12th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP12
Vancouver, Canada, July 12-15, 2015

Targeted Random Sampling for Time-invariant Reliability Analysis

Michael D. Shields
Assistant Professor, Department of Civil Engineering, The Johns Hopkins University, Baltimore, MD, USA

V. S. Sundar
Post doctoral researcher, Department of Civil Engineering, The Johns Hopkins University, Baltimore, MD, USA

ABSTRACT: The Targeted Random Sampling (TRS) method is extended to moderate-to-high dimensional time invariant reliability analysis. TRS is an adaptive sampling method rooted in stratified sampling variance reduction technique and represents a specific algorithm for Refined Stratified Sampling. In this work, a Markov Chain is used to sample the space densely in the vicinity of the limit stat and multi-dimensional stratum division is proposed along with a neural network based approximation of the performance function. The method is shown to perform efficiently and accurately for estimation of probability of failure for a 20 dimensional problem.

In a probabilistic reliability analysis, the uncertainties associated with the external loading, structural and material properties are quantified in terms of a $n-$dimensional random vector $\boldsymbol{X}$ with joint probability density function (pdf) $p_{\boldsymbol{X}}(\boldsymbol{x})$. In general, the components of the random vector are taken to be mutually correlated and non-Gaussian.When random fields or random process models are used, it is typically assumed that these can be adequately represented in terms of random variables through a suitable discretization scheme (Ghanem and Spanos (1991); Shinozuka and Deodatis (1991); Sudret and Der Kiureghian (2000)). Defining a function $g(\boldsymbol{X})$ that serves as a metric for structural performance such that $g(\boldsymbol{X}) < 0$ represents unsatisfactory performance of the structure under consideration and $g(\boldsymbol{X}) = 0$ is referred to as the limit surface separating the "safe" and "failure" domains, the probability of failure of the structural system with respect to $g(\boldsymbol{X})$ is defined as

$$\begin{aligned} P_F &= \int_{g(\boldsymbol{x})\leq 0} p_{\boldsymbol{X}}(\boldsymbol{x})d\boldsymbol{x} \\ &= \int_{\mathbb{R}^n} I[g(\boldsymbol{x}) \leq 0] p_{\boldsymbol{X}}(\boldsymbol{x})d\boldsymbol{x} \\ &= \Big\langle I[g(\boldsymbol{x}) \leq 0] \Big\rangle \end{aligned} \tag{1}$$

where $I[\bullet]$ is the indicator function and $< \bullet >$ is the expectation operator.

There exist broadly two alternative approaches for solving Eq. (1), namely, Taylor series expansion-based methods utilizing reliability indices and sampling-based methods (e.g. Monte Carlo simulation - MCS) (Melchers (1999)). In the first class of methods, the performance function, in the standard normal space, is expanded as a first order Taylor series (first order reliability method, FORM) and the shortest distance from the origin to the limit surface is taken as a measure of reliability - termed as the Hasofer-Lind reliability index, β_{HL} (Hasofer and Lind (1974)). Denoting by $\boldsymbol{U}$ the vector of random variables in the standard normal space and by $G(\boldsymbol{U})$ the corresponding performance function, the problem of reliability analysis is posed as: minimize $\sqrt{\boldsymbol{U}^t\boldsymbol{U}}$ subject to $G(\boldsymbol{U}) = 0$. The probability of failure is then expressed as $P_F = \Phi(-\beta_{HL})$, where $\Phi(\bullet)$ is the $n-$dimensional standard normal pdf.

The transformation to the standard normal space from the basic random variable space is carried out either using the Nataf transformation or Rosenblatt transformation (Melchers (1999), Madsen et al. (2006)). For structures with simple limit states,

these approaches also lead to the determination of a design point (e.g. the point of most probable failure) and measures of sensitivity for failure probability with respect to the basic random variables. An important feature of these methods is that the problem of evaluation of a multi-dimensional integral over an irregular domain is posed as an equivalent problem in constrained non-linear optimization. In nearly all practical applications, however, the method remains approximate in nature and they produce poor approximations for problems with complex (e.g. nonlinear or discontinuous) limit states. These are the problems of specific interest in this work.

On the other hand, sampling-based methods utilizing MCS produce a point estimate of probability of failure that is exact as the sample size becomes large. The basic idea with MCS is to simulate realizations of the random vector $\boldsymbol{X}$ according to the prescribed joint pdf and use statistical tools to estimate the probability of failure. Denoting a set of N realizations of the basic random variable vector by $\{\boldsymbol{x}^{(i)}\}_{i=1}^{N}$, an estimate for the probability of failure is obtained as

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

It can be shown that $\hat{P}_F$ is an unbiased and consistent estimator for P_F with variance being equal to $P_F(1-P_F)/N$ (Rubinstein and Kroese (2008)). It should be noted that the sampling variance is independent of n, the size of the random vector $\boldsymbol{X}$. Also, by virtue of the strong law of large numbers, it can be shown that $\hat{P}_F \to P_F$ with probability 1 as $N \to \infty$. It is clear that the variance of the estimator is inversely proportional to the sample size; hence in order to estimate a probability of failure as low as 10^{-4}, a minimum of 10^5 to 10^6 samples are required. This poses severe computational issues, especially if the performance function is defined implicitly using, for example, a finite element code, as the computational time for one structural analysis may be huge depending on the size and type of problem being solved.

On the footings laid by the FORM and MCS, several improvements have been proposed. In case of reliability index based methods, the focus has been on accounting for multiple design points and implicitly defined performance functions(e.g. Bucher and Bourgund (1990), Der Kiureghian and Dakessian (1998)). Improving the performance of MCS equates to reducing the variance of the statistical estimator given in Eq. (2), i.e. obtaining a better estimate with a smaller sample size. Generally such methods are referred to as variance reduction techniques.

Numerous variance reduction methods have been developed for reliability analysis - the most widely used being importance sampling (Engelund and Rackwitz (1993)), subset simulation (Au and Beck (2001)), and line sampling (Koutsourelakis et al. (2004)). While these methods produce significant improvement over classical MCS, they often produce estimates with large coefficient of variation (Schuëller and Pradlwarter (2007)). The motivation for the Targeted Random Sampling (TRS) methodology proposed in this work is to produce sample-based reliability estimates for moderate-dimensional problems with complex limit states in a small number of samples with reduced coefficient of variation when compared to existing methods.

The TRS method (Shields and Sundar (2014)) is based upon the application of a variance reduction technique called stratified sampling to reliability analysis. In particular, the Refined Stratified Sampling (RSS) method developed by Shields et al. (2014) - wherein the strata that decompose the probability space are divided to efficiently add samples to a stratified design - underpins the TRS method. The method presents a distinct break from classical Monte Carlo approach in certain regards. In particular, it operates by decomposing the probability space of the random vector into a space-filling set of disjoint regions (strata) that are defined from a carefully selected ("targeted") set of random samples (i.e. samples are targeted in the vicinity of the limit state). Probability of failure is then estimated by determining the total volume associated with strata in the failure domain. In other words, rather than utilizing classical unbiased statistical estimators (e.g. Eq. (2) where the denominator N must be

very large), the method operates by targeting samples in such a way that the limit state can be approximated through an efficient partitioning of the probability space.

The focus of the present work is on extending the original TRS method to problems in high dimensions. The method is illustrated for a 20 dimensional time invariant reliability problem with explicitly defined limit state function. Note, however, that the method provides a generalized framework for any complex limit state (explicit or implicit). We begin with a brief background of the Refined Stratified Sampling concepts upon which TRS is constructed. A few drawbacks of the original TRS method are discussed along with possible solution strategies. An algorithmic description of the proposed multicut TRS method is then provided, followed by numerical illustration and conclusions.

1. STRATIFIED AND REFINED STRATIFIED SAMPLING

Stratified sampling operates by dividing the probability space Ω into a set of M subdomains (strata) Ω_i that are space filling ($\bigcup_{i=1}^{M} \Omega_i = \Omega$) and disjoint ($\Omega_i \cap \Omega_j = \emptyset$). Samples are then drawn from within each of the strata and the quantity of interest, Θ, is estimated by

$$\hat{\Theta} = \sum_{i=1}^{M} \frac{p_i}{M_i} \sum_{j=1}^{M_i} y_{ij} \tag{3}$$

where p_i is the volume of the i^{th} stratum, M_i is the number of samples drawn from the i^{th} stratum, and y_{ij} is the response quantity evaluated from the j^{th} sample drawn from stratum i.

In this setting, the Refined Stratified Sampling method can be employed to add samples to sparsely sampled regions of the space by dividing specifically selected strata to obtain a more desirable distribution of samples in the space. In its initial implementation (Shields et al. (2014)), an initial stratification of the space is provided with a single sample drawn from each stratum. Strata are then selected for division based upon their associated probability weights (i.e. larger strata are divided first). To clarify this, consider N number of $n-$dimensional vectors scattered in a $n-$dimensional space. Assume that there exists an initial rectilinear stratification. The volume of the strata are given as $p_i = \prod_{j=1}^{n} L_i^j$, where L_i^j is the length of stratum i in the j^{th} dimension. The stratum corresponding to $p_{max} = \max_{1 \le k \le M} (p_i)$ is chosen as the stratum to divide. An orthogonal direction is defined as $\{j : L_j = \max_{1 \le k \le p} L_k\}$. The stratum is now divided in half along this direction and a sample is drawn randomly from the newly created empty stratum. In this way, the RSS method creates an even distribution of samples across the space and has been proven to reduce variance when compared to the method of simply adding samples to existing strata without stratum division. Additional details can be found in Shields et al. (2014).

In general, an even distribution of samples throughout the space may reduce variance for certain statistics (like estimates of the mean). However, for reliability analysis the desire is to produce samples that are concentrated in the vicinity of the limit state. This can be accomplished by modifying the RSS algorithm to divide strata according to a combination of the stratum size and the value of the performance function $g(\mathbf{X})$. That is, the stratum to break is chosen to be one with highest probability content amongst the strata that lie in the vicinity of the limit surface. Specifically, the TRS method represents a specific RSS algorithm that divides strata that are near the limit state in proportion to the stratum size (i.e. large strata in the vicinity of the limit state are divided first). In this way, the definition of the strata are refined to partition the space into "failure strata" and "safe" strata. It is to be noted that the stratification is performed in the probability space where the random variables are independent and uniformly distributed in $[0,1]$, thus calculating volumes of rectilinear strata in high dimensions do not possess a problem. Also, it is assumed that there exists a transformation, T, that relates the basic correlated non-Gaussian random variable space to the probability space, i.e. $\boldsymbol{X} = \mathrm{T}(\boldsymbol{Z})$, and vice-versa. Here $\boldsymbol{Z} \sim \mathrm{U}[0,1]$ (e.g. Nataf Transformation).

2. TARGETED RANDOM SAMPLING

The TRS method detailed in Shields and Sundar (2014) is briefly described below. The drawbacks are discussed and an improved algorithm that overcomes these deficiencies is presented in the next section. The steps for the implementation of the TRS method are as follows:

- Produce an initial stratified design. Note that this initial stratified design must possess at least one sample in each disjoint failure region.
- Identify all "fail-safe" pairs composed of points adjacent to one another but on opposite sides of the limit state. The strata corresponding to the "fail-safe" pairs make up the so-called target sampling subdomain. This is the domain from which future sample are drawn.
- Select the stratum from the target sampling subdomain that has the maximum probability content (i.e. largest volume). Divide this stratum in one direction by interpolating to identify the approximate location of the limit state.
- Generate a random sample from the newly defined empty stratum.
- The probability of failure is reported as the sum of the probability content (strata volumes) associated with all the failure samples.
- Repeat until satisfactory convergence is achieved.

2.1. Drawbacks

A few shortcomings of the original TRS method are discussed along with the solution strategies.

2.1.1. Locating the failure region

The TRS method requires at least one sample to lie within each disjoint failure region in order to initiate the algorithm. Obtaining these samples may not be trivial task, possibly due to the implicit and high non-linear nature of the limit state function. Intuitively, one may sample from the tails of the pdf's of the random variables. However, such intuition fails to address problems whose failure domains are not associated with the rare events in a given random vector. Instead, failure may be associated with some combination of undesirable values of random variables - that need not lie in the tails of the distribution. This is certainly the case in the "failure island" scenario - a case where the limit surface forms a closed region in the probability space. Additionally, this strategy may be computationally inefficient when applied to high dimensional problems because the number of strata grow exponentially with dimension. Consider, for example, a simple stratification where each dimension is cut once to produce two strata. In 2-dimensions, this produces $2^2 = 4$ strata. But in 20-dimensions, this produces $2^{20} = 1,048,576$ strata - clearly an unreasonable initial stratification.

To identify an initial sample set possessing the requisite samples in the failure domain, it may be possible to use the end results from a FORM analysis as a starting point. Using the design point obtained from FORM, importance sampling may be performed to obtain a small number of samples (say 10 or less) in the failure domain. However, this does not address the challenge of identifying multiple failure domains. To solve this problem in a robust and computationally efficient manner, we propose to initialize multiple Markov Chains in the standard normal space that are designed to converge to the failure domain - much like the approach used in the subset simulation method. Multiple chains are initiated in order to facilitate convergence to multiple failure domains. This process is described in Section 3.1.

2.1.2. Curse of dimensionality

For high dimensional problems, dividing strata in a single dimension, as proposed in the original TRS method results in slow convergence due to the poor resolution of the performance function at the limit surface. One way to overcome this problem is by employing simultaneous cuts along all the dimensions. Though this approach seems intuitive, the associated computational cost is high because at each cut, the number of times we need to sample (and thus the number of performance function evaluations) is often equal to the dimension of the problem. In order to alleviate this burden, an approach utilizing neural networks is used. The neural network is trained and validated using the samples obtained as states of the initial Markov Chains. Once the neural network is trained, the performance function value at each of the added samples is approx-

imated by the trained neural network. Additional "real" samples can be strategically placed within large strata where it is expected that the neural network may be poorly resolved.

3. MULTI-CUT TARGETED RANDOM SAMPLING

This section provide a step-wise procedure for implementing the multi-cut TRS algorithm designed to address the drawbacks of the original TRS method outlined above.

3.1. Locating failure domains

A set of n_c suitably designed Markov Chains are used to identify and explore the failure domain. The Markov Chain is constructed in such a way that the forward propagation of the chain would imply that the chain is marching toward the failure domain, i.e. if $\bar{g}_1, \bar{g}_2, \cdots, \bar{g}_i$ denote the performance function values evaluated at the states of the chain, then $\bar{g}_1 > \bar{g}_2 > \cdots > \bar{g}_i$. Denote the number of Markov Chains by n_c. For each of the Markov Chain, starting from points $\boldsymbol{u}_i; i = 1, 2, \cdots, n_c$, run the chain such that n_f number of samples are obtained near the vicinity of the limit surface. The propagation of the Markov Chain from i^{th} state to $(i+1)^{\text{th}}$ state is provided below:

- Choose a proposal density function $q(\bullet)$, say $N(\boldsymbol{x}_i, \boldsymbol{\sigma})$. Let $g(\boldsymbol{x}_i) = g_i$. The proposal density function is, in general, chosen to be a pdf which is easy to sample from; for example, normal or uniform. Choosing a normal/uniform pdf as proposal simplifies the calculation of the acceptance ratio, α, due to their symmetric form(Rubinstein and Kroese (2008)).
- Sample $\boldsymbol{y} \sim q(\boldsymbol{x}_i)$
- Calculate the acceptance probability $\alpha = \min[1, \phi(\boldsymbol{y})/\phi(\boldsymbol{x}_i)]$ where $\phi(\bullet)$ is the $n-$dimensional standard normal pdf.
- Evaluate $g_y = g(\boldsymbol{y})$.
- Generate $z \sim$U[0,1].
- Acceptance criteria prior to the Markov chain crossing the limit surface (i.e. $g(\boldsymbol{x}_i) > 0 \; \forall i$)
$$\boldsymbol{x}_{i+1} = \begin{cases} \boldsymbol{y}, & \alpha > z \text{ and } g_y < g_i \\ \boldsymbol{x}_i, & \text{otherwise.} \end{cases}$$
- Acceptance criteria after the Markov chain crosses the limit surface
$$\boldsymbol{x}_{i+1} = \begin{cases} \boldsymbol{y}, & \alpha > z \text{ and } g_f < g_{i+1} < g_s \\ \boldsymbol{x}_i, & \text{otherwise.} \end{cases}$$
where g_s is the value of the performance function at the sample state prior to the Markov Chain leaving the safe domain for the first time, and g_f is the performance function value of the sample state after the Markov Chain crosses the safe domain for the first time.
- Through a suitable post-stratification technique (here using rectilinear grids), assign strata to each of the samples generated from the Markov Chains. One possible methodology is explained in section 4.

Each Markov Chain initialized above is design to converge toward the limit state. Once the limit state is crossed, the Markov Chain is redesigned to concentrate samples in the immediate vicinity of the limit state. In this way, the samples converge to the failure domain and then "trace" its boundary. In any case, the motivation behind this step is to obtain as many number of samples possible near the vicinity of the limit surface. Post-stratification of the samples allows the volume of the failure domain to be assessed and an initial estimate of probability of failure to be produced.

3.2. Multiple cuts

The process for utilizing multiple cut planes for stratum refinement is now described. In the same manner as the original TRS algorithm, identify "fail-safe" sample pairs and the associated strata (i.e. the target sampling subdomain). From these pairs, let $x_a^f, a = 1, 2, \cdots, n_f$, and $x_b^s, b = 1, 2, \cdots, n_s$ denote the n_f samples in the failure domain and n_s samples in the safe domain respectively. Denote their corresponding strata by $\Omega_a^f, a = 1, 2, \cdots, n_f$ and $\Omega_b^s, b = 1, 2, \cdots, n_s$. For the sake of demonstration, consider a failure point x^f whose stratum is adjacent to k strata corresponding to safe points $x_i^s, i = 1, 2, \cdots, k$. For each $x_i^s, i = 1, 2, \cdots, k$ let d_i denote the dimension along which its stratum shares a boundary with the stratum possessing to x^f. For each pair of x^f and x_i^s, the stratum pair is divided along the dimension d_i at the interpo-

lated approximate location of the limit state function $g(\boldsymbol{x}) = 0$. This creates k new strata from which new random samples are drawn. As previously, the probability of failure is determined as the sum of the strata weights corresponding to all failure samples. The iterations are terminated after a satisfactory level of convergence in the desired estimate is observed.

3.3. Neural network-based sampling

The above multi-cut TRS method necessitates k new samples for every existing sample in the failure domain - which will be computationally quite expensive. To mitigate this cost, an artificial neural network is trained to approximate the performance function. Using this artificial neural network, it is possible to approximate the performance function at the randomly generated points in the new strata - thus replacing the expensive performance function evaluation through, for example, a finite element model.

In general, the neural network based performance function approximation will be very accurate in the vicinity of the limit state given the concentration of samples in this important region produced from the Markov Chain. The neural network is likely to be far less accurate in regions some distance from the limit state. However, given the nature of the stratification process - specifically the strata refinement in the vicinity of the limit state only - the neural network will not be utilized for these inaccurate regimes. Additionally, sparsity of samples in certain regions near the limit state may cause local inaccuracies in the neural network. In such cases, it is possible to produce a small number of true performance function evaluations to locally refine the approximation and improve the strata refinement.

4. NUMERICAL ILLUSTRATION

The proposed multi-cut TRS algorithm is illustrated through a $20-$dimensional problem (Engelund and Rackwitz (1993)). The performance function is defined as

$$g(\boldsymbol{X}) = \sum_{i=1}^{20} X_i + C \tag{4}$$

where $C = -5.343$ producing a convex limit surface. The probability of failure is defined as $P_F = \mathrm{P}[g(\boldsymbol{X}) \leq 0]$.

Twenty Markov Chains are initiated in order to obtain approximately 400 samples in the vicinity of the limit surface. Around 70% of these samples are used to train the neural network and the remainder are used for validation and testing. On average, between 700-1000 total performance function evaluations are required to obtain these samples - including those necessary for convergence to the limit state and those rejected by the chain. Post-stratification of the Markov Chain sample set is performed by adding each sample sequentially. As a sample is added, it will fall within a stratum that is already occupied by a previously added sample. The stratum in which it falls is divided along the direction of the maximum distance between it and the other sample in that stratum. This method of post-stratification is admittedly simplistic and may not be well-suited for many problems. More advanced initial stratification methods that more accurately approximate the limit state are currently under development. The multi-cut TRS algorithm is run for 20 iterations. The number of Markov Chains is problem dependent and needs to be selected in such a way that at least once sample is obtained per dis-joint failure region. A heuristic approach is to consider the number of chains equal to the number of dimensions. This, however, may not always be the optimal number and, to the extent possible, should be inferred from the problem itself.

The proposed method is compared with subset simulation (SS) (Au and Beck (2001)) and importance sampling (IS) based on design point (Melchers (1999), Madsen et al. (2006)). The efficiency of the methods is compared in terms of the coefficient of variation, COV, (over 20 independent trials) and the number of performance function evaluations N_g as reported in Table 1. The exact probability of failure for the problem considered is 1×10^{-6} (Engelund and Rackwitz (1993). For the subset simulation method, 500 samples are used for determining the intermediate conditional probabilities of failure in increments of 0.1. Importance sampling estimate

is obtained using 1000 samples with the importance sampling density function as normal with mean located at the design point and identity covariance matrix.

As is witnessed from Table 1, the performance of the proposed multi-cut TRS method is comparable to importance sampling and produces approximately the same COV as subset simulation with approximately $1/3$ the number of function evaluations. The large COV in the TRS estimates is likely due to the limited number of TRS iterations along with the crudeness of the initial stratification. This can also be witnessed from Figure 1, which shows the convergence of the proposed method for 20 iterations and 10 independent trials. The estimate for probability of failure will converge to the true value as the number of iterations are increased beyond 20, provided, the Markov Chains have located all the regions that contribute significantly to the probability of failure. Given that most of the trials produce a good estimate for P_F, 20 iterations can be deemed as being suitable for the chosen example. The trials that do not converge to the true value are provided to emphasize the fact that the number of iterations, and the estimate for P_F crucially depends on the identification of all the failure regions by the Markov Chains.

Table 1: Comparison of estimates of P_F obtained using Target Random Sampling (TRS), Importance Sampling (IS), and Subset Simulation (SS).

Method	Estimate (COV)	N_g
TRS	1.2002×10^{-6} (56%)	1000 approx.
IS	8.9474×10^{-7} (40%)	1000
SS	1.0888×10^{-6} (51%)	2700
FORM	5.9942×10^{-4}	–
SORM	6.4885×10^{-4}	–

5. CONCLUSIONS

The Targeted Random Sampling method has been extended to handle moderate-high dimensional problems. The approach is based on application of multi-directional stratum division and neural network based performance function approximation. Identification of all disjoint failure domains and good initial stratification technique are crucial to the performance of the algorithm. The authors are working on development of improved initial stratification methodologies, and extending the method to time variant reliability problems.

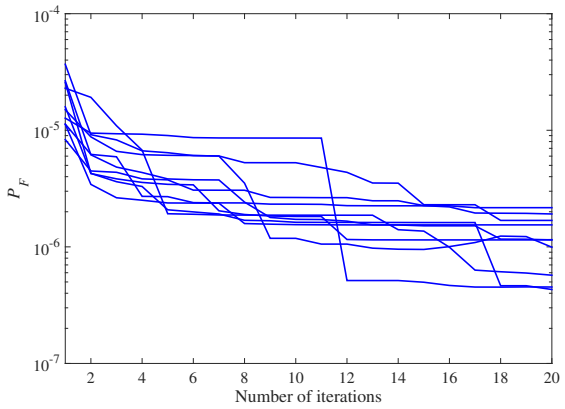

Figure 1: Convergence of the multicut TRS method for 10 independent trials.

6. REFERENCES

Au, S.-K. and Beck, J. L. (2001). "Estimation of small failure probabilities in high dimensions by subset simulation." *Probabilistic Engineering Mechanics*, 16(4), 263–277.

Bucher, C. and Bourgund, U. (1990). "A fast and efficient response surface approach for structural reliability problems." *Structural safety*, 7(1), 57–66.

Der Kiureghian, A. and Dakessian, T. (1998). "Multiple design points in first and second-order reliability." *Structural Safety*, 20(1), 37–49.

Engelund, S. and Rackwitz, R. (1993). "A benchmark study on importance sampling techniques in structural reliability." *Structural Safety*, 12(4), 255–276.

Ghanem, R. G. and Spanos, P. D. (1991). *Stochastic finite elements: a spectral approach*, Vol. 387974563. Springer.

Hasofer, A. M. and Lind, N. C. (1974). "Exact and invariant second-moment code format." *Journal of the Engineering Mechanics Division*, 100(1), 111–121.

Koutsourelakis, P., Pradlwarter, H., and Schuëller, G. (2004). "Reliability of structures in high dimensions, part i: algorithms and applications." *Probabilistic Engineering Mechanics*, 19(4), 409–417.

Madsen, H. O., Krenk, S., and Lind, N. C. (2006). *Methods of structural safety*. Courier Dover Publications.

Melchers, R. E. (1999). "Structural reliability analysis and prediction (civil engineering)." *Wiley*.

Rubinstein, R. Y. and Kroese, D. P. (2008). *Simulation and the Monte Carlo method*, Vol. 707. Wiley-interscience.

Schuëller, G. and Pradlwarter, H. (2007). "Benchmark study on reliability estimation in higher dimensions of structural systems–an overview." *Structural Safety*, 29(3), 167–182.

Shields, M., Teferra, K., Hapij, A., and Daddazio, R. (2014). "Refined stratified sampling for efficient monte carlo based uncertainty quantification." *Reliability Engineering and System Safety (under review)*.

Shields, M. D. and Sundar, V. S. (2014). "Targeted random sampling: A new approach for efficient reliability estimation for complex systems." *International Journal of Reliability and Safety (under review)*.

Shinozuka, M. and Deodatis, G. (1991). "Simulation of stochastic processes by spectral representation." *Applied Mechanics Reviews*, 44(4), 191–204.

Sudret, B. and Der Kiureghian, A. (2000). *Stochastic finite element methods and reliability: a state-of-the-art report*. Department of Civil and Environmental Engineering, University of California.