Structural Safety 62 (2016) 1–11

Contents lists available at ScienceDirect

Structural Safety

journal homepage: www.elsevier.com/locate/strusafe

Surrogate-enhanced stochastic search algorithms to identify implicitly defined functions for reliability analysis

V.S. Sundar, Michael D. Shields *

Department of Civil Engineering, The Johns Hopkins University, Baltimore, MD 21218, USA

ARTICLE INFO

Article history:
Received 30 August 2015
Received in revised form 2 May 2016
Accepted 6 May 2016
Available online 20 May 2016

Keywords:
Reliability
Response surface
Meta-model
Surrogate model
Monte Carlo simulation
Artificial neural networks
Stochastic search

ABSTRACT

The paper proposes two adaptive stochastic search algorithms to locate and trace an implicitly defined function with samples that are used to construct a surrogate model for reliability analysis. Both methods begin by propagating a series of surrogate-informed stochastic processes toward the implicit performance function. Having located the function, the first method conducts a "global" tracing of the function while the second traces the function by propagating samples locally. A key feature of both proposed algorithms is that the surrogate model evolves continuously with the sample selection and is used to inform the selection of new samples such that it converges rapidly to an accurate representation of the limit surface. In the present implementation, an artificial neural network surrogate model is employed but the method can, in principle, be applied with any surrogate model form. Performance of the algorithms is illustrated through problems with highly nonlinear limit state functions and high dimensional non-Gaussian random variables.

1. Introduction

With recent advancements in computational power, numerical tools are now, more than ever, being widely applied to solve complex problems in engineering and the sciences. However, the problem of evaluating reliability, or estimating probability of failure, for complex structures and systems remains intractable for many problems given the large computational expense associated with model evaluations and the need to evaluate the model a very large number of times. The advancement of variance reduction strategies such as importance sampling [1,2], stratified sampling [3,4], subset simulation [5], and line sampling [6] has made sample-based approaches increasingly tractable but these methods often still require an unacceptably large number of simulations or possess undesirably high coefficient of variation [7]. One way of addressing the problem of computational expense is to develop an explicit mathematical model that approximates the computer code, is computationally efficient, and accurately reflects the most important properties of the response. Broadly referred to as surrogate models, meta-models, or response surfaces, they need to be tailor made to suit specific problems (see, for example, [8,9]) and may then be used to estimate statistical properties of the desired quantity. Over the past twenty years, there have been numerous studies proposing new methodologies to construct surrogate models using principles of machine learning [10] (artificial neural networks [11–13] and support vector machines [14–17]), improvements to the existing response surface techniques (using, for example, high-order polynomial approximations [18] or accounting for multiple failure domains [19]) and Gaussian process/Kriging models [20,21]. A review of some available methods, their merits and demerits can be found in [22].

A key feature of most recent surrogate model developments is adaptivity in the sample point selection. For all surrogate model forms, adaptive methods have been employed at some level to identify "optimal" sample points for the construction of the model. In this spirit, Guan and Melchers [23] have highlighted the sensitivity of polynomial response surfaces to the selection of sample points – illustrating the importance of their appropriate selection. For conventional (first and second order) polynomial response surfaces, numerous adaptive sample design approaches have been proposed beginning with Kim and Na [24] who proposed a sequential method using the gradient projection method to place points close to the true limit surface. Following a similar approach, Zheng and Das [25] adaptively add complexity (second order square and interaction terms) to identify a suitable response surface approximation. Nguyen et al. [26] proposed a method that uses complementary experimental designs to adaptively select the support points and fits the response surface using a weighted regression where weights are determined based on the distance from the limit surface and the sample point [27]. Steenackers et al. [28]

* Corresponding author.
E-mail addresses: svelkur2@jhu.edu (V.S. Sundar), michael.shields@jhu.edu (M.D. Shields).

http://dx.doi.org/10.1016/j.strusafe.2016.05.001
0167-4730/

integrated an adaptive response surface directly into an computation driven design optimization. In the context of support vector machines, Richard et al. [16] utilize a rotated basis for the experimental design that is constructed from gradients of the surrogate model while Dai et al. [17] utilize an adaptive Markov Chain Monte Carlo (MCMC) approach to draw samples from the target optimal importance sampling density [29]. Meanwhile, for Kriging models Echard et al. [30] utilize a learning function to select the optimal point to add to the sample design based on considerations from the current state of the Kriging surrogate. Recently, Balesdent et al. [31] coupled a Kriging-based surrogate with adaptive importance sampling to estimate failure probabilities.

The proposed approach belongs to this general family of adaptive surrogate model designs. The basic idea is to utilize the surrogate model adaptively to obtain samples that "trace" the implicitly defined function. We employ a two-step strategy in which the first step involves a surrogate-enhanced search of the space to identify the failure domain(s). This search utilizes "pseudo-Markov Chains" designed to propagate samples toward the failure domain. To do so, a Markov Chain Monte Carlo-like algorithm is applied wherein the conventional Metropolis–Hastings acceptance/rejection criterion is replaced with a surrogate-based criterion that accepts a sample if the surrogate model shows the new sample decreases the performance function and rejects the sample otherwise. The surrogate model is updated after every step of the chain and the chains are propagated until each crosses the limit surface.

In the next step, two methods – a global method and a local method – are proposed to trace the limit surface. The global method defines a band around the surrogate limit surface and performs random draws from the known probability densities conditioned upon existence within the band. From these conditional random samples, the point is selected that maximizes a performance-adjusted distance from all other existing surrogate design points. The local method utilizes a similar point selection methodology but on a locally confined domain surrounding each individual surrogate design point. Again, the surrogate model is updated after each design point selection to gain increasing accuracy with the iterations. Both methods are applied to a set of benchmark reliability problems possessing different features such as strong nonlinearity, multiple failure domains, non-Gaussianity, and moderately high dimension. The methods are demonstrated to produce very accurate surrogate models in a small number of sample points. Given the heuristic nature of these algorithms, some discussion of each methodology is provided to highlight their relative strengths and weaknesses and address practical considerations for their implementation.

As an additional consideration, there seems to be no general strategy in the existing literature for determining the appropriate type of surrogate model to be employed for a given problem. In the present implementation and example problems, an artificial neural network surrogate model is employed. However, to the extent possible, the methodology is presented such that it is independent of the surrogate model form as the method is (at least in principle) applicable to any general form of surrogate.

2. Surrogate models for reliability analysis in structural mechanics

Consider that all the uncertainties associated with a structural system are quantified in terms of an $n-$dimensional random variable $\mathbf{X}$ with joint-probability density function (pdf) $p_{\mathbf{X}}(\mathbf{x})$. When random fields or random process models are used, it is assumed that these can be adequately represented in terms of random variables through a suitable discretization scheme [32,33]. The probability of failure with respect to a performance function, $g(\mathbf{X})$, is computed through the n-fold integral:

$$P_F = \mathrm{P}[g(\mathbf{X}) \leqslant 0] = \int_{g(\mathbf{x})\leqslant 0} p_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} \tag{1}$$

Here $g(\mathbf{X}) = 0$ is called the limit surface, and $g(\mathbf{X}) < 0$ denotes the unsatisfactory performance of the structural system. Solving Eq. (1) is often not straightforward due to (i) the correlated non-Gaussian nature of the random variable $\mathbf{X}$, (ii) partial specification of $p_{\mathbf{X}}(\mathbf{x})$, (iii) the non-linear nature of the performance function, $g(\mathbf{X})$, and (iv) difficulty in evaluating $g(\mathbf{X})$, or unavailability of $g(\mathbf{X})$ in explicit form. Typically Eq. (1) is solved by converting the correlated non-Gaussian random variables into uncorrelated standard normal random variables and posing the reliability problem in the standard normal space [34]. Given the often complex (nonlinear, discontinuous, etc.) and implicit nature of $g(\mathbf{X})$, approximate or simulation-based approaches are typically employed to solve Eq. (1) [29,35]. In such cases, the classical Monte Carlo estimator for P_F is given by:

$$\hat{P}_F = \frac{1}{N}\sum_{i=1}^{N} I[g(\mathbf{x}_i) \leqslant 0] \tag{2}$$

where $I[\bullet]$ is the indicator function.

Considering the large computational expense associated with the estimate in Eq. (2), it is often desirable to construct a surrogate model for the performance function $g(\mathbf{X})$. A surrogate model is an explicit mathematical expression that is computationally inexpensive to evaluate and is designed to approximate the implicitly defined function. This surrogate model is denoted $\hat{g}(\mathbf{X})$ and is considered sufficiently accurate if

$$|g(\mathbf{x}) - \hat{g}(\mathbf{x})| \leqslant \epsilon, \quad \forall \mathbf{x} \in \mathcal{S} \tag{3}$$

where $\mathcal{S}$ denotes a subset of the sample space over which the response is deemed to be important, and ϵ is a suitable accuracy threshold. Having established the surrogate model, the Monte Carlo estimator for P_F can be obtained using the approximate function as

$$\bar{P}_F = \frac{1}{N}\sum_{i=1}^{N} I[\hat{g}(\mathbf{x}_i) \leqslant 0] \tag{4}$$

For problems concerning determination of reliability as considered here, the strict constraint given in Eq. (3) need not necessarily be enforced. A sufficient condition for acceptance of the surrogate model can be established such that $g(\mathbf{x}) < 0 \Longleftrightarrow \hat{g}(\mathbf{x}) < 0$, and $g(\mathbf{x}) > 0 \Longleftrightarrow \hat{g}(\mathbf{x}) > 0,\ \forall \mathbf{x}; \mathbf{x} \sim p_{\mathbf{X}}(\mathbf{x})$.

Several types of surrogate models have been proposed for approximate reliability analysis with the most commonly utilized methods being polynomial response surfaces, Kriging or Gaussian process models, support vector machines, and artificial neural networks (ANNs). Utilizing a polynomial response surface, the true limit surface $g(\mathbf{X}) = 0$ is approximated by a kth-order polynomial with the coefficients estimated from regression. Much of the research related to polynomial response surfaces has concentrated on identifying the "best" experimental design that balances the restrictions of large computational cost associated with performance evaluation and accuracy in the coefficient estimation. In recent years, this has inspired numerous adaptive methods (e.g. [24–26]). The Kriging or Gaussian process surrogate model builds upon the polynomial response surface method by considering an explicit mathematical model (usually constructed from orthogonal polynomial basis functions [21]) to be the mean of a stationary and Gaussian stochastic process with assumed covariance structure [20]. The coefficients associated with the basis functions are then estimated using regression analysis or Bayesian methods.

A different class of methods are those using principles of machine learning such as support vector machines or artificial neural networks (which are utilized herein). Support vector machines [36] are binary classification tools, wherein a surrogate hyperplane