Structural Safety 62 (2016) 66–75

Contents lists available at ScienceDirect

Structural Safety

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

Sequential importance sampling for structural reliability analysis

Iason Papaioannou [a,*], Costas Papadimitriou [b], Daniel Straub [a]

[a] Engineering Risk Analysis Group, Technische Universität München, Arcisstr. 21, 80290 München, Germany
[b] Department of Mechanical Engineering, University of Thessaly, Pedion Areos, 38334 Volos, Greece

ARTICLE INFO

Article history:
Received 31 July 2015
Received in revised form 27 May 2016
Accepted 2 June 2016
Available online 17 June 2016

Keywords:
Reliability analysis
Simulation method
Importance sampling
MCMC
High dimensions

ABSTRACT

This paper proposes the application of sequential importance sampling (SIS) to the estimation of the probability of failure in structural reliability. SIS was developed originally in the statistical community for exploring posterior distributions and estimating normalizing constants in the context of Bayesian analysis. The basic idea of SIS is to gradually translate samples from the prior distribution to samples from the posterior distribution through a sequential reweighting operation. In the context of structural reliability, SIS can be applied to produce samples of an approximately optimal importance sampling density, which can then be used for estimating the sought probability. The transition of the samples is defined through the construction of a sequence of intermediate distributions. We present a particular choice of the intermediate distributions and discuss the properties of the derived algorithm. Moreover, we introduce two MCMC algorithms for application within the SIS procedure; one that is applicable to general problems with small to moderate number of random variables and one that is especially efficient for tackling high-dimensional problems.

1. Introduction

Structural reliability analysis requires the evaluation of the probability of failure, defined by the following n-fold integral:

$$P_f = \int_{g(\mathbf{x})\leqslant 0} f_{\mathbf{X}}(\mathbf{x})d\mathbf{x} \tag{1}$$

where $\mathbf{X}$ is an n-dimensional random vector and models the system variables that are expected to present an uncertain behavior, $f_{\mathbf{X}}(\mathbf{x})$ is the joint probability density function (PDF) of $\mathbf{X}$ and $g(\mathbf{x}) \leqslant 0$ defines the failure event. The function $g(\mathbf{x})$ is usually termed limit-state function and it can include one or several distinct failure modes [1].

It is common to transform the random variables $\mathbf{X}$ to a probability space $\mathbf{U}$ consisting of independent standard normal random variables. This is achieved by an isoprobabilistic transformation $\mathbf{U} = \mathbf{T}(\mathbf{X})$ [2,3]. The probability of failure can be expressed in the transformed space as

$$P_f = \int_{G(\mathbf{u})\leqslant 0} \varphi_n(\mathbf{u})d\mathbf{u} \tag{2}$$

where φ_n is the n-variate standard normal PDF and $G(\mathbf{u}) = g(\mathbf{T}^{-1}(\mathbf{u}))$ is the limit-state function in $\mathbf{U}$-space.

The integral in Eq. (2) can be evaluated by a variety of existing approaches [1,4]. These include approximation methods such as the first/second order reliability method (FORM/SORM), response surface approaches and simulation techniques based on the Monte Carlo method. Among these, simulation methods are often preferred because of their robustness in dealing with complex engineering models. The probability integral can be expressed as the expectation of the indicator function $I(G(\mathbf{u}) \leqslant 0)$, where $I(G(\mathbf{u}) \leqslant 0) = 1$ if $G(\mathbf{u}) \leqslant 0$ and $I(G(\mathbf{u}) \leqslant 0) = 0$ otherwise. Standard Monte Carlo estimates P_f by generating n_s independent samples $\{\mathbf{u}_k, k = 1, \ldots, n_s\}$ from the PDF $\varphi_n(\mathbf{u})$ and taking the sample mean of $I(G(\mathbf{u}) \leqslant 0)$, i.e.

$$\hat{P}_f = \hat{\mathrm{E}}_{\varphi_n}[I(G(\mathbf{u}) \leqslant 0)] = \frac{1}{n_s}\sum_{k=1}^{n_s} I(G(\mathbf{u}_k) \leqslant 0) \tag{3}$$

The estimate of Eq. (3) is unbiased and has coefficient of variation:

$$\delta_{\hat{P}_f} = \sqrt{\frac{1-P_f}{n_s P_f}} \tag{4}$$

$\delta_{\hat{P}_f}$ is a measure of the statistical accuracy of $\hat{P}_f$. Although $\delta_{\hat{P}_f}$ does not depend on the dimension of the random variable space n, it is inversely proportional to the target probability P_f. Hence for a probability in the order of 10^{-k}, crude Monte Carlo requires approximately 10^{k+2} samples to achieve an accuracy of $\delta_{\hat{P}_f} = 10\%$.

* Corresponding author
E-mail address: iason.papaioannou@tum.de (I. Papaioannou).

http://dx.doi.org/10.1016/j.strusafe.2016.06.002
0167-4730/

Several methods have been proposed that aim at reducing the variance of the crude Monte Carlo estimate. These include importance sampling (IS) and its adaptive variants [5–7], line sampling [8,9] and subset simulation (SuS) [10]. All of the above methods are based on producing samples that explore the failure region. In this paper, we discuss a sampling method that adaptively samples the failure region, termed sequential importance sampling (SIS). SIS was developed in the statistical community for exploring posterior distributions and estimating normalizing constants in the context of Bayesian analysis [11–13]. Although the published variants of the method diverge in their implementation, they are all based on the same principle of gradually transforming samples from a prior distribution to samples from a posterior distribution through a sequential reweighting operation. A variant of the method was introduced in the engineering community as transitional Markov chain Monte Carlo [14]. Initial ideas for application of SIS to structural reliability have been presented in [15]. Therein, estimation of the intermediate distributions was performed following an algorithm proposed in [16]. This algorithm was further developed for application to reliability-based optimization in [17].

Here, we present the principle of SIS for structural reliability and discuss ingredients for its efficient implementation. These include two Markov chain Monte Carlo (MCMC) algorithms. The first is an independent Metropolis–Hastings (M–H) algorithm combined with a Gaussian mixture proposal distribution that is suitable for application to low to moderate dimensional problems; the second is a conditional sampling M–H algorithm for application in high dimensions. The performance of the proposed algorithms is demonstrated through numerical examples.

2. Sequential importance sampling for reliability analysis

In this section, we introduce SIS for structural reliability. We first review standard IS; then we describe SIS for sampling from a sequence of distributions; a particular sequence of distributions for application to structural reliability is discussed next; subsequently we introduce two MCMC algorithms as important ingredients of SIS for application to different structural reliability problems; we finally draw a connection between SIS and SuS.

2.1. Importance sampling

Let $h(\mathbf{u})$ be a positive density referred to as the IS function. The integral in Eq. (2) can be rewritten as:

$$P_f = \int_{\mathbb{R}^n} I(G(\mathbf{u}) \leqslant 0) w(\mathbf{u}) h(\mathbf{u}) d\mathbf{u} = \mathrm{E}_h[I(G(\mathbf{u}) \leqslant 0) w(\mathbf{u})] \tag{5}$$

where $w(\mathbf{u}) = \frac{\varphi_n(\mathbf{u})}{h(\mathbf{u})}$ is the so-called importance weight function. An estimate of P_f can be obtained by generating samples $\{\mathbf{u}_k, k = 1, \ldots, n_s\}$ from $h(\mathbf{u})$ and taking the sample mean of $I(G(\mathbf{u}) \leqslant 0) w(\mathbf{u})$, i.e.

$$\hat{P}_f = \hat{\mathrm{E}}_h[I(G(\mathbf{u}) \leqslant 0) w(\mathbf{u})] = \frac{1}{n_s} \sum_{k=1}^{n_s} I(G(\mathbf{u}_k) \leqslant 0) w(\mathbf{u}_k) \tag{6}$$

The probability estimate of Eq. (6) is unbiased provided that the support of $h(\mathbf{u})$ contains the failure domain $G(\mathbf{u}) \leqslant 0$. An appropriate choice of the IS function can lead to significantly smaller variance compared to the one of the crude Monte Carlo estimate. The theoretically optimal IS is given by the following expression:

$$h_{\mathrm{opt}}(\mathbf{u}) = \frac{1}{P_f} I(G(\mathbf{u}) \leqslant 0) \varphi_n(\mathbf{u}) \tag{7}$$

The IS function of Eq. (7) leads to a variance of the IS estimate of zero. However, the optimal IS function cannot be used in practice since it requires knowledge of P_f. Common choices are unimodal [18,5] or multimodal [6] densities based on initial sampling or other type of calculations. However, it has been discussed [19,20] that in settings involving a large number of random variables, IS based on such densities may fail to describe the important region, leading to a dramatic increase of the variance of the resulting estimate.

2.2. Sequential importance sampling

Consider a sequence of distributions $\{h_j(\mathbf{u}), j = 0, \ldots, M\}$, where each distribution is known up to a normalizing constant, i.e.

$$h_j(\mathbf{u}) = \frac{\eta_j(\mathbf{u})}{P_j} \tag{8}$$

where $\eta_j(\mathbf{u})$ is known pointwise and the normalizing constant P_j is unknown. We assume that $\eta_0(\mathbf{u}) = h_0(\mathbf{u})$ and hence $P_0 = 1$. We further postulate that $h_0(\mathbf{u})$ is easy to sample from. We are interested in obtaining samples from $h_M(\mathbf{u})$ and estimating the normalizing constant P_M. The idea of SIS is to sample the distributions $\{h_j(\mathbf{u}), j = 0, \ldots, M\}$ in a step-wise manner and estimate each normalizing constant P_j by IS using as IS density the function $h_{j-1}(\mathbf{u})$. Assume that at step $j-1$ samples $\{\mathbf{u}_k, k = 1, \ldots, n_s\}$ from $h_{j-1}(\mathbf{u})$ are available. The constant P_j can be written as:

$$P_j = \int_{\mathbb{R}^n} \eta_j(\mathbf{u}) d\mathbf{u} = P_{j-1} \int_{\mathbb{R}^n} w_j(\mathbf{u}) h_{j-1}(\mathbf{u}) d\mathbf{u} = P_{j-1} \mathrm{E}_{h_{j-1}}[w_j(\mathbf{u})] \tag{9}$$

where $w_j(\mathbf{u}) = \frac{\eta_j(\mathbf{u})}{\eta_{j-1}(\mathbf{u})}$. An estimate of the ratio of normalizing constants $S_j = \frac{P_j}{P_{j-1}}$ is given by:

$$\hat{S}_j = \frac{\hat{P}_j}{\hat{P}_{j-1}} = \hat{\mathrm{E}}_{h_{j-1}}[w_j(\mathbf{u})] = \frac{1}{n_s} \sum_{k=1}^{n_s} w_j(\mathbf{u}_k) \tag{10}$$

To obtain an accurate estimate $\hat{S}_j$, we need to ensure that the two densities $h_{j-1}(\mathbf{u})$ and $h_j(\mathbf{u})$ do not vary significantly. This can be controlled by selecting $\eta_j(\mathbf{u})$ such that the variance of the importance weights is small. Given samples from $h_{j-1}(\mathbf{u})$, we can obtain samples from $h_j(\mathbf{u})$ applying the following resample-move scheme. First, we apply a resampling method that selects randomly with replacement samples from $\{\mathbf{u}_k, k = 1, \ldots, n_s\}$ with probability assigned to each kth sample proportional to $w_j(\mathbf{u}_k)$ [21]. We then move the resulting samples in regions of high probability mass of $h_j(\mathbf{u})$ by applying MCMC with invariant distribution $h_j(\mathbf{u})$. This procedure is repeated for each subsequent step and an estimate of P_M is obtained as:

$$\hat{P}_M = \prod_{j=1}^{M} \hat{S}_j \tag{11}$$

2.3. Choice of intermediate distributions

In the context of structural reliability, SIS can be applied to obtain samples from an approximation of the optimal IS density of Eq. (7). The indicator function $I(G(\mathbf{u}) \leqslant 0)$ can be expressed by the following limit (e.g. see [22])

$$I(G(\mathbf{u}) \leqslant 0) = \lim_{\sigma \to 0} \Phi\left(-\frac{G(\mathbf{u})}{\sigma}\right) \tag{12}$$

where Φ is the standard normal cumulative distribution function (CDF). It is noted that this is only one of several equivalent definitions of the indicator function through limits of smooth functions. Choosing $\sigma = \sigma_M$, with σ_M small enough, we can approximate $I(G(\mathbf{u}) \leqslant 0)$ by the following expression

$$I(G(\mathbf{u}) \leqslant 0) \approx \Phi\left(-\frac{G(\mathbf{u})}{\sigma_M}\right) \tag{13}$$