Journal of Computational Physics 341 (2017) 61–75

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Combining multiple surrogate models to accelerate failure probability estimation with expensive high-fidelity models

Benjamin Peherstorfer *, Boris Kramer, Karen Willcox

Department of Aeronautics & Astronautics, MIT, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:
Received 7 July 2016
Received in revised form 11 January 2017
Accepted 4 April 2017
Available online 6 April 2017

Keywords:
Multifidelity
Model reduction
Uncertainty quantification
Failure probability estimation
Surrogate modeling
Rare event simulation

ABSTRACT

In failure probability estimation, importance sampling constructs a biasing distribution that targets the failure event such that a small number of model evaluations is sufficient to achieve a Monte Carlo estimate of the failure probability with an acceptable accuracy; however, the construction of the biasing distribution often requires a large number of model evaluations, which can become computationally expensive. We present a mixed multifidelity importance sampling (MMFIS) approach that leverages computationally cheap but erroneous surrogate models for the construction of the biasing distribution and that uses the original high-fidelity model to guarantee unbiased estimates of the failure probability. The key property of our MMFIS estimator is that it can leverage multiple surrogate models for the construction of the biasing distribution, instead of a single surrogate model alone. We show that our MMFIS estimator has a mean-squared error that is up to a constant lower than the mean-squared errors of the corresponding estimators that uses any of the given surrogate models alone—even in settings where no information about the approximation qualities of the surrogate models is available. In particular, our MMFIS approach avoids the problem of selecting the surrogate model that leads to the estimator with the lowest mean-squared error, which is challenging if the approximation quality of the surrogate models is unknown. We demonstrate our MMFIS approach on numerical examples, where we achieve orders of magnitude speedups compared to using the high-fidelity model only.

1. Introduction

System inputs are often modeled as random variables to account for uncertainties in the inputs due to measurement errors, noise, or small perturbations in the manufacturing processes. A high-fidelity model of the system of interest maps inputs onto model outputs that approximate the system outputs with the accuracy required by the problem at hand. With an input random variable, the model output becomes a random variable as well. An important task in the context of reliability analysis is the estimation of probabilities of failure events. Basic Monte Carlo estimation draws realizations from the input random variable, evaluates the high-fidelity model at the realizations, and derives a failure probability estimate from the corresponding outputs. Because failure probabilities are typically small, basic Monte Carlo estimation typically requires a large number of high-fidelity model evaluations to derive estimates of an acceptable accuracy. Therefore, basic Monte Carlo estimation can become computationally intractable if the high-fidelity model is expensive to solve.

* Corresponding author.
E-mail address: pehersto@mit.edu (B. Peherstorfer).

http://dx.doi.org/10.1016/j.jcp.2017.04.012
0021-9991/

One class of methods that accelerate Monte Carlo estimation derives problem-dependent sampling strategies such that fewer samples are necessary to obtain an acceptable estimate of the failure probability than in basic Monte Carlo estimation. One such method is importance sampling. The first step of importance sampling is constructing a problem-dependent biasing distribution in which the failure event has a larger probability than in the original nominal distribution. In the second step, the biasing distribution is used to derive an estimate of the failure probability, where the change in distribution is compensated via re-weighting the samples of the biasing distribution. The key challenge of importance sampling is constructing a suitable biasing distribution. A general overview of importance sampling is given in [1–3]. Adaptive importance sampling techniques iterate between the biasing distribution construction and the estimation step to iteratively construct and adapt the biasing distribution as the computation proceeds [4–8]. The cross-entropy method [9,10] can be used to iteratively build the biasing distribution. Another popular approach to accelerate Monte Carlo estimation of failure probabilities is the subset method [11–13]. The subset method starts with an event that includes the failure event and that has a large probability compared to the failure probability. The subset method then successively refines the event until the actual failure event of interest is recovered. During the refinement, a problem-dependent sampling strategy is constructed with which the failure event can be estimated with few samples. A multilevel extension to the subset method is introduced in [14].

Another class of methods that accelerate Monte Carlo estimation uses surrogate models so that expensive high-fidelity model evaluations can be replaced with cheap surrogate evaluations. Examples of surrogate models include projection-based reduced models [15,16], data-fit interpolation and regression models [17], machine-learning-based models such as support vector machines [18], and other simplified models [19]. However, simply replacing the high-fidelity model with a surrogate model in a basic Monte Carlo estimator can lead to significantly biased estimates of the failure probability [20,21]. Multifidelity methods combine surrogate model and high-fidelity model outputs to control, or even avoid, the bias that is introduced by a surrogate model. In [22,23], the surrogate model is successively adapted during the Monte Carlo estimation so that the surrogate model becomes an accurate approximation of the high-fidelity model at the failure event. In [21,24], an approach is presented that switches between a surrogate model and the high-fidelity model depending on the error of the surrogate model. The obtained estimators are unbiased if tolerance parameters are chosen well. The work [25] uses *a posteriori* error estimators of a surrogate model to decide whether the surrogate model or the high-fidelity model should be used.

The two-fidelity importance sampling approach introduced in [26] uses the surrogate model for the construction of the biasing distribution and the high-fidelity model to derive the estimate of the failure probability. The key property of the two-fidelity approach [26] is that the two-fidelity estimator is an unbiased estimator of the failure probability, independent of the availability of error guarantees on the surrogate model. A surrogate model that is a poor approximation of the high-fidelity model can lead to a two-fidelity estimator that requires more evaluations of the high-fidelity model than a basic Monte Carlo estimator with the same mean-squared error (MSE); however, the unbiasedness of the two-fidelity estimator is guaranteed. If the surrogate model provides a reasonable approximation of the high-fidelity model, then the two-fidelity approach can lead to significant speedups compared to basic Monte Carlo estimation, see the numerical examples in [26].

The two-fidelity approach presented in [26], as well as the methods [22,23,21,24,25] discussed above, can leverage only a single surrogate model. If multiple surrogate models are given, one of the surrogate models has to be selected. Error bounds for the surrogate model outputs are available only in limited problem settings, which means that it is challenging in many situations to make an informed decision which of the surrogate models to use in a two-fidelity approach. Even if error bounds for the surrogate model outputs are available, it often remains challenging to propagate the bounds for the outputs onto bounds for the MSE of two-fidelity estimators that use the surrogate model. We present here a mixed multifidelity importance sampling (MMFIS) method that builds on the mixed importance sampling methodology of [27] and the two-fidelity approach of [26] to derive a biasing distribution from multiple surrogate models, instead of using a single surrogate model alone. We show that the MSE of our MMFIS estimator is up to a constant as low as the MSE of the best two-fidelity estimator (i.e., the one that uses the surrogate model that leads to the lowest MSE). Our MMFIS estimator avoids the problem of selecting one of the given surrogate models and is still guaranteed to achieve an MSE that is up to a constant as low as the MSE of any of the two-fidelity estimators derived from the given surrogate models. The constant is equal to the number of surrogate models and therefore is typically small. Furthermore, our MMFIS estimator is an unbiased estimator of the failure probability.

The outline of the presentation is as follows. Section 2 gives a problem definition and briefly discusses basic Monte Carlo estimation, importance sampling, and the two-fidelity importance sampling method of [26]. Section 3 introduces our MMFIS method that combines outputs of multiple surrogate models for the construction of a biasing distribution. Numerical results in Section 4 demonstrate our MMFIS method on a structural reliability example and a shock propagation problem. The conclusions in Section 5 close the presentation.

2. Failure probability estimation with importance sampling

This section gives the problem formulation and briefly discusses Monte Carlo approaches to estimate failure probabilities. Section 2.1 defines the problem, and Section 2.2 introduces basic Monte Carlo estimation and importance sampling. Section 2.3 introduces two-fidelity importance sampling, which fuses outputs of a surrogate model with outputs of the high-fidelity model to accelerate the estimation of failure probabilities compared to basic Monte Carlo estimation and importance sampling.