\pi_m(c)\not\to 0} |\phi(c,\omega) - E_{\pi_m}\phi(c,\omega)|^2 = 0$$

From the fact that $\pi_m(c_0) \geq \pi_m(c)\ \forall c \in C$ and $\phi(c_0,\omega) = G(\omega)\ \forall \omega \in \Omega$, we have

$$|G(\omega) - E_{\pi_m}\phi(c,\omega)|^2 = |\phi(c_0,\omega) - E_{\pi_m}\phi(c,\omega)|^2$$
$$\leq \#\mathcal{C} \sum |\phi(c,\omega) - E_{\pi_m}\phi(c,\omega)|^2 \pi_m(c) = \#\mathcal{C}\,\mathrm{Var}_{\pi_m}\left[\phi(c,\omega)\right] \to 0$$

as $m \to \infty$. Hence

$$\lim_{m\to\infty} E_{\pi_m}\phi(c,\omega) = G(\omega) \qquad \forall \omega \in \Omega$$

□

4. NUMERICAL RESULTS

4.1 A Model of the Acute Inflammatory Response to Infection

We applied our method to a model of the acute inflammatory response to infection. This 4-equation, 22-parameter model was presented in [13], where the state variables P, N_A, D, and C_A, correspond to the amounts of pathogen, pro-inflammatory mediators (e.g., activated neutrophils), tissue damage, and anti-inflammatory mediators (e.g., cortisol and interleukin-10), respectively.

The response is a complex process that exhibits both pro- and anti-inflammatory behaviors. The pro-inflammatory elements are responsible for eliminating the pathogen, but pathogen killing can cause collateral tissue damage. This in turn triggers an escalation in the pro-inflammatory response, creating a positive feedback cycle. The anti-inflammatory elements counteract this cycle, minimizing tissue damage and promoting healing. However, in cases of extreme infection, this delicate balance is destroyed, which leads to a potentially lethal amount of tissue damage.

In this example, we validated our method by reproducing results previously obtained in [13], and later in [4]. Death is defined as a sustained amount of tissue damage (D) above a specified threshold value and constitutes the undesirable outcome we wish to avoid. The two unknown parameters considered are the initial amount of pathogen, P_0, and the initial amount of anti-inflammatory mediators, CA_0. The parameter space Ω is defined relative to a nominal parameter vector; that is, for each component of the nominal vector, we define a range of 10 times smaller to 10 times larger than this component. The growth rate of pathogen was set to 0.3 and other parameters to their nominal values as in [4].

For boundary detection, we evaluate the model on the grid points of the 16-point, 36-point, and 80-point two-dimensional uniform sparse grid. The result is presented in Fig. 1(a), where the parameter space is scaled to the unit square. The expected boundary is computed using a 10^5-point Monte Carlo Markov Chain on a 13-dimensional coefficient space (that is, $N = 13$). The result is similar to what achieved in [4], but the number of model evaluations needed to produce this result is significantly lower.

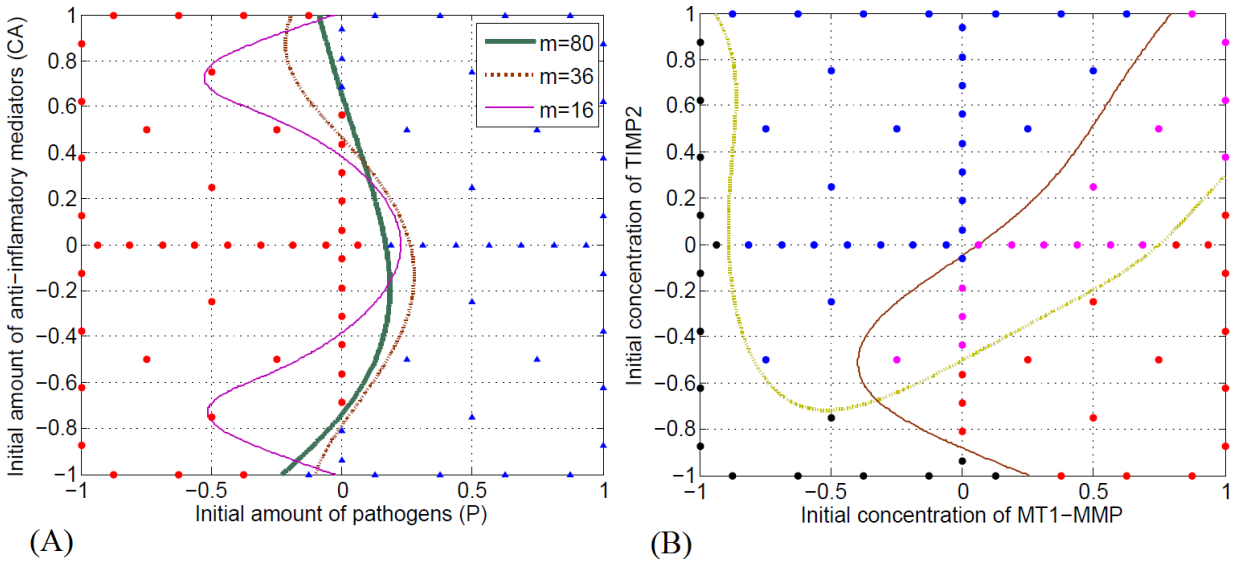

FIG. 1: (a) Model of the acute inflammatory response to infection: The predicted boundaries computed by sparse sampling at the 16-point, 36-point and 80-point sparse grid nodes. (b) Model of collagen degradation: the design points and predicted boundaries computed at the nodes of the 80-point sparse grid. In both figures, the expected boundaries are computed using a 10^5-point MCMC on a 13-dimensional coefficient space.

4.2 A Model of Collagen Degradation

Our second example is performed on the biochemical network adapted from [14] and later extended in [5], which models the loosening of the extra-cellular matrix, a crucial process in angiogenesis, the sprouting of new blood vessels as a reaction to signals that indicate the need for additional oxygen in certain tissues. The system consists of 12 differential equations that integrates on a long time scale, which makes model evaluations become very expensive.

We investigate the relative contribution of MT1-MMP and MMP2 on collagen proteolysis using a combination of constraints imposed on (1) the amount of collagen that has been degraded after a given time and (2) the respective contribution of any of the two enzymes onto collagen degradation. Studying those two properties, we are able to create the division maps in Fig. 1(b). The region above the solid curve corresponds to the case when the amount of collagen degraded by MT1-MMP is greater than that degraded by MMP2; the region on the right of the other curve corresponds to the case when the system does (does not) manage to degrade 90% of the collagen before 12 h. Similar to the previous example, the boundaries in Fig. 1(b) were computed using data at the first 80 grid points of a uniform sparse grid. This division map replicates the result about collagen proteolysis previously achieved in [5] with a lower number of evaluations.

In Fig. 2(a), we employ the sequential sampling scheme to approximate the blue boundary in Fig. 1(b) with higher accuracy using the same number of model evaluations. Starting with a prior data set collected at the first 16 points of a uniform sparse grid, we use the relaxed variance criteria to choose the next sample point until 80 data points are obtained. We can see that the boundary generated by the sequential sampling method converges quickly to the true boundary in less than a hundred model evaluations. This advantage comes from the fact that after a burn-in period, the algorithm selects points near the true boundary and focuses the probability distribution on $\mathcal{C}$ on a few regions that contain potential candidates.

As we emphasized earlier in the Introduction, the framework proposed in this method provides a natural probabilistic representation of the classifying boundary, which enables further uncertainty analysis of the system. This is a feature that distinguishes the algorithm from other approaches. In Fig. 2(b), we provide a contour map of the variance in prediction with respect to the data-dependent probability distribution π_m constructed on the coefficient space. The variance map can be considered as a representation of the uncertainty in discrimination, or a measure of confidence

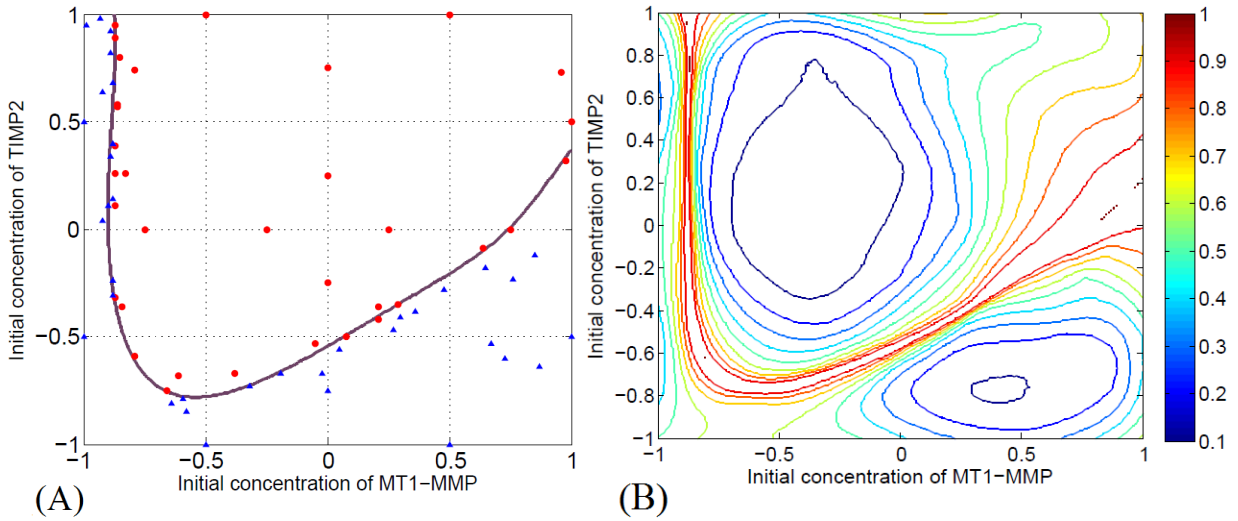

FIG. 2: Relative contribution of MT1-MMP and MMP2 to collagen degradation. The boundary separates points for which the contribution of MT1-MMP is dominant from points for which MMP2's contribution is dominant: (a) Design points and predicted boundary derived by the sequential sampling scheme. (b) A characterization of uncertainty in discrimination by variance. Notice that the points with high variance lie around the true boundary, which explains why the data sampled on the figure on the left also tend to focus around the true boundary.

one has in classification. From the figure, we also notice that the region with high variance encloses around the true boundary, illustrating the fact that after a burn-in period, the sequential sampling scheme only selects points that are close to the boundary.

4.3 A Model of the T-Cell Signaling Pathway

Our next example is a mathematical model of the T-cell signaling pathway proposed by Lipniacki et al. in [15]. This is a system of ODEs with 37 state variables, 19 parameters, and fixed initial conditions. We seek to design a sampling scheme to identify the boundary between the region of the parameter space where pERK, a state variable of the system, stabilizes at a high level of concentration, and the region at which pERK's concentration is less than a threshold level. In this example, the parameter space is defined relative to a nominal parameter vector. That is, for each component of the nominal vector, we define a range of 5 times smaller to 5 times larger than this component. We focus our attention on the eight most sensitive parameters determined by a global sensitivity analysis algorithm. The whole parameter space is the eight-dimensional set formed by the product of these 8 intervals. The time interval for analysis is $[0, 6000]$ (seconds).

For the sake of clarity in illustrations, we first investigate the performance of the algorithm in the case when all but the three most sensitive parameters are fixed. Starting with a prior data set collected at the first 44 points of a three-dimensional uniform sparse grid, we use the relaxed variance criteria to choose the next sample point until 144 data points are obtained. Figure 3(a) shows the boundary surface where the region above the surface corresponds to parameter values that will stabilize the concentration of pERK at a level higher than a fixed threshold. Similar to the previous example, samples selected by the sequential scheme tends to focus more and more along the boundary surface.

We then consider the case when eight most sensitive parameters are varied in the interval mentioned above. To evaluate the performance of of different sampling schemes in learning the structure of the boundary, we look at three different scenarios. In the first scenario, the Latin hypercube sampling is employed to collect 400 data points for inference. In the second scenario, starting with 100 data points chosen uniformly at random on the parameter space, the sequential sampling scheme is performed until 400 samples are obtained. Finally, to compare with the best possible performance, in the last scenario, we consider the ideal case when the sampling scheme is designed by an omniscient

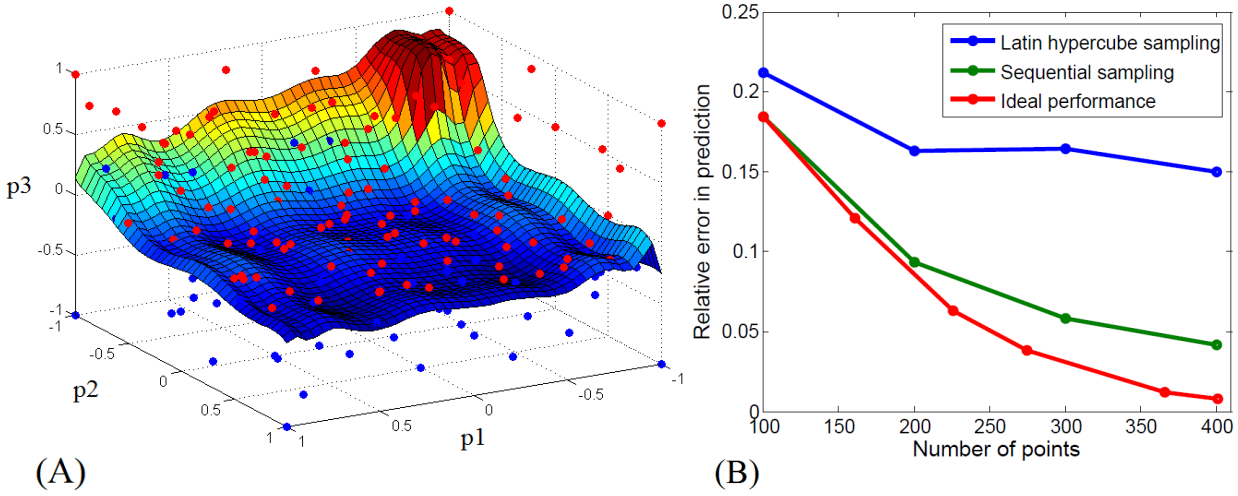

FIG. 3: A model of the T-cell signaling pathway with discrimination based on a threshold value for pERK at the final simulation time: (a) three-dimensional case: Design points and predicted boundary derived by the sequential sampling schemes with $N = 57$. (b) eight-dimensional case with $N = 321$: Error in prediction as the number of samples increase in three different scenarios: Latin hypercube sampling, sequential sampling scheme and oracular sampling.

oracle that at each step knows all the points for which current prediction is incorrect. Starting with 100 data points chosen uniformly at random (the same as in the second scenario), the oracle uses the estimated boundary as above to derive a expected prediction function. Upon testing the expected prediction function on thousands of samples, the oracle would add the misclassified samples into the training data set and continue the process.

Since the same algorithm of boundary computation is employed in all three scenarios, this examples provides a fair evaluation of the performance of these sampling schemes. The results are provided in Fig. 3(b). We can see that the sequential sampling scheme significantly outperforms the Latin hypercube sampling, and has comparable performance to the ideal case.

This result indicates that the sequential sampling scheme can effectively approximate a complex surface in 8-dimensional space with reasonable accuracy using only 400–500 samples. We note that this result is not likely to be produced by traditional methods for behavior discrimination. For example, if the discrimination is done by iteratively subdividing cubes that neither satisfy nor violate the property of interest into smaller hypercubes, every time a subdivision is performed, $2^8 - 1$ new elements are created and the number of samples needed to check if each of those elements satisfies/violates the property is almost as large as the number of samples we used in this example.

5. ADDITIONAL PROPERTIES

In this section, through 2D examples, we validate the theoretical results established in Section 3, as well as illustrate other properties of our algorithm. In most examples (except Example 5.3), the synthetic discontinuous output $u_1, u_2, ..., u_m$ is generated by evaluating the signum function along some discontinuity curve $r(x, y) = 0$:

$$u = \text{sign}\,(r(x, y))$$

Throughout this section, except example 5.2, the boundaries are computed by Griddy Gibbs Markov chains of 10^5 points on the coefficient space, while the number of terms in the approximation is set at 13 (i.e., we consider the first 13 terms in the sparse grid expansion of Γ).

5.1 Convergence

Figure 4 demonstrates the convergence of the algorithm in both low-discrepancy and sequential settings. In this experiment, data $u_1, u_2, ..., u_m$ are generated by evaluating the signum function along the discontinuity curve $r(x, y) = 0$:

$$u = \text{sign}\,(r(x, y)) \tag{18}$$

The parameter space in this case is $[-1, 1] \times [-1, 1]$, whereas the discontinuity function is described by $r(x, y) = y - (x - 1/4)^2$.

To compute the prediction error of each estimated classification, we evaluate the model at 10^6 points collected at random (uniformly) and compute the predicted results using one of the two sampling schemes. The blue curve corresponds to the case when data are sampled at the grid points of the sparse grid with uniform grid points, while for the red curve, data are collected using the sequential scheme with prior data given by the first 16 sparse gird points (low-discrepancy case). In both cases, we can see that when the number of sampling points increases, the error of prediction converges to zero. However, the convergence rate of the sequential scheme is significantly faster than that of the sparse grid sampling.

5.2 Dependence on the Number of Terms in Polynomial Expression

In Fig. 5(a), we consider a case when the discontinuity curve cannot be expressed as the level set of a polynomial function, in which

$$r(x, y) = y^2 - (x^3 - 2x - 1 - 3e^x) \tag{19}$$

and the input space is $[-6, 6] \times [-6, 6]$.

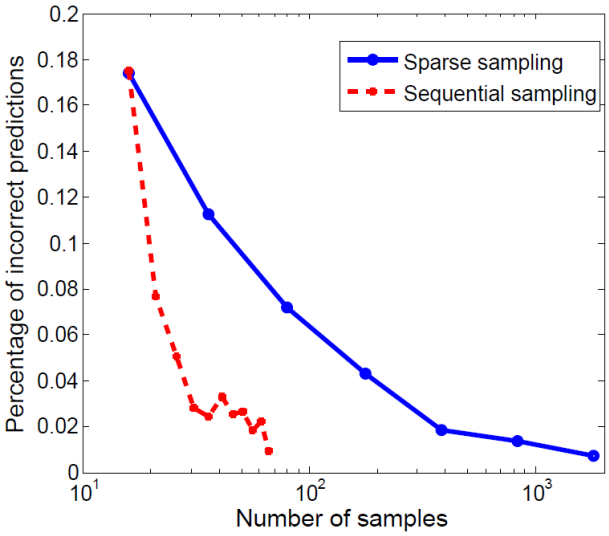

FIG. 4: Convergence rate of prediction error for the function in Eq. (18) when number of samples increases. In both cases, when the number of sampling points increase, the error of prediction converges to zeros. However, the convergence rate of the sequential scheme is significantly faster than that of the sparse grid sampling.

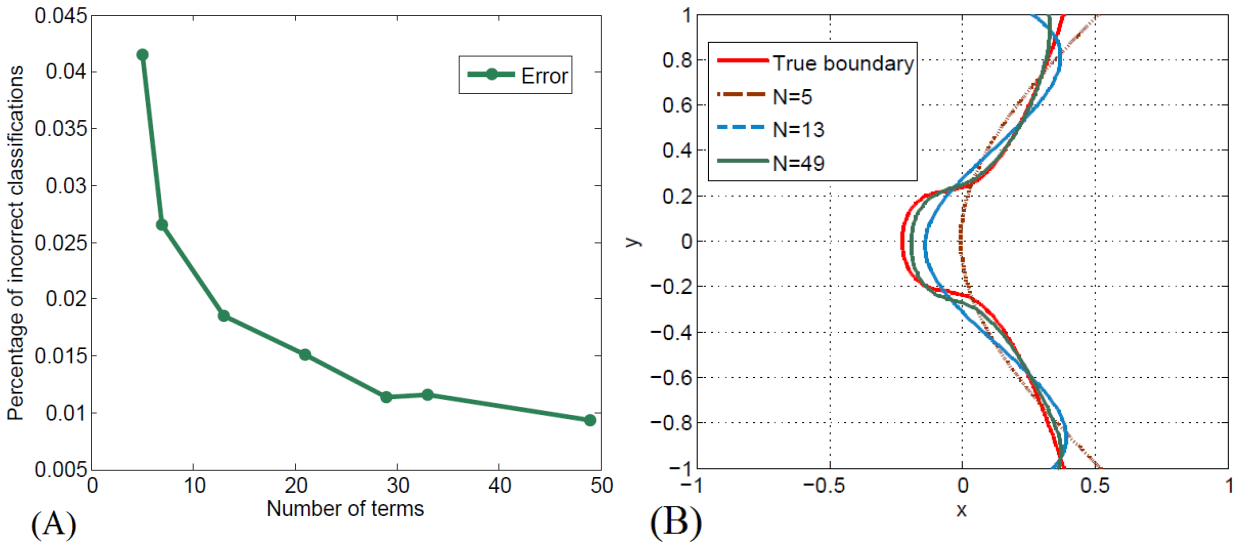

FIG. 5: (a) Decrease in prediction error when the number of terms (N) in the approximation increases for the function in Eq. (19). (b) Sample predicted boundary with different value of N. The boundaries are computed using 500 samples collected uniformly at random, while the error rates are estimated by the empirical prediction error on 10^6 uniformly distributed random points.

In this example, to reconstruct the boundary with high accuracy, the model is evaluated at 500 points, which are chosen uniformly at random in the input space. The algorithm to identify the boundary is employed with increasing number of terms in the polynomial approximation expressions. The expected boundaries were computed by a Markov Chain of 10^5 points using the Griddy Gibbs sampling method. The error rates were derived by the empirical prediction error on 10^6 points collected at random (uniformly). On the left panel, we plot the error rates in terms of the number

of basis functions used in the approximation. On the right panel, the approximated boundaries with various degree of approximation were illustrated. As expected, when the number of terms used in the approximation increases, the predicted boundary converges to the true boundary of discrimination.

5.3 Boundary with Multiple Components

We illustrate the fact that our algorithm can deal with cases when the boundary of interest has disconnected components. In this example, the model response is computed by the elliptic function

$$r(x,y) = \frac{y^2}{4} - \left(\frac{x^3}{8} - \frac{x}{2}\right) \tag{20}$$

on $[-2,2] \times [-2,2]$.

The constructed boundary is plotted in Fig. 6(a), using 200 sample chosen uniformly at random.

5.4 Robustness

In this particular example, we consider the behavior of our algorithm when some of the assumptions are not met. In this case, the synthetic discontinuous data $u_1, u_2, ..., u_m$ are generated by evaluating a bivariate error function with discontinuity strength parameter γ, discontinuity curve $r(x,y)$, and an additional global oscillatory structure with amplitude δ:

$$u_i = \text{erf}\left(\gamma(r(x_i,y_i))\right) + \delta \sin\left(\frac{\pi}{3}(y_i + x_i)\right) \tag{21}$$

This example violates (mildly) some conditions of our algorithm: (1) the response function is not discontinuous, but changes sharply across a curve, (2) the curve of discontinuity is not a perfect zero level set of any polynomial functions (but approximately is), and (3) the responses on both sides of the discontinuity curve are not flat, but have some additional oscillatory structure. Despite those violations, the function itself still resembles behavior of a yes/no response, and hence would work well with our algorithm for behavior discrimination.

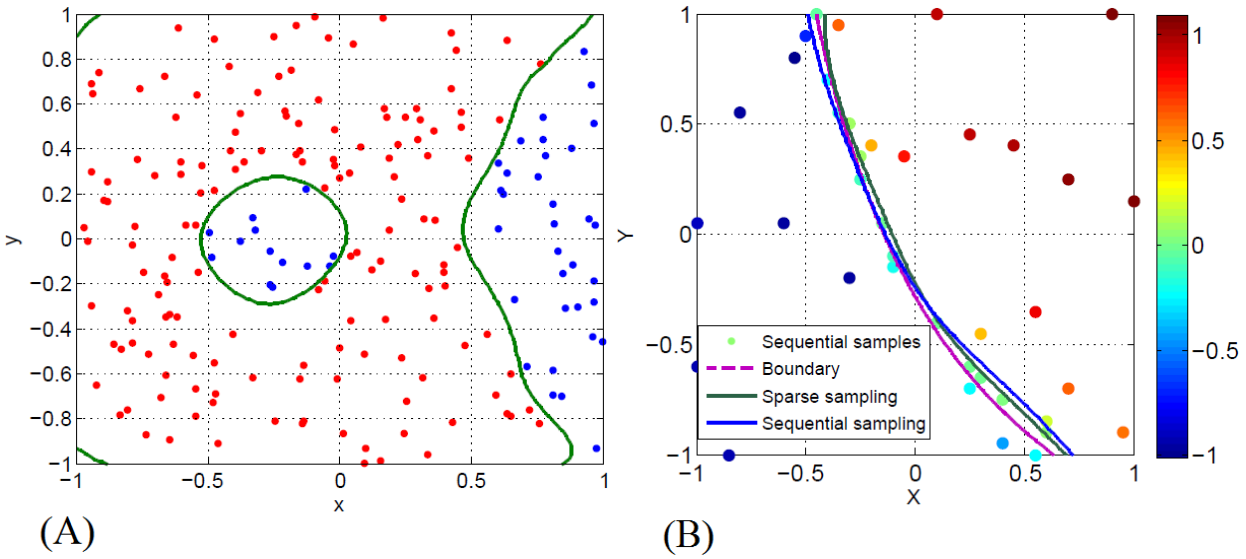

FIG. 6: (a) Example of a boundary with multiple components: discrimination between the region of positive and negative values of the elliptic function in Eq. (20) (b) Performance of the algorithm when multiple assumptions of the setting are mildly violated [Eq. (21)].

In Fig. 6(b), we replicate this example as used previously in [3], where the input parameter space is $[0.5, 6] \times [0, 2]$, while the discontinuity curve is a shifted and rescaled inverse function that goes through $(2, 2)$ and $(5, 0)$, i.e., $r(x, y) = y - 20/3\,(1/x - 1/5)$. The amplitude of the oscillatory structure is set to $\delta = 0.1$, and the steepness parameter is $\gamma = 2$. In the figure, the first 40 points generated by the sequential sampling are shown along with a expected boundary using 40 samples. We also computed the expected boundary derived by using first 80 nodes of a uniform sparse grid to make it convenient to compare the two methods. We can see that both methods of sampling provide good estimates of the boundary. The boundary is reconstructed with better quality than in [3], where data are collected at random locations. This example confirms the robustness of our method under perturbation of the system of interest. This feature is important for studies of reaction networks, where in some cases, the qualitative behavior of interest can only be determined within a certain level of confidence.

6. CONCLUSIONS AND DISCUSSION

In this work, we investigated the problem of choosing effective data sampling schemes for behavior discrimination of nonlinear systems. Using a probabilistic framework to estimate the boundary and quantify the uncertainty in discrimination, we give results about two classes of sampling schemes: the low-discrepancy sampling, and the sequential sampling. In both cases, we successfully derived theoretical results about the convergence of the expected boundary to the true boundary of interest. We then demonstrate the efficacy of the method in different application contexts. The method proves to be effective to study high-dimensional and expensive systems.

Nevertheless, there are some limitations to this method that are worth mentioning. First, although we are able to give an estimate of the convergence rate for the low-discrepancy setting, no such estimate is provided for the sequential setting (although we did provide a consistency result in that case). Second, one of our main assumptions in the paper is that the boundary of interest is well approximated by a finite polynomial expansion. In practice, however, the number of terms needed for approximation is difficult to be determined in advance. A more practical extension of the method should include an adaptive scheme to determine the number of terms in an effective way. These will be subjects of future work.

ACKNOWLEDGMENT

This research was partially supported by NSF Grant DMS-0900277.

REFERENCES

1. Summers, S., Jones, C. N., Lygeros, J., and Morari, M., A multiresolution approximation method for fast explicit model predictive control, *IEEE Trans. Automatic Control*, 56(11):2530–2541, 2011.
2. Chakrabarty, A., Dinh, V., Buzzard, G. T., Zak, S. H., and Rundell, A. E., Robust explicit nonlinear model predictive control with integral sliding mode, In *American Control Conference (ACC)*, pp. 2851–2856, IEEE, 2014.
3. Sargsyan, K., Safta, C., Debusschere, B., and Najm, H., Uncertainty quantification given discontinuous model response and a limited number of model runs, *SIAM J. Sci. Comput.*, 34(1):B44–B64, 2012.
4. Donzé, A., Clermont, G., and Langmead, C. J., Parameter synthesis in nonlinear dynamical systems: Application to systems biology, *J. Comput. Biol.*, 17(3):325–336, 2010.
5. Donzé, A., Fanchon, E., Gattepaille, L. M., Maler, O., and Tracqui, P., Robustness analysis and behavior discrimination in enzymatic reaction networks, *PLoS One*, 6(9):e24246, 2011.
6. Niederreiter, H., Quasi-Monte Carlo methods and pseudo-random numbers, *Bull. Am. Math. Soc.*, 84(6):957–1041, 1978.
7. Freund, Y., Mansour, Y., and Schapire, R. E., Generalization bounds for averaged classifiers, *Ann. Stat.*, pp. 1698–1722, 2004.
8. MacKay, D. J., Introduction to Monte Carlo methods, In *Learning in Graphical Models*, pp. 175–204, Springer, Berlin, 1998.
9. Settles, B., Active learning literature survey, *Computer Sciences Technical Report 1648*, University of Wisconsin, Madison, 2010.
10. Bierstone, E. and Milman, P. D., Semianalytic and subanalytic sets, *Publications Math. l'IHÉS*, 67(1):5–42, 1988.

11. Pronzato, L., Asymptotic properties of nonlinear estimates in stochastic models with finite design space, *Stat. Probability Lett.*, 79(21):2307–2313, 2009.
12. Pronzato, L., One-step ahead adaptive D-optimal design on a finite design space is asymptotically optimal, *Metrika*, 71(2):219–238, 2010.
13. Reynolds, A., Rubin, J., Clermont, G., Day, J., Vodovotz, Y., and Bard Ermentrout, G., A reduced mathematical model of the acute inflammatory response: I. Derivation of model and analysis of anti-inflammation, *J. Theor. Biol.*, 242(1):220–236, 2006.
14. Karagiannis, E. D. and Popel, A. S., A theoretical model of type I collagen proteolysis by matrix metalloproteinase (MMP)2 and membrane type 1 MMP in the presence of tissue inhibitor of metalloproteinase 2, *J. Biol. Chem.*, 279(37):39105–39114, 2004.
15. Lipniacki, T., Hat, B., Faeder, J. R., and Hlavacek, W. S., Stochastic effects and bistability in T cell receptor signaling, *J. Theor. Biol.*, 254(1):110–122, 2008.