



Hamiltonian MCMC methods for estimating rare events probabilities in high-dimensional problems

Konstantinos G. Papakonstantinou ^{*}, Hamed Nikbakht, Elsayed Eshra

Department of Civil and Environmental Engineering, The Pennsylvania State University, University Park, PA 16802, United States of America



ARTICLE INFO

Keywords:

Hamiltonian MCMC
Quasi-Newton
Rare event probability
Reliability estimation
High-dimensional spaces
Inverse importance sampling

ABSTRACT

Accurate and efficient estimation of rare events probabilities is of significant importance, since often the occurrences of such events have widespread impacts. The focus in this work is on precisely quantifying these probabilities, often encountered in reliability analysis of complex engineering systems, based on an introduced framework termed *Approximate Sampling Target with Post-processing Adjustment* (ASTPA), which herein is integrated with and supported by gradient-based Hamiltonian Markov Chain Monte Carlo (HMC) methods. The developed techniques in this paper are applicable from low- to high-dimensional stochastic spaces, and the basic idea is to construct a relevant target distribution by weighting the original random variable space through a one-dimensional output likelihood model, using the limit-state function. To sample from this target distribution, we exploit HMC methods, a family of MCMC methods that adopts physical system dynamics, rather than solely using a proposal probability distribution, to generate distant sequential samples, and we develop a new *Quasi-Newton mass preconditioned* HMC scheme (QNP-HMC), which is particularly efficient and suitable for high-dimensional spaces. To eventually compute the rare event probability, an original post-sampling step is devised using an *inverse importance sampling* procedure based on the already obtained samples. The statistical properties of the estimator are analyzed as well, and the performance of the proposed methodology is examined in detail and compared against Subset Simulation in a series of challenging low- and high-dimensional problems.

1. Introduction

In this work, we develop a framework for estimation of rare events probabilities, a commonly encountered important problem in several engineering and scientific applications, often observed in the form of probability of failure (P_F) estimation or, alternatively, reliability estimation. In many practical applications, failure probabilities are fortunately very low, from 10^{-4} to even 10^{-9} and lower, and calculating such small probabilities presents many numerical and mathematical challenges, particularly in cases with high dimensional random spaces and/or expensive computational models, that practically limit the afforded number of model calls. The number of model calls is thus of great importance in these problems and one of the critical parameters that limits or prohibits use of several available techniques in the literature.

The reliability estimation problem has a long history in the engineering community [1–5]. One of the significant early successes was the discovery of the so called First Order Reliability Method (FORM) [6,7], long investigated by Der Kiureghian, Ditlevsen and co-workers [8,9], and many others, e.g., Shinozuka [10], providing also several enhancements, including second order effects (SORM)

by Breitung [11]. In FORM/SORM methods, the search for the most probable failure point (MPP) is usually performed by gradient-based optimization methods [12,13]. Although these asymptotic approximation methods are usually computationally inexpensive, they have several limitations and may involve considerable errors, particularly in high-dimensional problems or in problems with highly nonlinear limit-state functions [6,14]. As such, various sampling-based methods have also been suggested by the reliability community, e.g., Schuëller and Pradlwarter [15], to tackle the problem in its utmost generality, with crude Monte Carlo approaches being prohibitive for this type of problems due to their excessive computational demands. Only some of many notable contributions can be seen in [16–20], describing and studying the state-of-the-art Subset Simulation and its enhancements, originally presented in [21], and in [22–24] utilizing importance sampling schemes, often also combined with the cross-entropy method [25–29]. Alternative approaches include directional and line sampling [30–33], the PHI2 method for time-variant reliability [34], and asymptotic sampling strategies [35,36], among others. The problem of estimating rare event probabilities has also attracted a great deal of

^{*} Corresponding author.

E-mail addresses: kpapakon@psu.edu (K.G. Papakonstantinou), hun35@psu.edu (H. Nikbakht), eme5375@psu.edu (E. Eshra).

attention in other relevant communities and in mathematical literature, with several suggested methods sharing many similarities with FORM/SORM approaches, e.g., [37], and Subset Simulation, such as in approaches involving Sequential Monte Carlo samplers for rare events [38–40], interacting particle system methodologies [41,42], multilevel splitting methods [43,44] and forward flux sampling [45], to name but a few.

In this paper, we are presenting a new solution approach to the problem by combining gradient-based approaches, already familiar to the engineering community and often available in computational tools, with Markov Chain Monte Carlo (MCMC) sampling methods in the form of Hamiltonian MCMC. MCMC methods [46] are plausibly the most broadly accepted ones to generate samples from target distributions, in cases where direct sampling is not possible. Despite notable successes, many MCMC methods scale poorly with the number of dimensions and can become inefficient. For complicated multivariate models, classic methods such as random-walk Metropolis–Hastings [47] and Gibbs sampling [48] may require an unacceptably long time and number of samples to adequately explore the target distribution. Originally developed by Duane et al. [49] and to a large extent understood and popularized through the works of Neal [50,51], Hamiltonian Markov Chain Monte Carlo (HMC), usually called Hamiltonian Monte Carlo (HMC) in the literature, produces Markov chain samples based on Hamiltonian dynamics principles, is characterized by much better scalability [51,52] and faster mixing rates, is capable of generating samples with much weaker auto-correlation, even in complex high-dimensional random spaces [53], and has enjoyed broad-spectrum successes in most general settings [54]. Balanced against these features and achievements is of course the need for multiple gradient evaluations in each HMC sampling iteration, making the method more computationally intensive per iteration than other algorithms, such as random-walk Metropolis–Hastings and Metropolis-adjusted Langevin [55,56], for example. Girolami and Calderhead introduced a Riemannian Manifold Hamiltonian Monte Carlo (RMHMC) approach in [57,58] that has demonstrated significant successes in many challenging problems but requires computing higher-order derivatives of the target distribution. Overall, the two impediments to using Hamiltonian MCMC methods are the required gradients, since analytical formulas are not always available and numerical techniques are computationally costly, particularly in high dimensions, and the heedful tuning of the involved parameters [51]. The first issue can in certain cases be solved by automatic differentiation (e.g. [54,59]) and stochastic gradient approaches [60], while for the second a fully automated state-of-the-art HMC sampling algorithm has been developed by Hoffman and Gelman, known as the No-U-Turn Sampler (NUTS [61]). NUTS introduces, among others, an expensive tree building procedure, in order to trace when the Hamiltonian trajectory turns back on itself. Many of these approaches are not however relevant and/or applicable to the analyzed problem in this work, since, in general, many rare event and reliability problems involve complex, computationally expensive models, complicating and/or precluding use of automatic differentiation and data-based stochastic gradient techniques, as well as methodologies that require a considerably high number of model calls per problem, such as NUTS.

A new computationally efficient sampling framework for estimation of rare events probabilities is thus presented in this work, having exceptional performance in quantifying low failure probabilities for any type of reliability problems described in both low and high dimensional stochastic spaces. The introduced methodology is termed *Approximate Sampling Target with Post-processing Adjustment* (ASTPA) and comprises a sampling and a post-processing phase. The sampling target in ASTPA is constructed by appropriately combining the multi-dimensional random variable space with a cumulative distribution function that utilizes the limit-state function. Having acquired the samples, an adjustment step is then applied, in order to account for the fact that the samples are drawn from an approximate target distribution, and to thus correctly quantify the rare event probability. An original *inverse importance sampling* procedure is devised for this adjustment step, taking its name from

the fact that the samples are already available. Although the ASTPA framework is general and can be combined with any appropriate Monte Carlo sampling method, it becomes substantially efficient when directly supported by gradient-based Hamiltonian Markov Chain Monte Carlo (HMC) samplers. To address the scalability issues a typical HMC sampler may manifest in high-dimensional spaces, a new *Quasi-Newton mass preconditioned HMC* approach is also developed. This new sampling scheme follows an approximate Newton direction and estimates the pertinent Hessian in its burn-in stage, only based on gradient information and the BFGS approximation, and eventually utilizes the computed Hessian as a preconditioned mass matrix in the main non-adaptive sampling stage. An approximate analytical expression for the uncertainty of the computed estimation is also derived, showcasing significant accuracy with numerical results, and all involved user-defined parameters of ASTPA are thoroughly analyzed and general default values are suggested. Finally, to fully examine the capabilities of the proposed methodology, its performance is demonstrated and compared against Subset Simulation in a series of challenging low- and high-dimensional problems.

2. Failure probability estimation

The failure probability P_F for a system, that is the probability of a defined unacceptable system performance, can be expressed as a d -fold integral, as:

$$P_F = \mathbb{E}[I_F(\theta)] = \int_{g(\theta) \leq 0} I_F(\theta) \pi_\Theta(\theta) d\theta \quad (1)$$

where θ is the random vector $[\theta_1, \dots, \theta_d]^T$, $F \subset \mathbb{R}^d$ is the failure event, $g(\theta)$ is the limit-state function that can include one or several distinct failure modes and defines the system failure by $g(\theta) \leq 0$, $I(\cdot)$ denotes the indicator function with $I_F(\theta) = 1$ if $\theta \in g(\theta) \leq 0$ and $I_F(\theta) = 0$ otherwise, \mathbb{E} is the expectation operator, and π_Θ is the joint probability density function (PDF) for Θ . As is common practice for problems of this type, in this work the joint PDF of Θ is the standard normal one, due to its rotational symmetry and exponential probability decay. In most cases this is not restrictive, since it is uncomplicated to transform the original random variables \mathbf{X} to Θ , e.g. [62]. When this is not the case, however, and the probabilistic characterization of \mathbf{X} can be defined in terms of marginal distributions and correlations, the Nataf distribution (equivalent to Gaussian copula) is commonly used to model the joint PDF, and the mapping to the standard normal space can be then accomplished [8,63].

The focus in this work is to analyze the described integration in Eq. (1) under very general settings, including the following challenging sampling context: (i) Computation of Eq. (1) can only be done in approximate ways; (ii) the relationship between θ and I_F is not explicitly known and for any θ we can merely check whether it belongs to the failure set or not, i.e. calculate the value $I_F(\theta)$; (iii) the computational effort for evaluating $I_F(\theta)$ for each value of θ is assumed to be quite significant, so that it is essential to minimize the number of such function evaluations (model calls); (iv) the probability of failure P_F is assumed to be very small, e.g. in order of $P_F \sim 10^{-4} - 10^{-9}$; (v) the parameter space \mathbb{R}^d is assumed to be high-dimensional, in the order of 10^2 and more, for example. Under these general settings, several sampling methods, including direct Monte Carlo approaches and NUTS [61], become highly inefficient and fail to address the problem effectively. Subset Simulation (SuS) [21] has however proven successful and robust in dealing with problems of this type and is shown to outperform other relevant methods in numerous papers, e.g. [24,31]. SuS relies on a modified component-wise Metropolis MCMC method that can successfully work in high dimensions and does not require a burn-in sampling stage. A notable adaptive conditional sampling (aCS) methodology within the SuS framework is also introduced by Papaiouannou et al. in [16], providing important advantages and enhanced performance in several cases. Relevant SuS variants are thus utilized in this work, for validation and comparison purposes with our presented methodology that completely deviates from SuS and is efficiently supported by a direct Hamiltonian MCMC sampling approach [64].

3. Hamiltonian Markov Chain Monte Carlo

3.1. Standard HMC with leapfrog integrator

Based on the aforementioned discussion in the previous sections, Hamiltonian dynamics can be used to produce distant Markov chain samples, thereby avoiding the slow exploration of the state space that results from the diffusive behavior of simple random-walk proposals. This Hamiltonian approach was firstly introduced to molecular simulations by Alder and Wainwright in [65], in which the motion of the molecules was deterministic. Duane et al. in [49] united the MCMC and molecular dynamics approaches. Given d -dimensional variables of interest θ with (unnormalized) density $\pi_\theta(\cdot)$, the Hamiltonian Monte Carlo method introduces d -dimensional auxiliary momentum variables \mathbf{z} and samples from the joint distribution characterized by:

$$\pi(\theta, \mathbf{z}) \propto \pi_\theta(\theta) \pi_{Z|\theta}(\mathbf{z}|\theta) \quad (2)$$

where $\pi_{Z|\theta}(\cdot|\theta)$ is proposed to be a symmetric distribution. With $\pi_\theta(\theta)$ and $\pi_{Z|\theta}(\mathbf{z}|\theta)$ being uniquely described up to normalizing constants, the functions $U(\theta) = -\log \pi_\theta(\theta)$ and $K(\theta, \mathbf{z}) = -\log \pi_{Z|\theta}(\mathbf{z}|\theta)$ are introduced as the potential energy and kinetic energy, owing to the concept of the canonical distribution [51] and the physical laws which motivate the Hamiltonian Markov Chain Monte Carlo algorithm. The total energy $H(\theta, \mathbf{z}) = U(\theta) + K(\theta, \mathbf{z})$ is often termed the *Hamiltonian* H . The kinetic energy function is unconstrained and can be formed in various ways according to the implementation. In most typical cases, the momentum is sampled by a zero-mean normal distribution [51, 53], and accordingly the kinetic energy can be written as: $K(\theta, \mathbf{z}) = -\log \pi_{Z|\theta}(\mathbf{z}|\theta) = -\log \pi_Z(\mathbf{z}) = \frac{1}{2} \mathbf{z}^T \mathbf{M}^{-1} \mathbf{z}$, where \mathbf{M} is a symmetric, positive-definite inverse covariance matrix, termed mass matrix.

HMC generates a Metropolis proposal on the joint state-space (θ, \mathbf{z}) by sampling the momentum and simulating trajectories of Hamiltonian dynamics in which the time evolution of the state (θ, \mathbf{z}) is governed by Hamilton's equations, expressed typically by:

$$\frac{d\theta}{dt} = \frac{\partial H}{\partial \mathbf{z}} = \frac{\partial K}{\partial \mathbf{z}} = \mathbf{M}^{-1} \mathbf{z}, \quad \frac{d\mathbf{z}}{dt} = -\frac{\partial H}{\partial \theta} = -\frac{\partial U}{\partial \theta} = \nabla_\theta \mathcal{L}(\theta) \quad (3)$$

where $\mathcal{L}(\theta)$ denotes here the log-density of the target distribution. Hamiltonian dynamics prove to be an effective proposal generation mechanism because the distribution $\pi(\theta, \mathbf{z})$ is invariant under the dynamics of Eq. (3). These dynamics enable a proposal, triggered by an approximate solution of Eq. (3), to be distant from the current state, yet with high probability acceptance. The solution to Eq. (3) is analytically intractable in general and thus the Hamiltonian equations need to be numerically solved by discretizing time using some small step size, ϵ . A symplectic integrator that can be used for the numerical solution is the *leapfrog* one and works as follows:

$$\begin{aligned} \mathbf{z}_{t+\epsilon/2} &= \mathbf{z}_t - \left(\frac{\epsilon}{2}\right) \frac{\partial U}{\partial \theta}(\theta_t), \quad \theta_{t+\epsilon} = \theta_t + \epsilon \frac{\partial K}{\partial \mathbf{z}}(\mathbf{z}_{t+\epsilon/2}), \\ \mathbf{z}_{t+\epsilon} &= \mathbf{z}_{t+\epsilon/2} - \left(\frac{\epsilon}{2}\right) \frac{\partial U}{\partial \theta}(\theta_{t+\epsilon}) \end{aligned} \quad (4)$$

The main advantage of using the leapfrog integrator is its simplicity, that is volume-preserving, and that it is reversible, due to its symmetry, by simply negating \mathbf{z} , in order to generate a valid Metropolis proposal. See Neal [51] and Betancourt [53] for more details on energy-conservation, reversibility and volume-preserving integrators and their connections to HMC. It is noted that in the above leapfrog integration algorithm, the computationally expensive part is the one model call per step to acquire the $\frac{\partial U}{\partial \theta}$ term. With τ the trajectory or else path length, taking $L = \tau/\epsilon$ leapfrog steps approximates the evolution $(\theta(0), \mathbf{z}(0)) \rightarrow (\theta(\tau), \mathbf{z}(\tau))$, providing the exact solution in the limit $\epsilon \rightarrow 0$.

Typically, a simple Gaussian momentum is used for the Hamiltonian, $\pi_{Z|\theta}(\mathbf{z}|\theta) = \pi_Z(\mathbf{z}) = \mathbf{N}(\mathbf{0}, \mathbf{M})$ (or $\mathbf{z} \sim \mathbf{N}(\mathbf{0}, \mathbf{M})$) and the mass matrix \mathbf{M} is often set to the identity matrix, \mathbf{I} . A generic procedure for drawing N_{Iter} samples via HMC is described in Algorithm 1, where again $\mathcal{L}(\theta)$ is the log-density of the target distribution of interest, θ^0 are initial

Algorithm 1 Hamiltonian Markov Chain Monte Carlo

```

1: procedure HMC( $\theta^0, \epsilon, L, \mathcal{L}(\theta), N_{Iter}$ )
2:   for  $m = 1$  to  $N_{Iter}$  do
3:      $\mathbf{z}^0 \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$   $\triangleright$  momentum sampling from standard normal
4:     distribution
5:      $\theta^m \leftarrow \theta^{m-1}, \tilde{\theta} \leftarrow \theta^{m-1}, \tilde{\mathbf{z}} \leftarrow \mathbf{z}^0$ 
6:     for  $i = 1$  to  $L$  do
7:        $\tilde{\theta}, \tilde{\mathbf{z}} \leftarrow \text{Leapfrog}(\tilde{\theta}, \tilde{\mathbf{z}}, \epsilon)$   $\triangleright$  leapfrog integration
8:     end for
9:     with probability:
10:     $\alpha = \min\left\{1, \frac{\exp(\mathcal{L}(\tilde{\theta}) - \frac{1}{2} \tilde{\mathbf{z}}^T \tilde{\mathbf{z}})}{\exp(\mathcal{L}(\theta^{m-1}) - \frac{1}{2} \mathbf{z}^{0T} \mathbf{z}^0)}\right\}$   $\triangleright$  Metropolis step
11:     $\theta^m \leftarrow \tilde{\theta}, \mathbf{z}^m \leftarrow -\tilde{\mathbf{z}}$ 
12:   end for
13: end procedure

```

values, and L is the number of leapfrog steps, as explained before. For each HMC step, the momentum is first resampled and then the L leapfrog updates are performed, as seen in Eq. (4), before a typical accept/reject MCMC Metropolis step takes place.

3.2. HMC parameters

The HMC performance and efficiency is well known to rely on selecting suitable values for the ϵ and L parameters. For a fixed trajectory length τ , the stepsize ϵ balances the trade-off between accuracy and computational cost. In this work, we select the stepsize ϵ in such a way so that the corresponding average acceptance rate is approximately 65%, as values between 60% and 80% are typically assumed optimal [51, 52, 61]. The *dual averaging* algorithm of Hoffman and Gelman [61] is adopted to perform this task, used here only in the burn-in phase, to tune ϵ . To determine the number of leapfrog steps, L , we estimate an appropriate to use trajectory length τ based on a few simulation runs, so as to have a sufficiently high so called normalized *Expected Square Jumping Distance* (ESJD), $\tau^{-1/2} \mathbb{E} \|\theta^{(m+1)}(\tau) - \theta^{(m)}(\tau)\|^2$, as introduced in [66], and then we randomly perturb each trajectory length $\tau^{(m)}$ in the range $[0.9\tau, 1.1\tau]$ to further avoid periodicity (m denotes the m th iteration of HMC). In all our numerical experiments herein, we determine L and τ in this manner, as we have found it to work well in practice. The role of these parameters and techniques for determining them have been quite extensively studied in the literature and for more details we refer to [51, 52, 61].

4. Quasi-Newton mass preconditioned HMC (QNp-HMC)

In complex high-dimensional problems, the performance of the typical HMC sampler, presented as Algorithm 1, may deteriorate and a prohibitive number of model calls could be required. A variety of methods have been proposed in the literature to address this issue. Among others, a Riemannian Manifold Hamiltonian Monte Carlo (RMHMC) has been suggested in [57] that takes advantage of the manifold structure of the variable space, at the cost of calculating second- and third-order derivatives of distributions and using a generalized leapfrog scheme, requiring additional model calls per leapfrog step. Although possible in some cases regarding second-order derivatives, e.g. [67], in the majority of cases higher-order derivatives are not provided by computational models, such as finite element models. In addition, the computational cost still increases importantly and extra model calls per leapfrog step are usually restrictive for computationally expensive models.

In this work, we instead address the high-dimensionality performance issue in a different, Newton-type context, without needing additional model calls per leapfrog step and by only using the Hessian

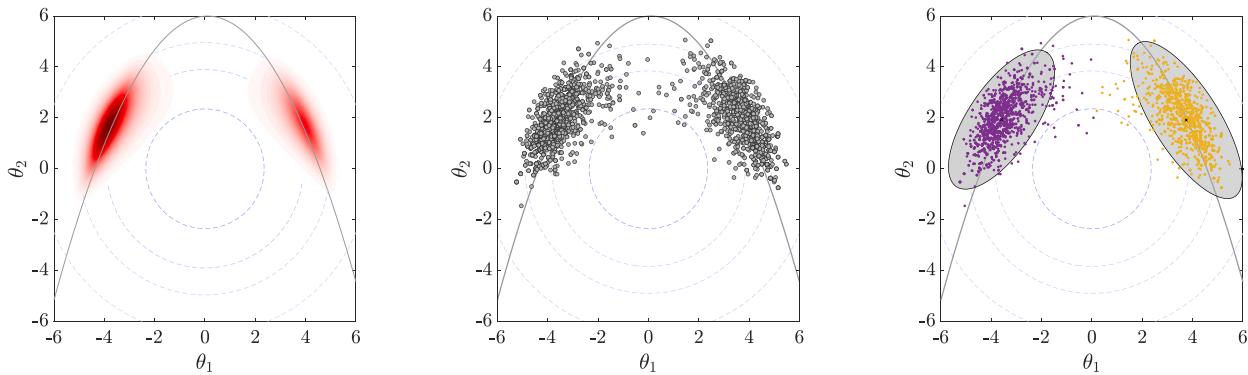


Fig. 1. An analytical target distribution, simulated target distribution samples based on our QNp-HMCMC method, and a fitted Gaussian Mixture Model describing the simulated samples, from left to right plot, respectively.

information of the target distribution, either the exact one, when relevant information is provided freely by the computational model, or, even more general, an approximate one that does not increase the computational cost. An approximate Hessian can be given in a systematic manner based on already available gradient information, similar to Quasi-Newton methods used in nonlinear programming [68]. The well-known BFGS approximation [68] is thus utilized for our Quasi-Newton type Hamiltonian MCMC approach and all numerical examples in this work are analyzed accordingly, based solely on this most general approximate Hessian case. Let $\theta \in \mathbb{R}^d$, consistent with the previous sections. Given the k th estimate \mathbf{W}_k , where \mathbf{W}_k is an approximation to the inverse Hessian at θ_k , the BFGS update \mathbf{W}_{k+1} can be expressed as:

$$\mathbf{W}_{k+1} = (\mathbf{I} - \frac{\mathbf{s}_k \mathbf{y}_k^T}{\mathbf{y}_k^T \mathbf{s}_k}) \mathbf{W}_k (\mathbf{I} - \frac{\mathbf{y}_k \mathbf{s}_k^T}{\mathbf{y}_k^T \mathbf{s}_k}) + \frac{\mathbf{s}_k \mathbf{s}_k^T}{\mathbf{y}_k^T \mathbf{s}_k} \quad (5)$$

where \mathbf{I} is the identity matrix, $\mathbf{s}_k = \theta_{k+1} - \theta_k$, and $\mathbf{y}_k = -\nabla \mathcal{L}(\theta_{k+1}) + \nabla \mathcal{L}(\theta_k)$ where $\mathcal{L} : \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the log density of the target distribution, as before. There is a long history of efficient BFGS updates for very large systems and several numerical techniques can be used, including sparse and limited-memory approaches.

Two relevant studies in the literature on Quasi-Newton extensions and connections to MCMC algorithms can be found in [69,70]. Our developed method, however, has fundamental differences that are summarized in that we are focusing on Hamiltonian methods, we are using two integrated coupled phases, an adaptive and a non-adaptive, and finally we consistently incorporate the Quasi-Newton outcomes in both stages of momentum sampling and simulation of Hamiltonian dynamics. In more detail, in the adaptive burn-in phase of the algorithm we are still sampling the momentum from an identity mass matrix, $\mathbf{M} = \mathbf{I}$, but the ODEs of Eq. (3) now become:

$$\dot{\theta} = \mathbf{W} \mathbf{M}^{-1} \mathbf{z}, \quad \dot{\mathbf{z}} = \mathbf{W} \nabla_{\theta} \mathcal{L}(\theta) \quad (6)$$

which is equivalent to the implicit linear transformation $\theta' = \mathbf{W} \theta$, and \mathbf{W} is given by Eq. (5). Accordingly, the leapfrog integrator is then reformulated as:

$$\begin{aligned} \mathbf{z}_{t+\epsilon/2} &= \mathbf{z}_t + \left(\frac{\epsilon}{2}\right) \mathbf{W} \nabla_{\theta} \mathcal{L}(\theta_t), \quad \theta_{t+\epsilon} = \theta_t + \epsilon \mathbf{W} \mathbf{z}_{t+\epsilon/2}, \\ \mathbf{z}_{t+\epsilon} &= \mathbf{z}_{t+\epsilon/2} + \left(\frac{\epsilon}{2}\right) \mathbf{W} \nabla_{\theta} \mathcal{L}(\theta_{t+\epsilon}) \end{aligned} \quad (7)$$

Hence, this new dynamic scheme efficiently and compatibly adjusts both the \mathbf{z} and θ evolutions based on the local structure of the target distribution, and also features a Quasi-Newton direction for the momentum variables, allowing large jumps across the state space. The final estimation of the approximated inverse Hessian matrix, \mathbf{W} , in the adaptive burn-in phase is then used in the subsequent non-adaptive phase of the algorithm as a preconditioned mass (inverse covariance) matrix, $\mathbf{M} = \mathbf{W}^{-1}$, used to sample the Gaussian momentum $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{M})$. As such, typical Hamiltonian dynamics are now used, albeit with this properly constructed mass matrix that takes into account the

scale and correlations of the position variables, leading to significant efficiency gains, particularly in high-dimensional problems. The BFGS procedure in Eq. (5) normally provides a symmetric, positive-definite \mathbf{W} matrix in an optimization context. However, in our case we are using BFGS under different settings that may not satisfy the curvature condition $\mathbf{s}_k^T \mathbf{y}_k > 0$, resulting in occasional deviations from positive-definiteness. Several standard techniques can be then implemented to ensure positive-definiteness, such as a damped BFGS updating [68] or the simple addition $\mathbf{W}_{new} = \mathbf{W}_{old} + \delta \mathbf{I}$, where $\delta \geq 0$ is some appropriate number. A straightforward method to determine δ is to choose it larger than the absolute value of the minimum eigenvalue of \mathbf{W}_{old} . Another technique involves utilizing a semidefinite programming approach to identify an optimized diagonal matrix to add to \mathbf{W}_{old} . Alternatively, \mathbf{W} can be updated only when the curvature condition is satisfied, which directly guarantees positive definiteness. To further ensure the stability of the sampler, a positive threshold can be introduced to the curvature condition instead of zero, e.g., $\mathbf{s}_k^T \mathbf{y}_k > 10$. This latter approach has been used and has worked well in this work. Since the final estimation of \mathbf{W} in the adaptive burn-in phase is then extensively utilized in the subsequent non-adaptive phase, we suggest use of a directly provided positive-definite matrix \mathbf{W} at this step. This can simply be accomplished by adding one more burn-in iteration step at the end of the burn-in phase, until an appropriate sample, directly supported by a positive definite \mathbf{W} matrix is identified.

Our derived *Quasi-Newton mass preconditioned Hamiltonian Markov Chain Monte Carlo* (QNp-HMCMC) method is concisely summarized and presented in Algorithm 2. Overall, QNp-HMCMC is a practical, efficient approach that only requires already available gradient information and provides important insight about the geometry of the target distribution, eventually improving computational performance and enabling faster mixing.

5. Approximate Sampling Target with Post-processing Adjustment (ASTPA)

In order for an appropriate number of samples to discover and enter the relevant regions, contributing to the rare event probability estimation, a suitable approximate target distribution is constructed in this work, as analyzed in Section 5.1, then sampled by Hamiltonian MCMC methods that can effectively reach regions of interest. For their initial stage, our HMCMC samplers have an adaptive annealed phase. This adaptive phase will be thoroughly explained later in Section 5.3. To estimate the pertinent probability, Eq. (1) needs to be then adjusted accordingly, since the samples are sampled from our constructed target distribution and not $\pi_{\theta}(\theta)$. An original post-sampling step is devised at this stage, using an inverse importance sampling procedure, described in Section 5.2, i.e. first sample, then choose the importance sampling density automatically, based on the samples. Fig. 1 concisely portrays the overall approach by using a bimodal target distribution with $P_F \sim$

3.95×10^{-5} . The gray curves represent the parabolic limit-state function $g(\theta)$ of this problem, with the failure domain being outside $g(\theta)$. The left figure displays the constructed target distribution, which in this simple 2D case can be visualized. The middle figure shows drawn samples from the target distribution by our suggested QNp-HMCMC algorithm, described in Section 4, and the right figure demonstrates the inverse importance sampling step. All these different steps will be discussed in detail in the following sections.

5.1. Target distribution formulation

The basic idea is to construct an approximate sampling target distribution that places higher importance to the failure regions, to efficiently guide the samples to these domains of interest, and then the probability of failure can be quantified using an inverse importance sampling procedure. Eq. (1) can be computed by directly sampling $I_F(\theta) \pi_\theta(\theta)$. However, this direct approach cannot be practically and effectively used in most general cases since the support domain of the non-smooth indicator function $I_F(\theta)$ is only the failure regions ($\theta \in g(\theta) \leq 0$), making it challenging of locating and adequately sampling the failure domains, especially in cases of high-dimensional and multi-modal spaces. In this work, the indicator function is hence approximated by a one-dimensional output likelihood function, that is based on the limit-state expression $g(\theta)$, supporting the entire domain Θ . This likelihood function, $\ell_{g\theta}$, is expressed as a logistic cumulative distribution function, F_{cdf} , with mean μ_g , and a dispersion factor σ , as:

$$\ell_{g\theta} = F_{cdf} \left(\frac{-g(\theta)}{g_c} \mid \mu_g, \sigma \right) = \frac{1}{1 + e^{-\left(\frac{\left(\frac{g(\theta)}{g_c} \right) + \mu_g}{\left(\frac{\sqrt{3}}{\pi} \right) \sigma} \right)}} \quad (8)$$

where both μ_g , σ , parameters are described in detail in Section 5.1.1, $g(\theta)$ is the limit-state function, and g_c is a scaling constant. Similar approaches to approximate the indicator function can be seen in [24, 29], whereas it was used in a completely different context, focusing on quantifying the probability of failure using a sequential importance sampling method, as well as in problems related to reliability-based design optimization, such as in [71], where several different options are compared, and in [72], where the indicator function is approximated by a cumulative distribution function based on the limit-state function and a related model prediction error. The non-normalized target PDF is then defined as:

$$\tilde{h}(\theta) = \ell_{g\theta} \pi_\theta(\theta) = F_{cdf} \left(\frac{-g(\theta)}{g_c} \mid \mu_g, \sigma \right) \pi_\theta(\theta) \quad (9)$$

and combining Eqs. (8) and (9), the $\tilde{h}(\theta)$ is finally expanded as:

$$\tilde{h}(\theta) = \left(1 + e^{-\left(\frac{\left(\frac{g(\theta)}{g_c} \right) + \mu_g}{\left(\frac{\sqrt{3}}{\pi} \right) \sigma} \right)} \right)^{-1} \pi_\theta(\theta) \quad (10)$$

where $\pi_\theta(\cdot)$ denotes the multivariate standard normal distribution PDF in this work, $\pi_\theta(\theta) = N(\mathbf{0}, \mathbf{I})$, describing the multidimensional variable space Θ . This approximate target distribution $\tilde{h}(\theta)$ is smooth and supports both the safe and failure domains, emphasizing on the failure one, and can be sampled efficiently, particularly by the gradient-based HMCMC samplers which can take informed large jumps across the state space, exploring all regions of interest. The $\mathcal{L}(\theta)$ in Sections 3 and 4 is thus the logarithmic form of the $\tilde{h}(\theta)$ distribution, $\mathcal{L}(\theta) = \ln(\tilde{h}(\theta))$.

The reason for the scaling $g(\theta)/g_c$ is mainly to properly and generally adjust the influence of the limit-state surface on the whole standard normal space Θ , and to standardize parameter values, aiming at a universal algorithm and settings. Eq. (11) defines recommended values for the scaling constant g_c and the suggested q range [3 5] was found to work well in practice. Nonetheless, it is worth mentioning that these

Algorithm 2 Quasi-Newton mass preconditioned Hamiltonian Markov Chain Monte Carlo (QNp-HMCMC)

```

1: procedure QNp-HMCMC( $\theta^0$ ,  $\epsilon$ ,  $L$ ,  $\mathcal{L}(\theta)$ ,  $BurnIn$ ,  $N_{Iter}$ )
2:    $\mathbf{W} = \mathbf{I}$ 
3:   for  $m = 1$  to  $N_{Iter}$  do
4:     if  $m \leq BurnIn$  then
5:        $\mathbf{z}^0 \sim \mathbf{N}(\mathbf{0}, \mathbf{M})$  ▷ where  $\mathbf{M} = \mathbf{I}$ 
6:        $\theta^m \leftarrow \theta^{m-1}$ ,  $\tilde{\theta} \leftarrow \theta^{m-1}$ ,  $\tilde{\mathbf{z}} \leftarrow \mathbf{z}^0$ ,  $\mathbf{B} \leftarrow \mathbf{W}$ 
7:       for  $i = 1$  to  $L$  do
8:          $\tilde{\theta}, \tilde{\mathbf{z}} \leftarrow$  Leapfrog-BurnIn( $\tilde{\theta}, \tilde{\mathbf{z}}, \epsilon, \mathbf{B}$ )
9:         Update  $\mathbf{W}$  using Eq. (5)
10:      end for
11:      with probability:
12:         $\alpha = \min \left\{ 1, \frac{\exp(\mathcal{L}(\tilde{\theta})) - \frac{1}{2} \tilde{\mathbf{z}}^T \tilde{\mathbf{z}}}{\exp(\mathcal{L}(\theta^{m-1})) - \frac{1}{2} \mathbf{z}^{0T} \mathbf{z}^0} \right\}$ 
13:         $\theta^m \leftarrow \tilde{\theta}$ ,  $\mathbf{z}^m \leftarrow -\tilde{\mathbf{z}}$  ▷ If proposal rejected:  $\mathbf{W} \leftarrow \mathbf{B}$ 
14:      else ▷ If  $m > BurnIn$ 
15:         $\mathbf{z}^0 \sim \mathbf{N}(\mathbf{0}, \mathbf{M})$  ▷ where  $\mathbf{M} = \Sigma^{-1} = \mathbf{W}^{-1}$ 
16:         $\theta^m \leftarrow \theta^{m-1}$ ,  $\tilde{\theta} \leftarrow \theta^{m-1}$ ,  $\tilde{\mathbf{z}} \leftarrow \mathbf{z}^0$ 
17:        for  $i = 1$  to  $L$  do
18:           $\tilde{\theta}, \tilde{\mathbf{z}} \leftarrow$  Leapfrog( $\tilde{\theta}, \tilde{\mathbf{z}}, \epsilon, \mathbf{M}$ )
19:        end for
20:        with probability:
21:           $\alpha = \min \left\{ 1, \frac{\exp(\mathcal{L}(\tilde{\theta})) - \frac{1}{2} \tilde{\mathbf{z}}^T \mathbf{M}^{-1} \tilde{\mathbf{z}}}{\exp(\mathcal{L}(\theta^{m-1})) - \frac{1}{2} \mathbf{z}^{0T} \mathbf{M}^{-1} \mathbf{z}^0} \right\}$ 
22:           $\theta^m \leftarrow \tilde{\theta}$ ,  $\mathbf{z}^m \leftarrow -\tilde{\mathbf{z}}$ 
23:        end if
24:      end for
25:    end procedure
26:
27:
28: function LEAPFROG-BURNIN( $\theta, \mathbf{z}, \epsilon, \mathbf{B}$ )
29:    $\tilde{\mathbf{z}} \leftarrow \mathbf{z} + (\epsilon/2) \mathbf{B} \nabla_\theta \mathcal{L}(\theta)$ 
30:    $\tilde{\theta} \leftarrow \theta + \epsilon \mathbf{B} \tilde{\mathbf{z}}$ 
31:    $\tilde{\mathbf{z}} \leftarrow \tilde{\mathbf{z}} + (\epsilon/2) \mathbf{B} \nabla_\theta \mathcal{L}(\tilde{\theta})$ 
32:   return  $\tilde{\theta}, \tilde{\mathbf{z}}$ 
33: end function
34:
35:
36: function LEAPFROG( $\theta, \mathbf{z}, \epsilon, \mathbf{M}$ )
37:    $\tilde{\mathbf{z}} \leftarrow \mathbf{z} + (\epsilon/2) \nabla_\theta \mathcal{L}(\theta)$ 
38:    $\tilde{\theta} \leftarrow \theta + \epsilon \mathbf{M}^{-1} \tilde{\mathbf{z}}$ 
39:    $\tilde{\mathbf{z}} \leftarrow \tilde{\mathbf{z}} + (\epsilon/2) \nabla_\theta \mathcal{L}(\tilde{\theta})$ 
40:   return  $\tilde{\theta}, \tilde{\mathbf{z}}$ 
41: end function

```

values may need further investigation when the proposed framework is directly used in non-Gaussian spaces.

$$g_c = \begin{cases} \frac{g(\mathbf{0})}{q}, q \in [3 5] & \text{if } \left((g(\mathbf{0}) > 7) \cup (0 < g(\mathbf{0}) < 2) \right) \\ 1, & \text{otherwise} \end{cases} \quad (11)$$

Fig. 2 illustrates the described approach in constructing the target distribution $\tilde{h}(\theta)$ in Eqs. (9) and (10) under the ASTPA framework, emphasizing also on the effect of the scaling constant g_c . The gray island-shaped curve is the well-known multi-modal Himmelblau's function (particularly popular in mathematical optimization), seen also in Eq. (30), and it characterizes our limit-state function $g(\theta)$, with the failure domain being inside of the curves and $P_F \sim 2.77 \times 10^{-7}$ in this case. Fig. 2(a) illustrates the bivariate normal distribution $\pi_\theta(\theta)$ which defines the (θ_1, θ_2) space. The limit state function here provides

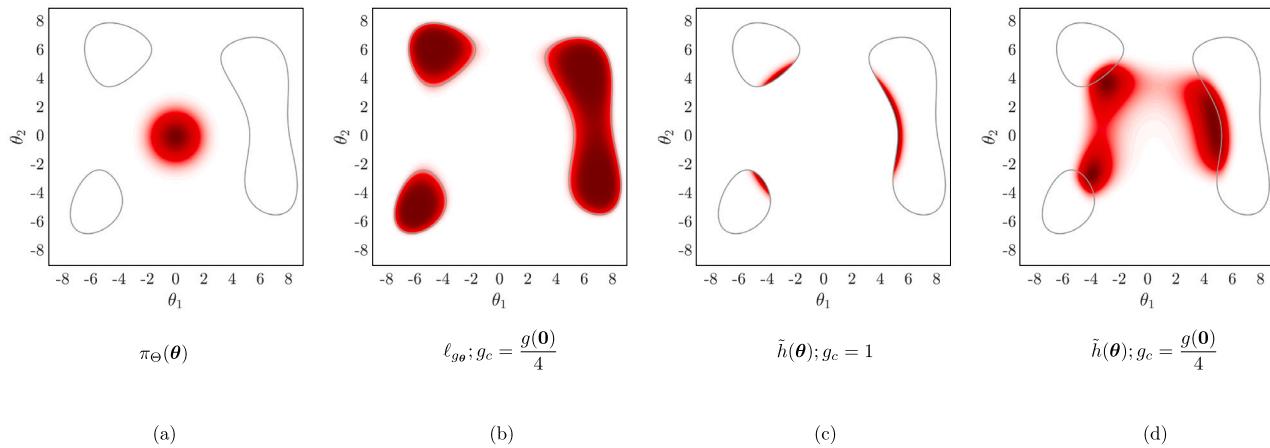


Fig. 2. Demonstrating the concept of constructing a target distribution (Eq. (9) or (10)) as a product of a 1-D likelihood function and a multi-dimensional variable space. From left to right plot, $\pi_\theta(\theta)$ represents the variable space, ℓ_{g_θ} is the likelihood function, $\tilde{h}(\theta)$ is the non-normalized target distribution computed without the proper value of g_c , and $h(\theta)$ is the correct non-normalized target distribution, respectively.

$g(\mathbf{0}) > 7$, so the scaling constant g_c can be defined as $g(\mathbf{0}) / 4$, following Eq. (11). Fig. 2(b) portrays our prescribed 1-D likelihood function ℓ_{g_θ} , described by a logistic cumulative distribution function, as shown in Eq. (8). Fig. 2(c) shows the non-normalized target PDF $\tilde{h}(\theta)$ determined as a product of $\pi_\theta(\theta)$ and ℓ_{g_θ} , without considering however the scaling constant g_c here. Fig. 2(d) instead represents the non-normalized target PDF $\tilde{h}(\theta)$ constructed with the suggested value of g_c . These two figures can thus clearly explain the role of the scaling constant g_c in attracting samples to the regions of interest.

5.1.1. Impact of mean (μ_g) and dispersion factor (σ)

The use of σ in this work, largely follows ideas presented in [24], albeit at a completely different context. The value of σ also has an important effect on the convergence and the computational efficiency of our HMCMC sampling methods. In Fig. 3 the effect of σ on the target distribution is displayed based on two 2D examples on the Θ space, with a small failure probability ($\sim 10^{-6}$) and a unimodal failure region. A higher value of σ can reduce the number of model calls as the constructed target becomes more dispersed compared to the case of lower σ values, thus being comparatively easier to be located and sampled, but an insufficient number of samples may then reach inside the failure regions. In contrast, reducing σ concentrates the target distribution within the regions of interest, making it, however, a more challenging sampling target, often requiring a higher number of model calls to be located and sampled. Choosing the value of the likelihood dispersion factor, σ , is therefore a trade-off. The suggested value of σ for our methodology is in the range [0.1 0.8]. Fine tuning higher decimal values of σ in that range is not usually necessary. It is generally recommended to use higher σ values (0.5 – 0.8) in multi-modal cases, enabling longer state jumps, even between modes, and lower values, (0.1–0.6), for cases with very small failure probability, e.g., around $P_F \leq 10^{-6}$ and lower, and/or high-dimensional problems, in order to attract more samples into the failure domain.

For the mean parameter, μ_g , we are interested in generally locating it into the failure region, $g(\theta) < 0$, to enhance the sampling efficiency. As such, we are describing μ_g through a percentile, p , of the logistic CDF and its quantile function, Y_g :

$$Y_g(p; \mu_g, \sigma) = \mu_g + \left(\frac{\sqrt{3}}{\pi} \sigma \right) \ln \left(\frac{p}{1-p} \right) \quad (12)$$

Placing a chosen percentile p of the logistic CDF on the limit-state surface $g(\theta) = 0$ results in $Y_g(p; \mu_g, \sigma) = 0$ and the μ_g can be then given as:

$$\mu_g = - \left(\frac{\sqrt{3}}{\pi} \sigma \right) \ln \left(\frac{p}{1-p} \right) \quad (13)$$

Therefore, as also shown in Fig. 4, by reducing p , the μ_g of the likelihood function further moves into the failure domain and consequently more samples can infiltrate into this region and can contribute to the failure probability estimation. In Fig. 4, p_{50} , for example, denotes the 50th percentile ($p = 0.5$). In all examples presented in this paper, the p_{10} percentile, $p = 0.1$, is used to define the value μ_g , which is found to yield good efficiency. Eq. (13) and parameter σ can now be used in Eq. (10) to fully define the non-normalized target distribution $\tilde{h}(\theta)$.

Fig. 5 displays the effect of σ on sampling the approximate target distribution, showcasing how smaller values of σ push the samples further inside the failure regions. It also compares the $\mathbb{E}[\hat{P}_F]$ and Coefficient of Variation (C.o.V), computed according to Eqs. (16) and (26), based on 500 independent simulations with a number of model calls ~ 900 in all cases. Smaller values of σ result in a more accurate estimate of $\mathbb{E}[\hat{P}_F]$ and a lower C.o.V, due to having the presence of more samples in the failure region.

5.2. Inverse importance sampling

An original post-sampling step is devised at this stage, termed *inverse importance sampling* (IIS), that can successfully employ the already acquired samples and the normalized target distribution $h(\theta)$ as the importance sampling density function. Hence, the probability of failure can be now computed as follows:

$$P_F = \int_{\theta} I_F(\theta) \pi_\theta(\theta) d\theta = \int_{\theta} I_F(\theta) \frac{\pi_\theta(\theta)}{h(\theta)} h(\theta) d\theta \quad (14)$$

$$\text{where } h(\theta) = \frac{\tilde{h}(\theta)}{C_h}, \quad \tilde{h}(\theta) = \ell_{g_\theta}(\theta) \pi_\theta(\theta), \quad (15)$$

$\pi_\theta(\theta)$ is the original distribution, $\tilde{h}(\theta)$ denotes the non-normalized sampling target distribution, and C_h is its normalizing constant. After some simplifications we get:

$$P_F = C_h \int_{\theta} I_F(\theta) \frac{1}{\ell_{g_\theta}(\theta)} h(\theta) d\theta = \left(\frac{1}{N} \sum_{i=1}^N \frac{I_F(\theta_i)}{\ell_{g_\theta}(\theta_i)} \right) C_h \quad (16)$$

with N the number of already acquired samples. To now calculate the normalizing constant C_h , once again we resort to an importance sampling scheme in the following manner:

$$C_h = \int_{\theta} \tilde{h}(\theta) d\theta = \int_{\theta} \frac{\tilde{h}(\theta)}{Q(\theta)} Q(\theta) d\theta = \frac{1}{M} \sum_{i=1}^M \frac{\tilde{h}(\theta'_i)}{Q(\theta'_i)} \quad (17)$$

where $Q(\cdot)$ can be a computed Gaussian Mixture Model (GMM), based again on the already available samples and the generic Expectation-Maximization (EM) algorithm [73], as indicatively seen in the right plot

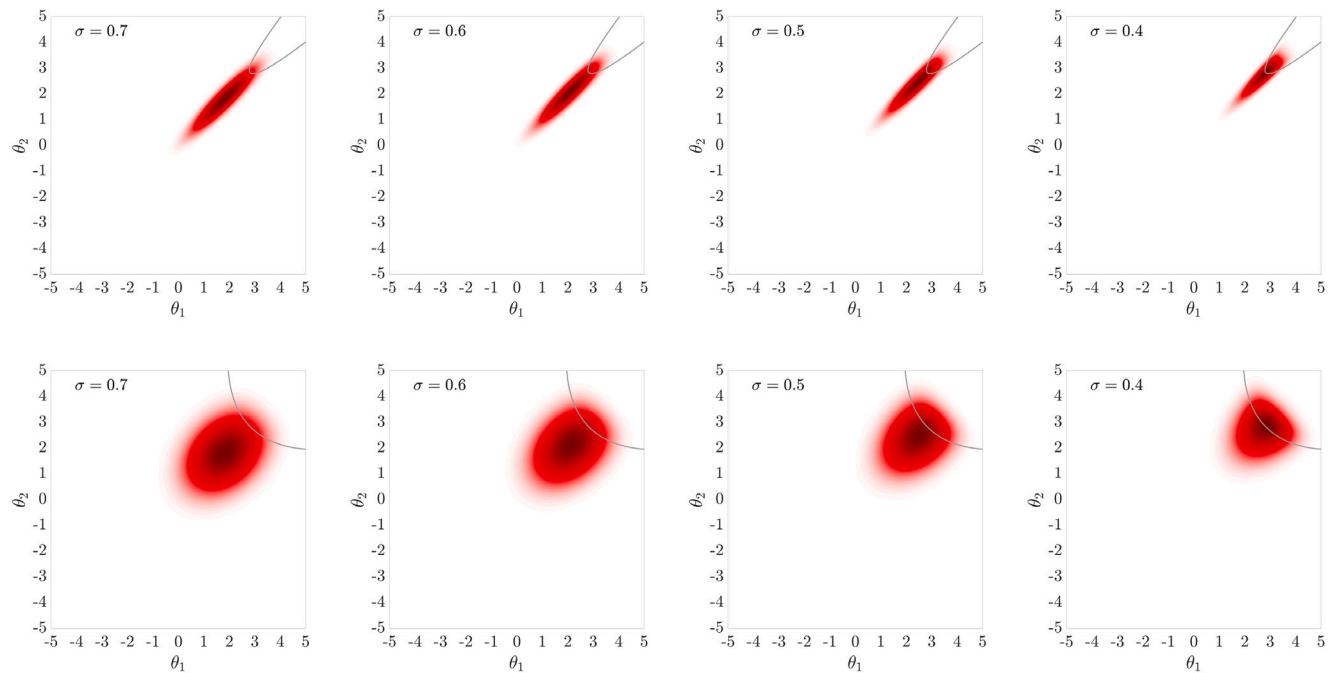


Fig. 3. Effect of the likelihood dispersion factor, σ , on the target distribution for two different unimodal limit-state functions, for the $p50$ percentile parameter of μ_g .

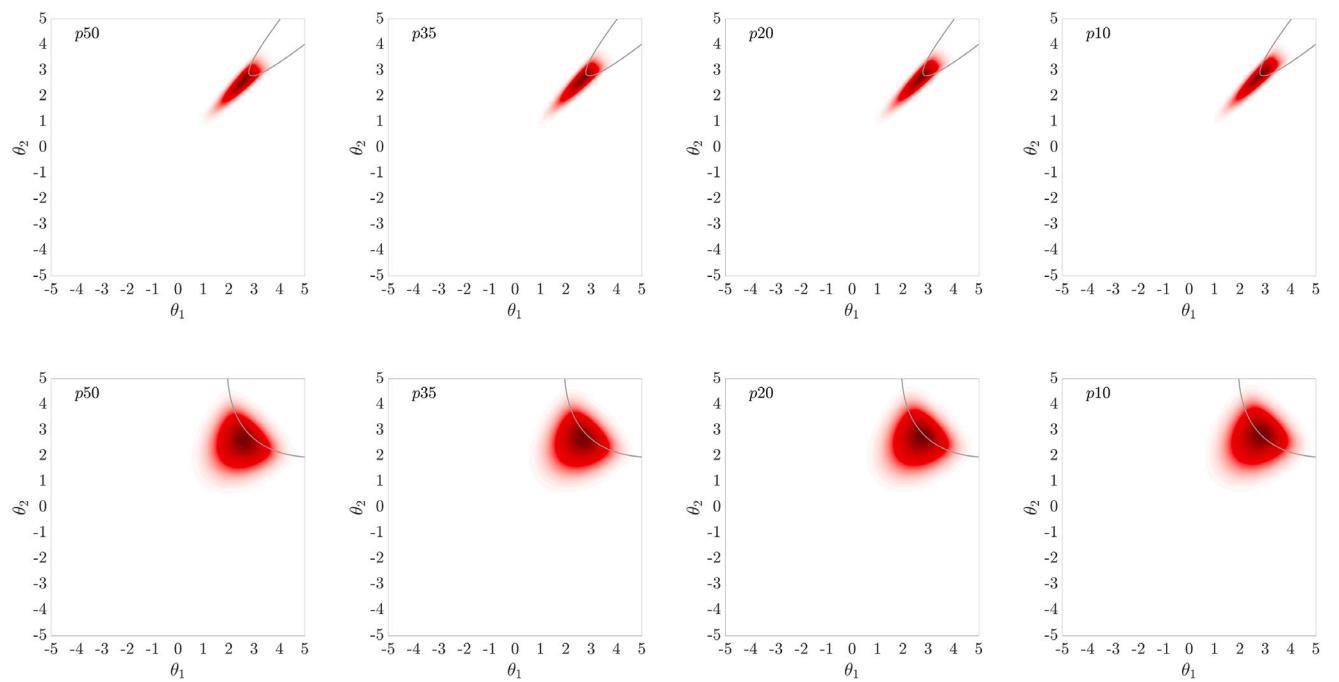


Fig. 4. Effect of μ_g of the likelihood function on the target distribution, expressed through its percentile parameter p , for two different unimodal limit-state functions, and for $\sigma = 0.4$.

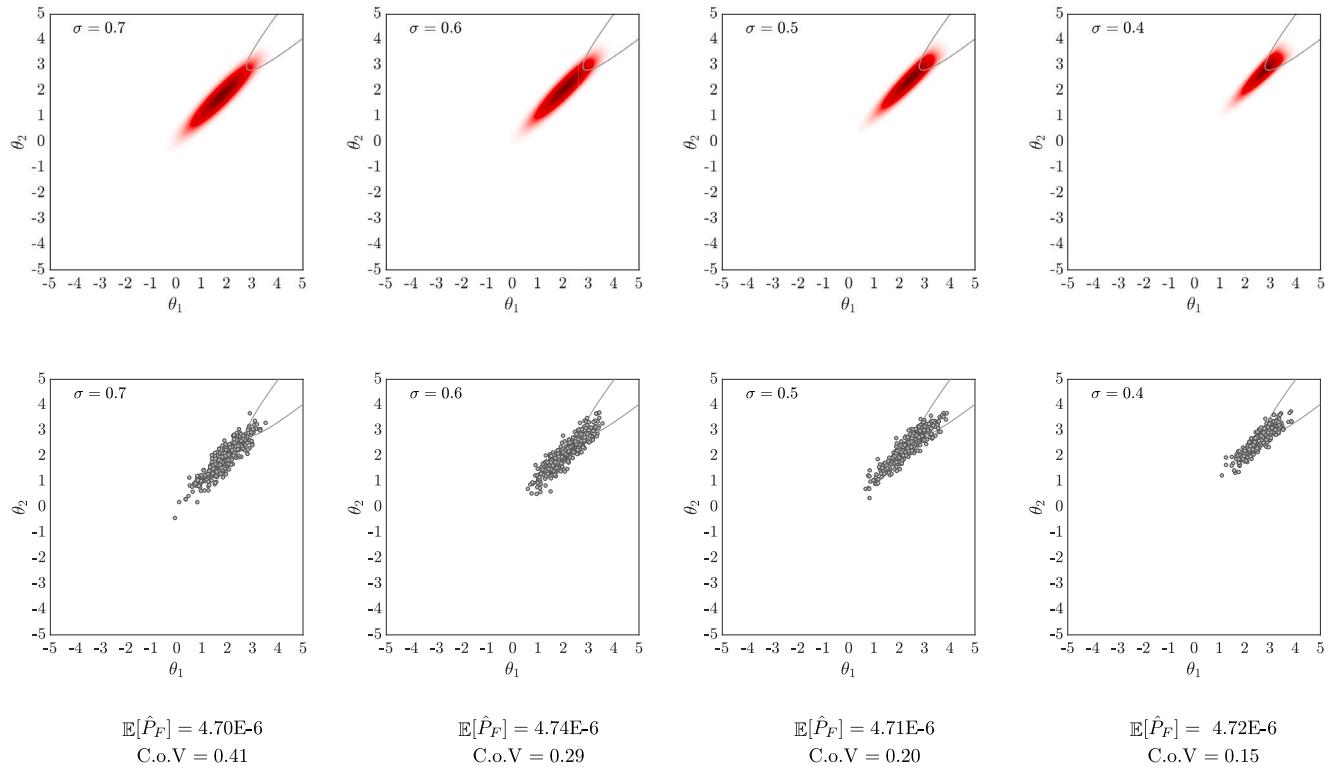


Fig. 5. Effect of the prescribed σ on sampling from the target distribution; μ_g is calculated based on $p10$; (Ref. $P_F \sim 4.73E-6$).

of Fig. 1. A typical GMM expression can then be given by:

$$Q(\theta) = \sum_{k=1}^K w_k \phi(\theta; \mu_k, \Sigma_k) \quad (18)$$

where $\phi(\cdot)$ is the PDF, w_k is the weight, μ_k is the mean vector and Σ_k is the covariance matrix of the k^{th} Gaussian component, that can all be estimated, for all components, by the EM algorithm [73] based on the samples. When a GMM can be accurately fitted, the original HMC samples can be used directly to compute the normalizing constant C_h , i.e., $\theta'_i \equiv \theta_i$, and $M = N$ in Eqs. (16) and (17). A precisely fitted GMM, with a sufficiently large number of components, adequately acts as a representative distribution for the original samples. However, given that the accuracy of GMMs may often deteriorate, particularly in high-dimensional and challenging multi-modal cases, due to the large number of parameters that need to be identified [29], similar to many other mixture models, additional M samples from the computed $Q(\cdot)$ can be required, just in order to accurately evaluate the normalizing constant C_h . In this latter, more general case, even a very approximately fitted GMM with diagonal covariance matrices, particularly appropriate for high-dimensional cases, works effectively, since it is solely used for computing the constant C_h , while the already obtained HMC samples are instead used in Eq. (16). Drawing each of the i.i.d. $Q(\cdot)$ samples requires only one model call and is called IIS sample in the following sections for the sake of clarity and simplicity. As a general guide, using IIS samples around 20% of the total number of samples results in a good estimate for C_h and eventually Eq. (16).

5.3. Adaptation during the burn-in phase

A burn-in sampling phase is required in the ASTPA framework, regardless of the used MCMC sampling scheme. For the presented HMC methods in this work, the burn-in samples are usually adequate to be around 5%–10% of the total number of samples. To improve

the overall sampling efficiency, we are taking advantage of this burn-in phase, that guides the chain to the appropriate target distribution, by also performing a series of adaptations and relevant parameter tuning steps. In particular, the stepsize ϵ and the trajectory length τ are established in this burn-in stage, as explained in Section 3.2, and then used throughout the rest of the analysis, while the preconditioned mass matrix \mathbf{M} , used in our QNp-HMC algorithm, is also computed in the burn-in phase and then utilized for the remainder of the sampling, as discussed in Section 4.

In addition to these sampling related parameters, we are also adjusting the target distribution in this sampling stage, through the parameters μ_g and σ . Following the discussion in Section 5.1.1, these two parameters are automatically evolving toward their constant values in this adaptive phase, to assist in guiding the samples to the final target distribution. The likelihood dispersion factor initiates with $\sigma_0 = 1$ and then follows an exponential decay until its prescribed constant value σ , as:

$$\sigma_{iter} = a_1 e^{\left(\frac{-iter}{a_2}\right)}, \quad \text{where } a_2 = \frac{(N_{BI} - 1)}{\ln(\frac{\sigma_0}{\sigma})}, \quad a_1 = \frac{\sigma_0}{e^{\left(\frac{-1}{a_2}\right)}} \quad (19)$$

where N_{BI} is the assigned number of total burn-in samples and $iter \geq 1$ denotes the iteration counter based on a fully completed leapfrog step. A similar exponential growth scheme is also used for μ_g , starting based on $p50$ in Eq. (13), i.e. $\mu_{p50} = 0 = g(\theta)$, and reaching the constant value $\mu_g = \mu_{p10}$ given by the prescribed $p10$ percentile in Eq. (13), as:

$$\mu_{iter} = b_1 e^{\left(\frac{-iter}{b_2}\right)} + \mu_{p50}, \quad \text{where } b_2 = \frac{(N_{BI} - 1)}{\ln(\frac{0.0001}{\mu_t})}, \quad (20)$$

$$b_1 = \frac{0.0001}{e^{\left(\frac{-1}{b_2}\right)}}, \quad \mu_t = \mu_{p10} + \mu_{p50}$$

Table 1
Summary of ASTPA parameters and generic values beyond tuning/optimization.

Dispersion factor (σ)	$\mu_{g(\theta)}$	Trajectory length (τ)	Step size (ϵ)	Total model calls	Burn-in	IIS samples
[0.1 0.8]	p10	0.7	<i>Dual averaging</i> [61]	5,000–10,000	10%	20%

5.4. Statistical properties of the estimator P_F

Based on Eq. (16) and given samples $[\theta_1, \dots, \theta_N]$, the sample estimator \hat{P}_F is expressed:

$$\hat{P}_F = \left(\frac{1}{N} \sum_{i=1}^N \frac{I_F(\theta_i)}{\ell_{g\theta}(\theta_i)} \right) \left(\frac{1}{M} \sum_{j=1}^M \frac{\tilde{h}(\theta'_j)}{Q(\theta'_j)} \right) = \tilde{P}_F \hat{C}_h \quad (21)$$

where

$$\tilde{P}_F = \left(\frac{1}{N} \sum_{i=1}^N \frac{I_F(\theta_i)}{\ell_{g\theta}(\theta_i)} \right) \text{ and } \hat{C}_h = \left(\frac{1}{M} \sum_{j=1}^M \frac{\tilde{h}(\theta'_j)}{Q(\theta'_j)} \right) \quad (22)$$

with variances:

$$\sigma^2(\tilde{P}_F) = \frac{1}{N_s(N_s - 1)} \sum_{i=1}^{N_s} \left(\frac{I_F(\theta_i)}{\ell_{g\theta}(\theta_i)} - \tilde{P}_F \right)^2 \quad (23)$$

$$\sigma^2(\hat{C}_h) = \frac{1}{M(M - 1)} \sum_{j=1}^M \left(\frac{\tilde{h}(\theta'_j)}{Q(\theta'_j)} - \hat{C}_h \right)^2 \quad (24)$$

where N_s denotes the used Markov chain samples, taking into account the fact that the samples are not independent and identically distributed (i.i.d) in this case.

Assuming \tilde{P}_F and \hat{C}_h are independent random variables, the variance of \hat{P}_F can be given as:

$$\sigma^2(\hat{P}_F) = \sigma^2(\tilde{P}_F) \sigma^2(\hat{C}_h) + \sigma^2(\tilde{P}_F) \hat{C}_h^2 + \tilde{P}_F^2 \sigma^2(\hat{C}_h) \quad (25)$$

The Coefficient of Variation (C.o.V) can then be provided as:

$$\text{C.o.V} \approx \frac{\sigma(\hat{P}_F)}{\hat{P}_F} \quad (26)$$

HMCMC samplers typically showcase low autocorrelation and thus thinning the sample size from N to N_s in Eq. (23), to enhance independence, is often not required. In any case, N_s can be easily determined by examining the sample autocorrelation and choosing an appropriate thinning lag, if needed. In this work, for the C.o.V calculation and N_s in Eq. (23), we used every 3rd sample for all examples. The same thinning process can also be used for \tilde{P}_F in Eq. (22), if wanted, although we have not done this in this work and all acquired N samples have been used for the probability of failure estimation.

5.5. Summary of the ASTPA parameters

The required input parameters for the presented methodology are summarized in Table 1, together with some suggested generic values for reliability estimation problems. The constant likelihood dispersion factor σ follows the suggestions in Section 5.1.1 and can generally be in the range of [0.1 0.8]. The mean, μ_g , is chosen to be provided by the p10 percentile in Eq. (13), and the trajectory length τ can be based on the ESJD metric, as explained in Section 3.2, with a generic value being 0.7. The dual averaging algorithm of [61] is adopted here to automatically provide the used constant step size ϵ in the non-adaptive sampling phase, and the required minimum number of model calls for the QNp-HMCMC method is roughly suggested to be 5,000–10,000 for high-dimensional problems and target probabilities lower than 10^{-4} . Naturally, the required number of model calls and samples is case dependent and convergence of the estimator can be checked through Eqs. (16), (23), and/or (24). For the burn-in phase less than 10% of the total number of samples are often required, so a 10% value can be generally suggested. Finally, to compute the normalizing constant C_h of

the approximate target, we can generally use IIS samples around 20% of the total number of samples.

6. Numerical results

Several numerical examples are studied in this section to examine the performance and efficiency of the proposed methods. In all examples, input parameters follow the provided guidelines in Section 5.5. To compute the normalizing constant C_h , IIS samples around 20% of the total number of samples have been drawn from a computed GMM with diagonal covariance matrices and, generally, 10 and 1 Gaussian components for low and high dimensional problems, respectively. Results are compared with the Component-Wise Metropolis-Hastings based Subset Simulation (CWMH-SuS) [21], with two proposal distributions, a uniform one of width 2 and a standard normal one. The adaptive Conditional Sampling (aCS) SuS variant introduced in [16] is also used in all examples, implemented based on the online provided code by the authors [74]. In all examples, the SuS parameters are chosen as $n_s = 1,000$ and $2,000$ for low- and high-dimensional problems, when needed, respectively, with n_s the number of samples in each subset level, and $p_0 = 0.1$, where p_0 is the samples percentile for determining the appropriate subsets. Comparisons are provided in terms of accuracy and computational cost, reporting the P_F estimation and the mean number of limit-state function calls. Analytical gradients are provided in all examples, hence one limit-state/model call can provide both the relevant function values and gradients. The number of limit-state function evaluations for the HMCMC-based methods has been determined based on reported C.o.V values $\in [0.1, 0.35]$, as estimated by 500 independent simulations. C.o.V values estimated by the analytical expression in Eq. (26) are also reported in parenthesis, for both HMCMC and QNp-HMCMC methods. The reference failure probabilities are obtained based on the mean estimation of 100 independent simulations using 10^8 crude Monte Carlo samples, where applicable for higher reference probabilities, and SuS method with $n_s = 100,000$ for smaller reference probabilities. The problem dimensions are denoted by d and all ASTPA parameters are carefully chosen for all examples but are not optimized for any one. Comparative and perhaps improved alternate performance might thus be achieved with a different set of parameters.

6.1. Example 1: Nonlinear convex limit-state function

The first example consists of a nonlinear convex limit-state function characterized by two independent standard normal random variables and a low failure probability ($P_F \sim 4.73\text{E-}6$):

$$g(\theta) = 4 - \frac{1}{\sqrt{2}} (\theta_1 + \theta_2) + 2.5 (\theta_1 - \theta_2)^2 \quad (27)$$

Table 2 compares the number of model calls, C.o.V and $\mathbb{E}[\hat{P}_F]$ obtained by SuS and the two HMCMC algorithms. For the HMCMC algorithms, the trajectory length τ is chosen equal to the default value of 0.7 and the likelihood dispersion factor, σ , is 0.4. The burn-in sample size is set to 150 samples. As shown, the suggested approaches perform noticeably better than SuS, and aCS-SuS results outperform the original SuS approach with two different proposal distributions. The QNp-HMCMC method exhibits an excellent performance with only 844 model calls in this problem, on average. The simulated QNp-HMCMC samples and the analytical target distribution based on the limit-state function of Eq. (27) can be seen in the two figures on the right hand side in Fig. 5.

Table 2Performance of various methods for the nonlinear convex limit-state function in Example 1 ($d = 2$).

500 independent simulations		CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
		$U(-1, 1)$	$N(0, 1)$			
$\sigma = 0.4$	Number of total model calls	5,462	5,552	5,453	1,873	836
	C.o.V	1.06	1.55	0.94	0.14(0.15)	0.15(0.15)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 4.73\text{E-}6$)	4.88E-6	5.11E-8	4.52E-6	4.73E-6	4.72E-6

Table 3Performance of various methods for the parabolic/concave limit-state function in Example 2 ($d = 2$).

500 independent simulations		CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
		$U(-1, 1)$	$N(0, 1)$			
$\sigma = 0.7$	Number of total model calls	4,559	4,565	4,562	3,306	3,306
	C.o.V	0.62	0.65	0.63	0.09(0.06)	0.09(0.06)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 3.95\text{E-}5$)	4.19E-5	4.14E-5	4.09E-5	3.93E-5	3.88E-5

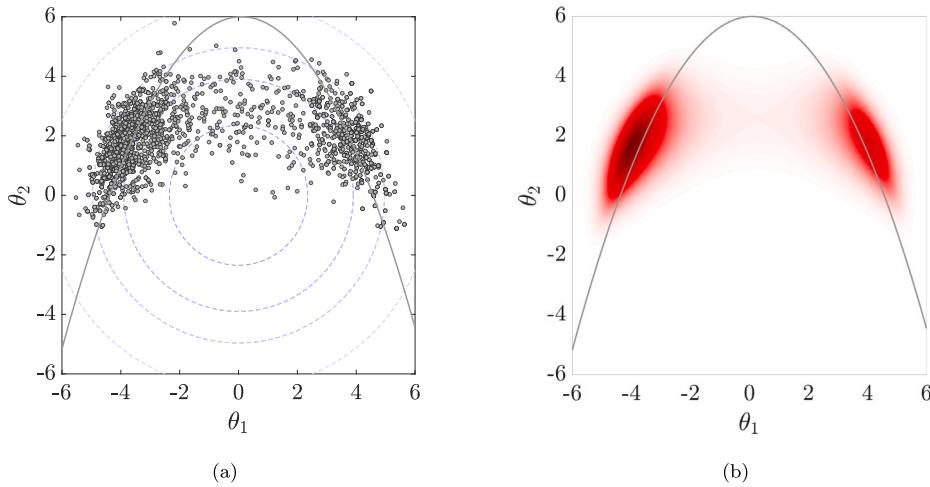


Fig. 6. Example 2: (a) Simulated samples from the target distribution. (b) Analytical target distribution.

6.2. Example 2: Parabolic/concave limit-state function

This example is based on the following limit state function with two standard normal random variables [75]:

$$g(\theta) = r - \theta_2 - \kappa(\theta_1 - e)^2 \quad (28)$$

where r , κ and e are deterministic parameters chosen as $r = 6$, $\kappa = 0.3$ and $e = 0.1$. The probability of failure is $3.95\text{E-}5$ and the limit-state function describes two main failure modes, as also seen in Fig. 6. For the HMCMC-based algorithms, the likelihood dispersion factor, σ , is 0.7, the burn-in sample size is 200 samples, and the trajectory length is set to $\tau = 1$.

Table 3 compares the number of model calls, the C.o.V and the $\mathbb{E}[\hat{P}_F]$ obtained by all methods. The HMCMC-based approaches provide significantly smaller C.o.V values than SuS, with fewer model calls, and the HMCMC in particular outperforms all other methods in this two dimensional problem. The QNp-HMCMC samples are also shown in Fig. 6 and can accurately represent the two important failure regions. The circular dash lines in Fig. 6(a) represent the probability contour of the standard normal space Θ , for probability levels of 10^{-2} , 10^{-4} , 10^{-6} and 10^{-10} , similarly provided for Figs. 7 and 8 as well.

6.3. Example 3: Quartic bimodal limit-state function

The third example is a quartic bimodal limit-state function with very low probability of failure ($P_F \sim 5.90\text{E-}8$), defined by the following limit-state function in the standard normal space:

$$g(\theta) = 6.5 - \frac{1}{\sqrt{2}}(\theta_1 + \theta_2) - 2.5(\theta_1 - \theta_2)^2 + (\theta_1 - \theta_2)^4 \quad (29)$$

Table 4 compares the performance of the HMCMC methods with the ones by SuS. The trajectory length is chosen $\tau = 0.7$, and the likelihood dispersion factor, σ , is 0.5. The burn-in sample size is set to 200 samples. HMCMC-based methods, particularly QNp-HMCMC, once more provide significantly lower C.o.V values with fewer model calls than SuS approaches, that perform rather poorly in this example, with aCS-SuS performing the best among them. Fig. 7 displays related target distribution based on the limit-state function of Eq. (29), as well as the samples generated by the QNp-HMCMC method.

6.4. Example 4: The Himmelblau function

In nonlinear optimization, a commonly used fourth order polynomial test function is the so-called Himmelblau [76] function. Here we adopt and modify this function, as:

$$g(\theta_1, \theta_2) = \left(\frac{(0.75\theta_1 - 0.5)^2}{1.81} + \frac{(0.75\theta_2 - 0.5)^2}{1.81} - 11 \right)^2 + \left(\frac{(0.75\theta_1 - 1)^2}{1.81} + \frac{(0.75\theta_2 - 0.5)^2}{1.81} - 7 \right)^2 - \beta \quad (30)$$

which is particularly suitable for reliability examples with multiple separated failure domains. θ_1 and θ_2 are assumed to be independent standard normal random variables and the constant β is used to define different levels of the failure probability. Table 5 compares the number of model calls, coefficient of variation and $\mathbb{E}[\hat{P}_F]$ obtained by SuS and the family of HMCMC algorithms. For the HMCMC algorithms, the trajectory length is chosen as $\tau = 1$ and the burn-in is set to 200 samples. The $g(\theta)$ is beyond the upper bound, $g(\theta) > 7$, and as discussed in Section 5, we perform the scaling with $g_c = g(\theta)/4$. The likelihood

Table 4Performance of various methods for the quartic bimodal limit-state function in Example 3 ($d = 2$).

	500 independent simulations	CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
		$U(-1, 1)$	$N(0, 1)$			
$\sigma = 0.5$	Number of total model calls	7,327	7,536	7,380	6,277	2,696
	C.o.V	1.64	2.45	1.59	0.28(0.27)	0.27(0.27)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 5.90\text{E-}8$)	6.13E-8	5.86E-8	6.12E-8	5.90E-8	5.89E-8

Table 5Performance of various methods for the Himmelblau limit-state function in Example 4 ($d = 2$).

	500 independent simulations	CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
		$U(-1, 1)$	$N(0, 1)$			
$\beta = 95$	Number of total model calls	3,833	3,833	3,821	3,100	3,100
	C.o.V	0.42	0.50	0.35	0.10(0.06)	0.13(0.06)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.65\text{E-}4$)	1.69E-4	1.78E-4	1.68E-4	1.64E-4	1.62E-4
$\beta = 50$	Number of total model calls	6,471	6,528	6,463	3,600	3,600
	C.o.V	0.87	1.67	0.54	0.16(0.10)	0.16(0.10)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 2.77\text{E-}7$)	3.14E-7	3.33E-7	2.89E-7	2.77E-7	2.77E-7

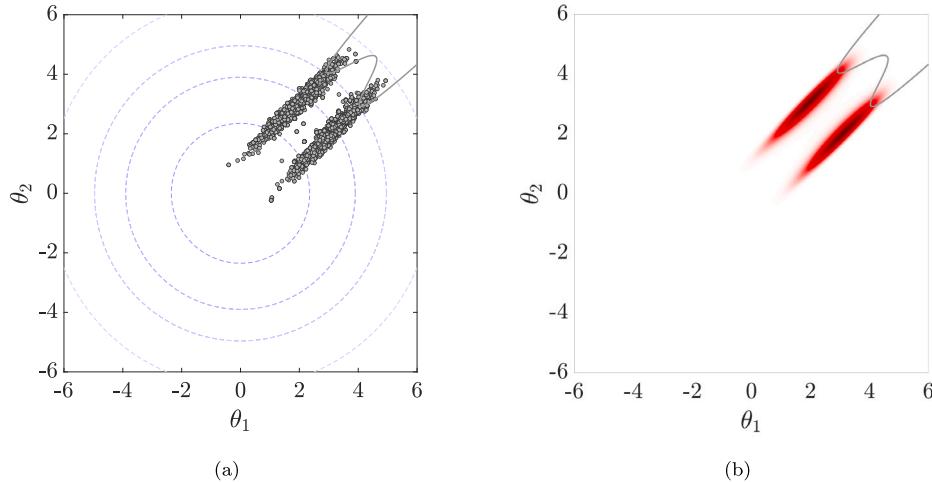


Fig. 7. Example 3: (a) Simulated samples from the target distribution. (b) Analytical target distribution.

dispersion factor, σ , used in this example is mentioned in Table 5, and has been chosen according to the guidelines in Section 5.1.1, considering both multimodality and the different levels of the failure probabilities. The Subset Simulation results are based on $n_s = 1,000$.

It is again shown here that the HMCMC approach gives significantly smaller C.o.V than SuS and also outperforms it in terms of the $\mathbb{E}[\hat{P}_F]$ estimation. Fig. 8 also demonstrates that the HMCMC samples accurately describe all the three important failure regions.

6.5. Example 5: Cantilever beam

In this last two-dimensional example, a cantilever beam problem is studied [77]. The beam is illustrated in Fig. 9, with cross-section width w , height t , beam length L , and transverse loads P_x and P_y . The beam failure mode in this case is the maximum deflection exceeding the allowable value, Y_0 , given by the limit-state function:

$$g(P_x, P_y) = Y_0 - \frac{4L^3}{Ewt} \sqrt{\left(\frac{P_y}{t^2}\right)^2 + \left(\frac{P_x}{w^2}\right)^2} \quad (31)$$

where $Y_0 = 4.2 \text{ in}$ and 4.5 in , $E = 30 \times 10^6 \text{ psi}$ is the elastic modulus, $L = 100 \text{ in}$, $w = 2 \text{ in}$ and $t = 4 \text{ in}$. P_x and P_y follow independent normal distributions $P_x \sim N(500, 100)/\text{lb}$ and $P_y \sim N(1000, 100)/\text{lb}$.

The normally distributed variables P_x and P_y are transformed into the standard normal space and the limit-state function in the Θ space

becomes:

$$g(\theta_x, \theta_y) = Y_0 - \frac{4L^3}{Ewt} \sqrt{\left(\frac{\mu_y + \theta_y \sigma_y}{t^2}\right)^2 + \left(\frac{\mu_x + \theta_x \sigma_x}{w^2}\right)^2} \quad (32)$$

For the HMCMC-based algorithms, the trajectory length is chosen $\tau = 0.7$ and the likelihood dispersion factor, σ , is 0.2. The burning-in sample size is set to 200 samples. Table 6 compares the number of model calls, C.o.V and $\mathbb{E}[\hat{P}_F]$ for two different failure probability levels obtained by SuS and the HMCMC-based methods. As shown, the HMCMC-based approaches noticeably again outperform the SuS methods.

6.6. Example 6: High-dimensional reliability example with linear limit-state function

In this first high-dimensional example, a linear limit-state function of independent standard normal random variables is considered:

$$g(\theta) = \beta - \frac{1}{\sqrt{d}} \sum_{i=1}^d \theta_i \quad (33)$$

where d denotes again the related problem dimensions, and the pertinent probability of failure is equal to $\Phi(-\beta)$, independent of d , with $\Phi(\cdot)$ the cumulative distribution function of the standard normal distribution. To investigate the effect of different failure probability levels,

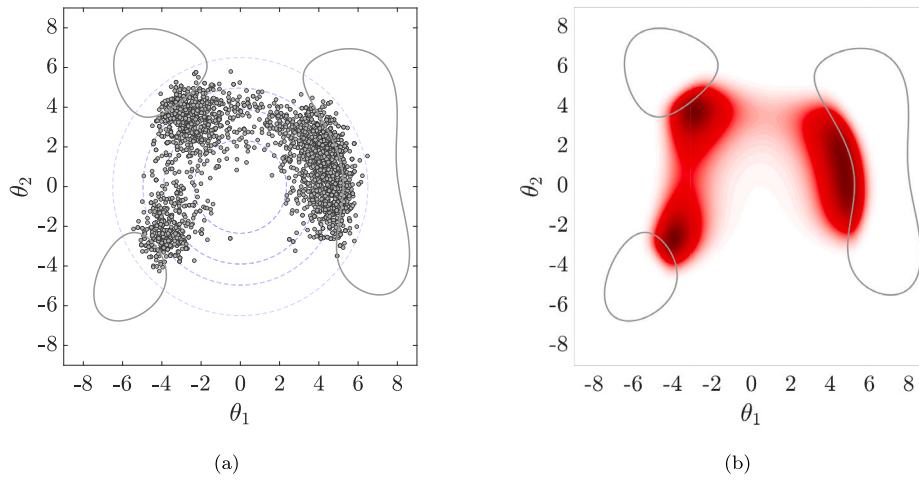


Fig. 8. Example 4: (a) Simulated samples from the target distribution. (b) Analytical target distribution.

Table 6

Performance of various methods for the cantilever beam in Example 5 ($d = 2$).

	500 independent simulations	CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
		$U(-1, 1)$	$N(0, 1)$			
$\sigma = 0.2$	Number of total model calls	6,056	6,069	6,062	1,900	1,900
$Y_0 = 4.2$	C.o.V	0.80	0.93	0.50	0.15(0.17)	0.14(0.17)
$\tau = 0.7$	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.01\text{E-}6$)	1.07E-6	1.08E-6	1.05E-6	1.01E-6	1.01E-6
$\sigma = 0.2$	Number of total model calls	7,561	7,586	7,569	3,200	3,200
$Y_0 = 4.5$	C.o.V	1.11	1.41	0.60	0.19(0.24)	0.19(0.24)
$\tau = 0.7$	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.97\text{E-}8$)	2.14E-8	2.22E-8	2.03E-8	1.97E-8	1.98E-8

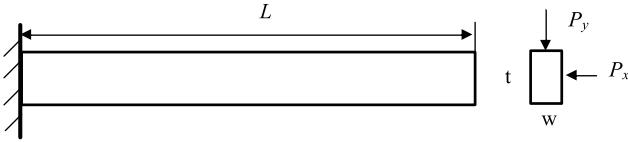


Fig. 9. A cantilever beam case study.

in relation to the HMCMC-based methods and SuS performance, a sequence of β values for $d = 100$ and 500 are studied. Table 7 summarizes the comparative results through the mean number of model calls, mean $\mathbb{E}[\hat{P}_F]$ and C.o.V of the estimated probabilities. For both HMCMC-based methods, the likelihood dispersion factor, σ , is equal to 0.3, the burn-in sample size is 300, and the trajectory length, τ , for all failure probability levels is chosen equal to 0.7, the default value. SuS results are again based on $n_s = 1,000$.

As reported in Table 7, the HMCMC-based approaches exhibit for all probability levels accurate and stable results in terms of C.o.V, outperforming all SuS results. Defining the “unit C.o.V” eff as $C.o.V = eff / \sqrt{N_{mc}}$, with N_{mc} the total number of model calls, an index can now be used that appropriately considers both accuracy and computational effort [78], with a lower eff value exhibiting higher efficiency. Fig. 10a displays the eff variation in relation to the reliability index for $d = 100$, confirming the HMCMC approaches efficiency for various failure probability levels. SuS-U and SuS-N in the figure stand for the SuS results with uniform and standard normal proposals, respectively. Fig. 10b showcases the mean estimate for $\beta = 5$ and $d = 100$, based on the number of model calls. As shown, the HMCMC methods provide a consistent unbiased estimator after a certain, relatively small, number of model evaluations. Finally, in Fig. 10c a similar plot is provided for the C.o.V results, with the proposed HMCMC-based framework exhibiting again excellent overall performance. The C.o.V-Anal curve in the figure represents the QNp-HMCMC C.o.V estimation based on the analytical expression in Eq. (26).

6.7. Example 7: High-dimensional problem with quadratic nonlinearity

This example involves a quadratic limit-state function expressed in the standard normal space, as:

$$g(\theta) = \lambda - \frac{1}{\sqrt{d}} \sum_{i=1}^d \theta_i + 2.5 \left(\theta_1 - \sum_{j=2}^d \theta_j \right)^2 \quad (34)$$

where d is the problem dimension, λ defines the level of the failure probability, and γ affects the level of nonlinearity. To investigate the effect of dimensionality and nonlinearity, in relation to the HMCMC-based methods and SuS performance, different γ values for $d = 100$ and 200 are studied. Table 8 presents the mean number of limit-state function evaluations, sample mean estimator, and C.o.V of the failure probabilities for various methods. For the HMCMC-based methods, the likelihood dispersion factor, σ , and the trajectory length, τ , are shown in Table 8, and the burn-in sample size is set to 500 samples. The Subset Simulation results for all approaches are based on $n_s = 2,000$.

As seen in Table 8, the QNp-HMCMC approach significantly outperforms all other methods, in all considered metrics, confirming its superiority, robustness, applicability and suitability in challenging, nonlinear, high-dimensional problems. By increasing the level of nonlinearity, through the parameter γ , HMCMC requires an excessive number of model calls in order to adequately explore the domain, while the QNp-HMCMC approach exhibits a largely invariant performance.

In Fig. 11a, the computed eff values with respect to γ for $d = 100$ are reported, based on Table 8 results, confirming the QNp-HMCMC efficiency and robust behavior in all considered cases of nonlinearity. Fig. 11b and Fig. 11c study the effect of the number of model calls (N_{mc}) on the mean estimate and C.o.V values for the case of $d = 100$ and $\gamma = 100$. The QNp-HMCMC method exhibits the best results and significantly outperforms all other approaches, while its results also consistently improve with increased sample sizes. HMCMC results are not reported in these two figures because the needed number of model calls in order to get meaningful results is quite high for this high-dimensional

Table 7

Performance of various methods for the high dimensional problem with linear limit-state function in Example 6.

d	500 independent simulations		CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
			U(-1, 1)	N(0, 1)			
$\beta = 5$ $\sigma = 0.3$ $\tau = 0.7$	100	Number of total model calls	6,418	6,443	6,409	2,225	2,225
		C.o.V	0.62	0.69	0.45	0.12(0.12)	0.12(0.11)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 2.87E-7$)	2.97E-7	2.94E-7	2.86E-7	2.86E-7	2.87E-7
$\beta = 6$ $\sigma = 0.3$ $\tau = 0.7$	100	Number of total model calls	8,711	8,798	9,279	2,226	2,228
		C.o.V	0.62	0.95	0.58	0.14(0.13)	0.14(0.13)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 0.99E-9$)	1.05E-9	1.01E-9	1.03E-9	0.98E-9	0.99E-9
$\beta = 7$ $\sigma = 0.3$ $\tau = 0.7$	100	Number of total model calls	11,458	11,473	11,922	2,736	2,735
		C.o.V	0.89	1.94	0.77	0.17(0.16)	0.17(0.16)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.28E-12$)	1.36E-12	1.32E-12	1.31E-12	1.28E-12	1.28E-12
$\beta = 6$ $\sigma = 0.3$ $\tau = 0.7$	500	Number of total model calls	8,760	8,808	9,271	5,439	5,532
		C.o.V	0.67	1.05	0.60	0.25(0.19)	0.24(0.19)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 0.99E-9$)	1.01E-9	1.01E-9	1.04E-9	1.00E-9	0.99E-9
$\beta = 7$ $\sigma = 0.3$ $\tau = 0.7$	500	Number of total model calls	11,334	11,870	11,908	5,634	5,583
		C.o.V	0.92	2.25	0.82	0.27(0.25)	0.30(0.24)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.28E-12$)	1.45E-12	1.15E-12	1.39E-12	1.16E-12	1.16E-12

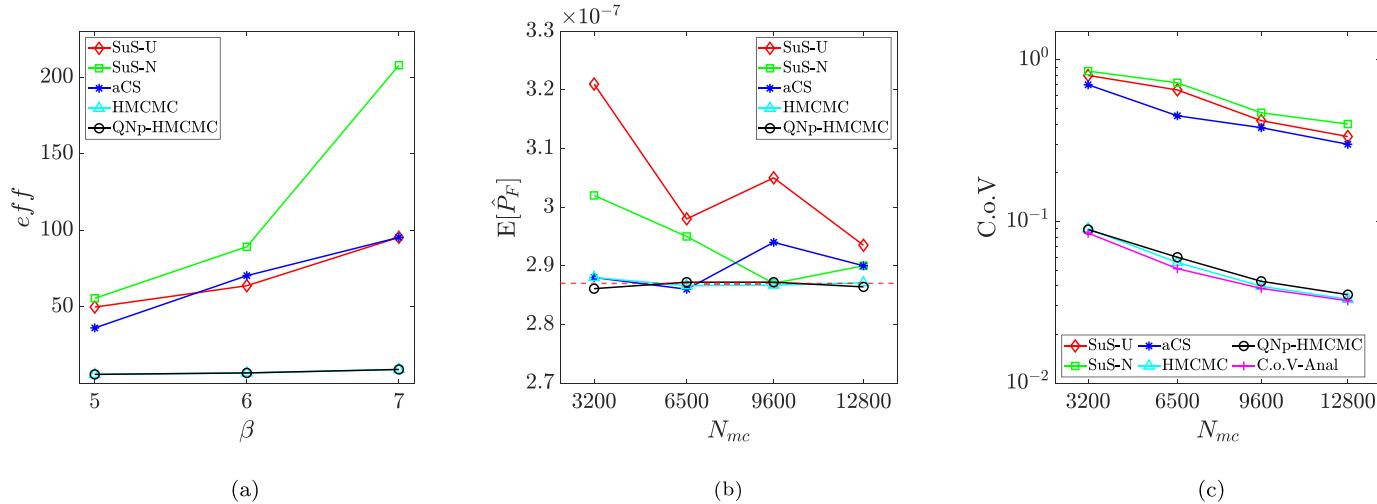


Fig. 10. Example 6: (a) eff versus different β values based on the results presented in Table 7 for $d = 100$; a low eff value indicates high efficiency. (b) Failure probability estimate against the number of model calls (N_{mc}) for $\beta = 5$ and $d = 100$ (dash line is the reference failure probability). (c) Coefficient of variation of the probability estimates plotted against the number of model calls for $\beta = 5$ and $d = 100$.

Table 8

Performance of various methods for the high dimensional problem with nonlinear limit-state function in Example 7.

d	500 independent simulations		CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
			U(-1, 1)	N(0, 1)			
$\lambda = 4.0$ $\gamma = 10$ $\sigma = 0.5$ $\tau = 0.7$	100	Number of total model calls	12,109	12,192	12,093	21,198	4,695
		C.o.V	1.43	1.71	2.18	0.16(0.15)	0.16(0.15)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.15E-6$)	1.22E-6	1.31E-6	1.23E-6	1.16E-6	1.16E-6
$\lambda = 3.0$ $\gamma = 50$ $\sigma = 0.5$ $\tau = 0.7$	100	Number of total model calls	13,758	14,200	14,442	53,664	5,924
		C.o.V	2.94	4.04	5.78	0.21(0.21)	0.21(0.21)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 5.63E-7$)	4.29E-7	4.70E-7	5.86E-7	5.69E-7	5.63E-7
$\lambda = 0.7$ $\gamma = 100$ $\sigma = 0.5$ $\tau = 0.7$	100	Number of total model calls	14,697	15,849	14,171	154,440	5,956
		C.o.V	7.76	5.19	5.63	0.26(0.26)	0.26(0.26)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 2.23E-6$)	2.40E-6	1.59E-6	2.83E-6	2.26E-6	2.24E-6
$\lambda = 2.5$ $\gamma = 100$ $\sigma = 0.6$ $\tau = 0.7$	200	Number of total model calls	11,785	12,098	13,872	78,414	6,510
		C.o.V	3.57	3.67	3.43	0.25(0.25)	0.25(0.24)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 5.06E-6$)	5.89E-6	5.33E-6	4.62E-6	5.08E-6	5.05E-6
$\lambda = 0.5$ $\gamma = 200$ $\sigma = 0.6$ $\tau = 0.7$	200	Number of total model calls	17,271	19,370	17,805	297,596	8,575
		C.o.V	8.09	11.85	6.64	0.31(0.31)	0.29(0.29)
		$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.19E-6$)	1.36E-6	1.40E-6	1.08E-6	1.18E-6	1.17E-6

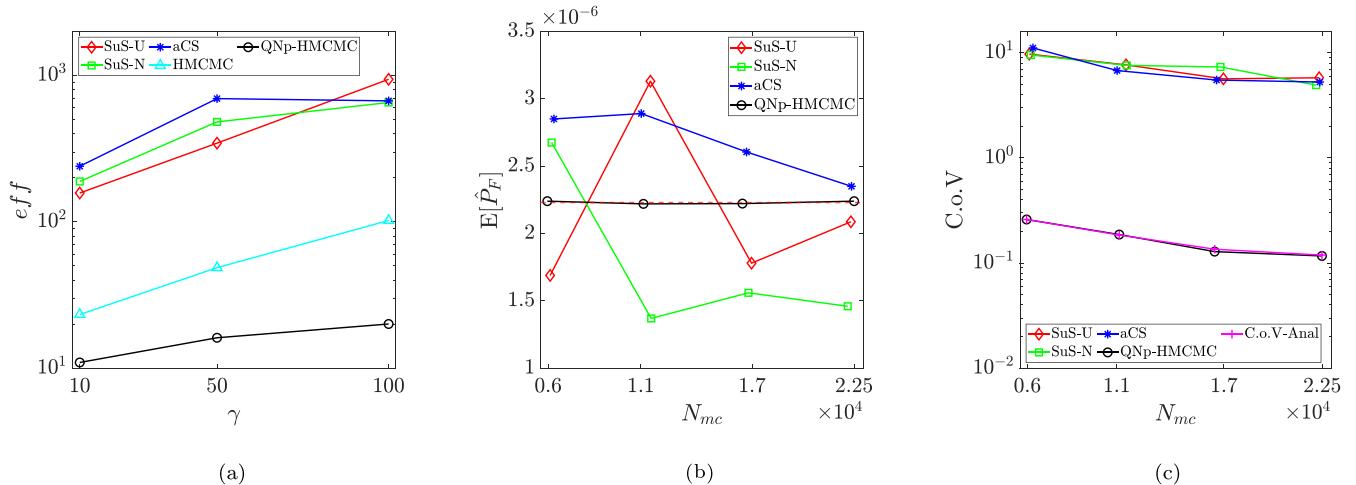


Fig. 11. Example 7: (a) eff versus different γ values based on the results presented in Table 8 for $d = 100$; a low eff value indicates high efficiency. (b) Failure probability estimate against the number of model calls (N_{mc}) for $d = 100$ and $\gamma = 100$ (the dash line is the reference failure probability). (c) Coefficient of variation estimates against the number of model calls for $d = 100$ and $\gamma = 100$.

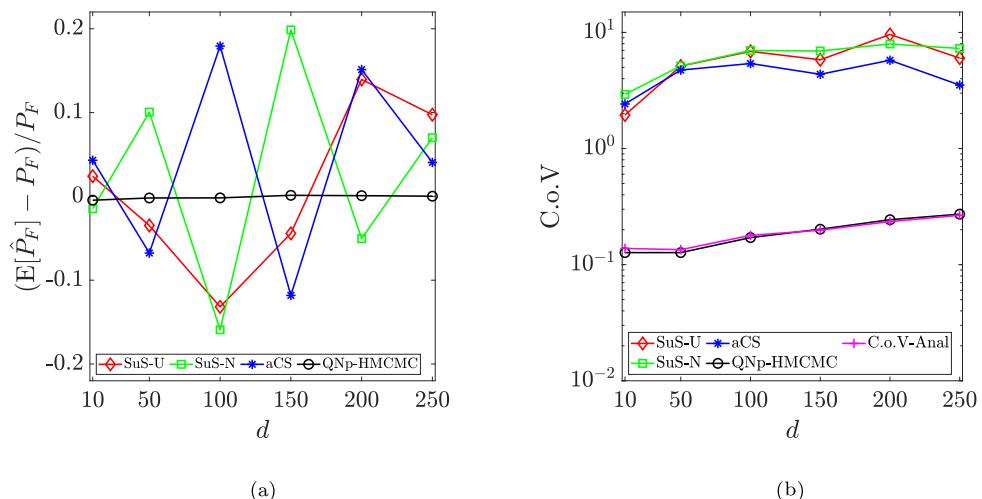


Fig. 12. (a) Relative bias of the $E[\hat{P}_F]$ estimator for problems with different number of dimensions d and the $\gamma = d$ cases in Example 7, based on the same number of model calls, approximately 11,000–12,000, for all methods. (b) Coefficient of variation of the estimates with respect to the number of dimensions d , for the $\gamma = d$ cases.

nonlinear problem, deeming this approach non-competitive in this case, and further noting the superior QNp-HMCMC performance and suitability in challenging high-dimensional spaces. The C.o.V-Anal curve in Fig. 11c again represents the QNp-HMCMC C.o.V estimation based on the analytical expression in Eq. (26), showcasing excellent agreement with numerical results.

Finally, Fig. 12 investigates the effect of the number of dimensions on the relative bias and C.o.V of the estimates, for the $\gamma = d$ cases. The threshold λ is adjusted to have a failure probability of around 10^{-6} . Again here, HMCMC results are not reported since they require a quite higher number of limit-state function evaluations, after $d \geq 50$ or so, than all other methods, to achieve meaningful and comparable results. The same number of model calls is used for all methods shown in Fig. 12, approximately 11,000–12,000 total model calls, which differs from Table 8 settings and results for the similar cases that are studied and reported there, and all results are based on 500 independent simulations. QNp-HMCMC results are shown to provide an essentially unbiased estimator and very low C.o.V. values for all analyzed cases, in contrast to SuS results, reporting very high C.o.V. values. The analytical C.o.V expression in Eq. (26) is also here in very good agreement with numerical results.

6.8. Example 8: High-dimensional highly nonlinear problem

To further investigate the QNp-HMCMC performance in challenging high-dimensional nonlinear problems, the limit-state function in this example is expressed in the standard normal space, as:

$$g(\theta) = Y_0 - \frac{1}{\sqrt{d}} \sum_{i=1}^d \theta_i + 2.5 \left(\theta_1 - \sum_{j=2}^{10} \theta_j \right)^2 + \left(\theta_{11} - \sum_{k=12}^{14} \theta_k \right)^4 + \left(\theta_{15} - \sum_{l=16}^{17} \theta_l \right)^8 \quad (35)$$

where d is the problem dimension, equal to 100. The parameters for the HMCMC approaches are chosen as $\sigma = 0.6$, $\tau = 0.7$ and 500 burn-in samples. Table 9 summarizes the computed results for various probability levels, acquired by adjusting the Y_0 threshold value, and as shown the QNp-HMCMC approach once more achieves excellent results, superior to all other methods. In Fig. 13a the eff metric is accordingly reported for the three different failure probability levels in Table 9, with QNp-HMCMC exhibiting the best and stable performance, with very significant efficiency gains in relation to other methods, and particularly for decreasing failure probability levels.

Table 9

Performance of various methods for the limit-state function of Eq. (35) in Example 8 ($d = 100$).

	500 Independent simulations	CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
		$U(-1, 1)$	$N(0, 1)$			
$Y_0 = 2.5$	Number of total model calls	9,380	9,593	10,593	28,026	7,295
$\sigma = 0.5$	C.o.V	0.86	1.13	0.86	0.22(0.21)	0.23(0.21)
$\tau = 0.7$	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 3.40E-5$)	3.38E-5	3.26E-5	3.62E-5	3.37E-5	3.41E-5
$Y_0 = 3.5$	Number of total model calls	12,948	13,704	13,965	32,009	7,924
$\sigma = 0.5$	C.o.V	2.60	3.27	2.25	0.23(0.22)	0.22(0.22)
$\tau = 0.7$	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 7.96E-7$)	8.26E-7	7.32E-7	7.99E-7	7.97E-7	7.96E-7
$Y_0 = 4.5$	Number of total model calls	17,304	17,793	17,974	31,948	7,889
$\sigma = 0.5$	C.o.V	6.33	5.85	6.91	0.26(0.25)	0.24(0.24)
$\tau = 0.7$	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 6.75E-9$)	8.84E-9	4.68E-9	7.32E-9	6.93E-9	6.86E-9

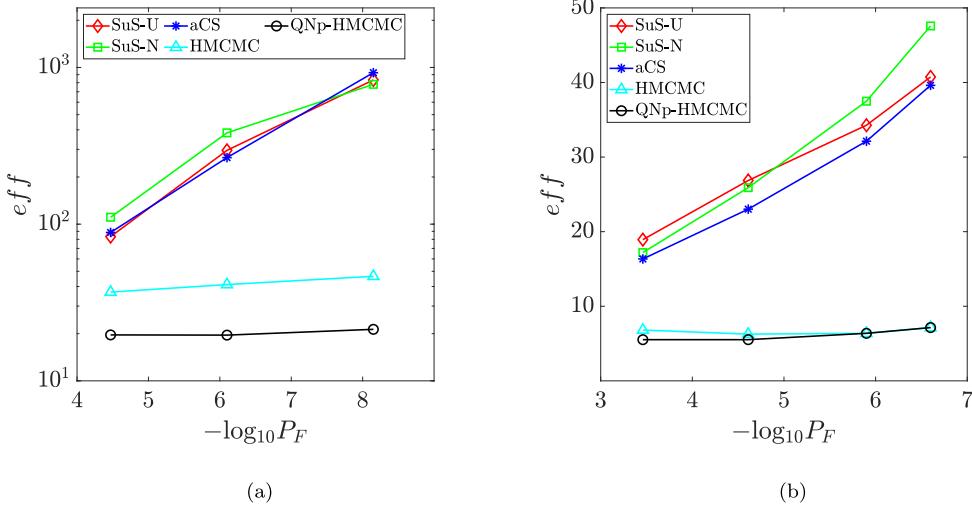
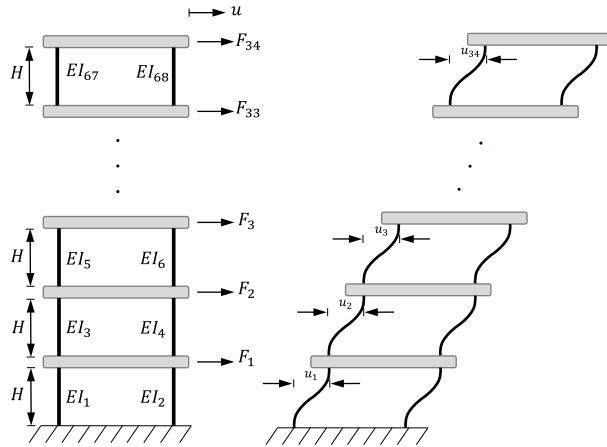
Fig. 13. Computed eff values for various failure probabilities, for (a) Example 8 with $d = 100$, and (b) Example 9 with $d = 102$. A low eff value exhibits high efficiency.

Fig. 14. A thirty four-story structure under static loads.

6.9. Example 9: A thirty four-story structural example

A thirty four-story structure is analyzed in this modified example from [79], as represented in Fig. 14, and is subjected to thirty four static loads $F_i, i = 1, 2, \dots, 34$. The floor slabs/beams are assumed to be rigid and all columns have identical length, $H = 4$ m, and different flexural stiffnesses $EI_k, k = 1, 2, \dots, 68$. Loads and stiffnesses are random variables and the total number of random variables is $d = 102$ in this problem. The loads are assumed normally distributed,

with a mean value of 2 kN and a C.o.V of 0.4, while stiffnesses are also normally distributed, with a mean value of 20 MNm² and a C.o.V of 0.2. Based on linear elastic behavior and excluding gravity effects, the top story displacement, u , can be calculated by adding the interstory relative displacements, u_i , as:

$$\begin{aligned}
 u_{34} &= \frac{F_{34}H^3}{12(EI_{67} + EI_{68})}, \quad u_{33} = \frac{(F_{33} + F_{34})H^3}{12(EI_{65} + EI_{66})}, \dots, \\
 u_2 &= \frac{(\sum_{i=2}^{34} F_i)H^3}{12(EI_3 + EI_4)}, \quad u_1 = \frac{(\sum_{i=1}^{34} F_i)H^3}{12(EI_1 + EI_2)} \\
 u &= \sum_{i=1}^{34} u_i, \quad g = Y_0 - u
 \end{aligned} \tag{36}$$

and g is the used limit-state function indicating failure when u exceeds a threshold value Y_0 , chosen as 0.21, 0.22, 0.23 and 0.235 m. All random variables are transformed into the standard normal space Θ for the analysis, and for the HMCMC-based methods the likelihood dispersion factor, σ , equals 0.3, and the burn-in samples are 400. In this example, $0 < g(\mathbf{0}) < 2$, and as discussed in Section 5, we perform the scaling with $g_c = g(\mathbf{0})/4$.

Table 10 summarizes all computed results, that once more validate the outstanding performance of the proposed HMCMC-based framework, particularly in high-dimensional, very low target probability problems. In Fig. 13b, these findings are further supported by the reported eff metric for the four considered failure probability levels in Table 10, showcasing again important advantages in relation to other methods.

Table 10Performance of various methods for the thirty four-story structure in Example 9 ($d = 102$).

	500 Independent simulations	CWMH-SuS		aCS-SuS	HMCMC	QNp-HMCMC
		$U(-1, 1)$	$N(0, 1)$			
$Y_0 = 0.21$	Number of total model calls	7,400	7,400	7,400	2,741	2,520
	C.o.V	0.22	0.20	0.19	0.13(0.12)	0.11(0.11)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 3.47E-4$)	3.56E-4	3.46E-4	3.39E-4	3.45E-4	3.47E-4
$Y_0 = 0.22$	Number of total model calls	9,200	9,200	9,200	2,723	2,519
	C.o.V	0.28	0.27	0.24	0.12(0.11)	0.11(0.10)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 2.48E-5$)	2.52E-5	2.54E-5	2.42E-5	2.48E-5	2.48E-5
$Y_0 = 0.23$	Number of total model calls	11,468	11,475	11,470	2,819	2,819
	C.o.V	0.32	0.35	0.30	0.12(0.11)	0.12(0.10)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 1.26E-6$)	1.27E-6	1.30E-6	1.22E-6	1.26E-6	1.26E-6
$Y_0 = 0.235$	Number of total model calls	12,804	12,836	12,815	3,019	3,019
	C.o.V	0.36	0.42	0.35	0.13(0.11)	0.13(0.11)
	$\mathbb{E}[\hat{P}_F]$ (Ref. $P_F \sim 2.56E-7$)	2.62E-7	2.53E-7	2.41E-7	2.52E-7	2.50E-7

7. Conclusions

A novel approach for estimation of rare event and failure probabilities, termed *Approximate Sampling Target with Post-processing Adjustment* (ASTPA), is presented in this paper, suitable for low- and high-dimensional problems, very small probabilities, and multiple failure modes. ASTPA can provide an accurate, unbiased probability estimation with an efficient number of limit-state function evaluations. The basic idea of ASTPA is to construct a relevant target distribution by weighting the high-dimensional random variable space through a likelihood model, using the limit-state function. Although this framework is general, to sample from this target distribution we utilize gradient-based Hamiltonian MCMC schemes in this work, including our newly developed *Quasi-Newton based mass preconditioned HMCMC algorithm* (QNp-HMCMC) that can sample very adeptly, particularly in difficult cases with high-dimensionality and very small rare event probabilities. Finally, an original post-sampling step is also devised, using the introduced *inverse importance sampling* procedure, based on the samples. As shown and discussed in the paper, a very approximately fitted Gaussian Mixture Model (GMM), even with one component and diagonal covariance matrix, is adequate at this step, thus avoiding the known scalability issues of GMMs, as also showcased by the high-dimensional numerical examples, with dimensions up to 500 presented in this work. Overall, the performance of the proposed methodology is carefully examined and compared very successfully against Subset Simulation in a series of low- and high-dimensional problems. As a general guideline, QNp-HMCMC is recommended to be used for problems with more than 20 dimensions, where traditional HMCMC schemes may not perform that well. However, even in lower dimensions QNp-HMCMC performs extremely well and is still a very competitive algorithm. Since we are utilizing gradient-based sampling methods in this work, all of our analyses and results are based on the fact that analytical gradients can be computed. In cases where numerical schemes are needed for the gradient evaluations, then HMCMC methods will not be competitive in relation to Subset Simulation. It should also be pointed out that different feature combinations of the HMCMC and QNp-HMCMC algorithms can be possible, based on problem-specific characteristics. Some of the ongoing and future works are directed toward exploring various ASTPA variants, including non-gradient based sampling approaches, sampling directly from non-Gaussian spaces [80], without the need for Gaussian transformations, and estimating high-dimensional first-passage problems under various settings [81].

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

This material is based upon work partially supported by the U.S. National Science Foundation under CAREER Grant No. 1751941. The authors would also like to thank Prof. Dr. Daniel Straub and Dr. Iason Papaioannou at the Technical University of Munich, for very fruitful scientific discussions in relation to the presented formulation, and particularly for their helpful insights on the construction of the target distribution.

References

- [1] R.E. Melchers, A.T. Beck, *Structural Reliability Analysis and Prediction*, Wiley, 2018.
- [2] O. Ditlevsen, H.O. Madsen, *Structural Reliability Methods*, Department of Mechanical Engineering, Technical University of Denmark, 2007, Available at <http://od-website.dk/books/OD-HOM-StrucRelMeth-Ed2.3.7.pdf>.
- [3] S.-K. Au, Y. Wang, *Engineering Risk Assessment with Subset Simulation*, Wiley, 2014.
- [4] E. Nikolaidis, D. Ghiocel, S. Singhal, *Engineering Design Reliability Handbook*, CRC Press, 2005.
- [5] M. Lemaire, A. Chateauneuf, J. Mitteau, *Structural Reliability*, Wiley, 2009.
- [6] R. Rackwitz, *Reliability analysis—A review and some perspectives*, Struct. Saf. 23 (4) (2001) 365–395.
- [7] K. Breitung, *40 Years FORM: Some new aspects?* Probab. Eng. Mech. 42 (2015) 71–77.
- [8] A. Der Kiureghian, P.-L. Liu, *Structural reliability under incomplete probability information*, J. Eng. Mech. 112 (1) (1986) 85–104.
- [9] O. Ditlevsen, P. Bjerager, *Methods of structural systems reliability*, Struct. Saf. 3 (3–4) (1986) 195–229.
- [10] M. Shinozuka, *Basic analysis of structural safety*, J. Struct. Eng. 109 (3) (1983) 721–740.
- [11] K. Breitung, *Asymptotic approximations for multinormal integrals*, J. Eng. Mech. 110 (3) (1984) 357–366.
- [12] P.-L. Liu, A. Der Kiureghian, *Optimization algorithms for structural reliability*, Struct. Saf. 9 (3) (1991) 161–177.
- [13] T. Haukaas, A. Der Kiureghian, *Strategies for finding the design point in nonlinear finite element reliability analysis*, Probab. Eng. Mech. 21 (2) (2006) 133–147.
- [14] M. Valdebenito, H. Pradlwarter, G. Schuëller, *The role of the design point for calculating failure probabilities in view of dimensionality and structural nonlinearities*, Struct. Saf. 32 (2) (2010) 101–111.
- [15] G. Schuëller, H. Pradlwarter, *Benchmark study on reliability estimation in higher dimensions of structural systems—An overview*, Struct. Saf. 29 (3) (2007) 167–182.
- [16] I. Papaioannou, W. Betz, K. Zwirglmaier, D. Straub, *MCMC algorithms for subset simulation*, Probab. Eng. Mech. 41 (2015) 89–103.
- [17] K.M. Zuev, L.S. Katafygiotis, *Modified Metropolis-Hastings algorithm with delayed rejection*, Probab. Eng. Mech. 26 (3) (2011) 405–412.

[18] K.M. Zuev, J.L. Beck, S.-K. Au, L.S. Katafygiotis, Bayesian post-processor and other enhancements of Subset Simulation for estimating failure probabilities in high dimensions, *Comput. Struct.* 92 (2012) 283–296.

[19] S.-K. Au, E. Patelli, Rare event simulation in finite-infinite dimensional space, *Reliab. Eng. Syst. Saf.* 148 (2016) 67–77.

[20] Z. Wang, M. Broccardo, J. Song, Hamiltonian Monte Carlo methods for subset simulation in reliability analysis, *Struct. Saf.* 76 (2019) 51–67.

[21] S.-K. Au, J.L. Beck, Estimation of small failure probabilities in high dimensions by subset simulation, *Probab. Eng. Mech.* 16 (4) (2001) 263–277.

[22] S.-K. Au, J.L. Beck, A new adaptive importance sampling scheme for reliability calculations, *Struct. Saf.* 21 (2) (1999) 135–158.

[23] S.-K. Au, J.L. Beck, Important sampling in high dimensions, *Struct. Saf.* 25 (2) (2003) 139–163.

[24] I. Papaioannou, C. Papadimitriou, D. Straub, Sequential importance sampling for structural reliability analysis, *Struct. Saf.* 62 (2016) 66–75.

[25] R.Y. Rubinstein, D.P. Kroese, *The Cross-Entropy Method: A Unified Approach to Monte Carlo Simulation, Randomized Optimization and Machine Learning*, Information Science & Statistics, Springer, 2004.

[26] N. Kurtz, J. Song, Cross-entropy-based adaptive importance sampling using Gaussian mixture, *Struct. Saf.* 42 (2013) 35–44.

[27] S. Geyer, I. Papaioannou, D. Straub, Cross entropy-based importance sampling using Gaussian densities revisited, *Struct. Saf.* 76 (2019) 15–27.

[28] D.Y. Yang, J. Teng, D.M. Frangopol, Cross-entropy-based adaptive importance sampling for time-dependent reliability analysis of deteriorating structures, *Struct. Saf.* 66 (2017) 38–50.

[29] I. Papaioannou, S. Geyer, D. Straub, Improved cross entropy-based importance sampling with a flexible mixture model, *Reliab. Eng. Syst. Saf.* 191 (2019) 106564.

[30] P. Koutsourelakis, H. Pradlwarter, G. Schueller, Reliability of structures in high dimensions, Part I: Algorithms and applications, *Probab. Eng. Mech.* 19 (4) (2004) 409–417.

[31] H. Pradlwarter, G. Schueller, P. Koutsourelakis, D. Charmpis, Application of line sampling simulation method to reliability benchmark problems, *Struct. Saf.* 29 (3) (2007) 208–221.

[32] P. Bjerager, Probability integration by directional simulation, *J. Eng. Mech.* 114 (8) (1988) 1285–1302.

[33] J. Nie, B.R. Ellingwood, Directional methods for structural reliability analysis, *Struct. Saf.* 22 (3) (2000) 233–249.

[34] C. Andrieu-Renaud, B. Sudret, M. Lemaire, The PHI2 method: A way to compute time-variant reliability, *Reliab. Eng. Syst. Saf.* 84 (1) (2004) 75–86.

[35] C. Bucher, Asymptotic sampling for high-dimensional reliability analysis, *Probab. Eng. Mech.* 24 (4) (2009) 504–510.

[36] A. Naess, B. Leira, O. Batsevych, System reliability analysis by enhanced Monte Carlo simulation, *Struct. Saf.* 31 (5) (2009) 349–355.

[37] S. Tong, G. Stadler, Large deviation theory-based adaptive importance sampling for rare events in high dimensions, *SIAM/ASA J. Uncertain. Quantif.* 11 (3) (2023) 788–813.

[38] A.M. Johansen, P. Del Moral, A. Doucet, Sequential Monte Carlo samplers for rare events, in: 6th International Workshop on Rare Event Simulation, 2006, pp. 256–267, <https://warwick.ac.uk/fac/sci/statistics/staff/academic-research/johansen/publications/JDD06c.pdf>.

[39] F. Cérou, P. Del Moral, T. Furon, A. Guyader, Sequential Monte Carlo for rare event estimation, *Stat. Comput.* 22 (3) (2012) 795–808.

[40] P. Del Moral, A. Doucet, A. Jasra, Sequential Monte Carlo samplers, *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)* 68 (3) (2006) 411–436.

[41] P. Del Moral, J. Garnier, Genealogical particle analysis of rare events, *Ann. Appl. Probab.* 15 (4) (2005) 2496–2534.

[42] F. Cérou, P. Del Moral, F. Le Gland, P. Lezaud, Genetic genealogical models in rare event analysis, *Alea* 1 (2006) 181–203.

[43] A. Guyader, N. Hengartner, E. Matzner-Löber, Simulation and estimation of extreme quantiles and extreme probabilities, *Appl. Math. Optim.* 64 (2) (2011) 171–196.

[44] C. Walter, Moving particles: A parallel optimal multilevel splitting method with application in quantiles estimation and meta-model based algorithms, *Struct. Saf.* 55 (2015) 10–25.

[45] R.J. Allen, C. Valeriani, P.R. ten Wolde, Forward flux sampling for rare event simulations, *J. Phys.: Condens. Matter* 21 (46) (2009) 463102.

[46] S. Brooks, A. Gelman, G. Jones, X.-L. Meng, *Handbook of Markov Chain Monte Carlo*, CRC Press, 2011.

[47] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* 21 (6) (1953) 1087–1092.

[48] S. Geman, D. Geman, Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intell.* (6) (1984) 721–741.

[49] S. Duane, A.D. Kennedy, B.J. Pendleton, D. Roweth, Hybrid Monte Carlo, *Phys. Lett. B* 195 (2) (1987) 216–222.

[50] R.M. Neal, *Bayesian Learning for Neural Networks*, Springer, 1995.

[51] R.M. Neal, MCMC using Hamiltonian dynamics, 2012, arXiv preprint [arXiv:1206.1901](https://arxiv.org/abs/1206.1901).

[52] A. Beskos, N. Pillai, G. Roberts, J.-M. Sanz-Serna, A. Stuart, Optimal tuning of the hybrid Monte Carlo algorithm, *Bernoulli* 19 (5A) (2013) 1501–1534.

[53] M. Betancourt, A conceptual introduction to Hamiltonian Monte Carlo, 2017, arXiv preprint [arXiv:1701.02434](https://arxiv.org/abs/1701.02434).

[54] B. Carpenter, A. Gelman, M.D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, A. Riddell, Stan: A probabilistic programming language, *J. Statistical Software* 76 (1) (2017).

[55] G.O. Roberts, J.S. Rosenthal, Optimal scaling for various Metropolis-Hastings algorithms, *Statistical Science* 16 (4) (2001) 351–367.

[56] G.O. Roberts, J.S. Rosenthal, Optimal scaling of discrete approximations to Langevin diffusions, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 60 (1) (1998) 255–268.

[57] M. Girolami, B. Calderhead, Riemann Manifold Langevin and Hamiltonian Monte Carlo methods, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 73 (2) (2011) 123–214.

[58] W. Chen, Z. Wang, M. Broccardo, J. Song, Riemannian Manifold Hamiltonian Monte Carlo based subset simulation for reliability analysis in non-Gaussian space, *Struct. Saf.* 94 (2022) 102134.

[59] A. Griewank, On automatic differentiation, *Mathematical Programming: Recent Developments and Applications* 6 (6) (1989) 83–107.

[60] T. Chen, E. Fox, C. Guestrin, Stochastic gradient Hamiltonian Monte Carlo, in: International Conference on Machine Learning, 2014, pp. 1683–1691, <https://proceedings.mlr.press/v32/chen14.html>.

[61] M.D. Hoffman, A. Gelman, The No-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.* 15 (1) (2014) 1593–1623.

[62] M. Hohenbichler, R. Rackwitz, Non-normal dependent vectors in structural safety, *J. Eng. Mech. Div.* 107 (6) (1981) 1227–1238.

[63] R. Lebrun, A. Dutfoy, An innovating analysis of the Nataf transformation from the copula viewpoint, *Probab. Eng. Mech.* 24 (3) (2009) 312–320.

[64] H. Nikbakht, K.G. Papakonstantinou, A direct Hamiltonian MCMC approach for reliability estimation, in: 3rd International Conference on Uncertainty Quantification in Computational Sciences and Engineering, 2019, <https://arxiv.org/abs/1909.03575>.

[65] B.J. Alder, T.E. Wainwright, Studies in molecular dynamics. I. General method, *J. Chem. Phys.* 31 (2) (1959) 459–466.

[66] Z. Wang, S. Mohamed, N. Freitas, Adaptive Hamiltonian and Riemann manifold Monte Carlo, in: International Conference on Machine Learning, 2013, pp. 1462–1470, <https://proceedings.mlr.press/v28/wang13e.html>.

[67] C.P. Andriots, K.G. Papakonstantinou, V.K. Koumousis, Nonlinear programming hybrid beam-column element formulation for large-displacement elastic and inelastic analysis, *J. Eng. Mech.* 144 (10) (2018) 04018096.

[68] J. Nocedal, S.J. Wright, *Numerical Optimization*, Springer, 2006.

[69] Y. Zhang, C.A. Sutton, Quasi-Newton methods for Markov Chain Monte Carlo, in: Advances in Neural Information Processing Systems, 2011, pp. 2393–2401, <https://proceedings.neurips.cc/paper/2011/hash/e702e51da2c0f5be4dd354bb3e295d37-Abstract.html>.

[70] T. Fu, L. Luo, Z. Zhang, Quasi-Newton Hamiltonian Monte Carlo, in: Uncertainty in Artificial Intelligence, 2016.

[71] S. Lacaze, L. Breault, S. Missoum, M. Balesdent, Probability of failure sensitivity with respect to decision variables, *Struct. Multidiscip. Optim.* 52 (2015) 375–381.

[72] A.A. Taflanidis, J.L. Beck, Reliability-based design using two-stage stochastic optimization with a treatment of model prediction errors, *J. Eng. Mech.* 136 (12) (2010) 1460–1473.

[73] G. McLachlan, D. Peel, *Finite Mixture Models*, Wiley, 2000.

[74] Engineering Risk Analysis group, Subset Simulation (SuS) with the adaptive conditional sampling (aCS) approach, Matlab code, Technical University of Munich. Available at <https://www.cee.ed.tum.de/era/software/reliability/subset-simulation/>.

[75] A. Der Kiureghian, T. Dakessian, Multiple design points in first and second-order reliability, *Struct. Saf.* 20 (1) (1998) 37–49.

[76] D.M. Himmelblau, *Applied Nonlinear Programming*, McGraw-Hill Companies, 1972.

[77] X. Du, First order and second reliability methods, University of Missouri – Rolla, 2005, pp. 1–33.

[78] S.-K. Au, J. Ching, J.L. Beck, Application of Subset Simulation methods to reliability benchmark problems, *Struct. Saf.* 29 (3) (2007) 183–193.

[79] C. Bucher, *Computational Analysis of Randomness in Structural Mechanics: Structures and Infrastructures Book Series*, vol. 3, CRC Press, 2009.

[80] K.G. Papakonstantinou, H. Nikbakht, E. Eshra, Quasi-Newton Hamiltonian MCMC sampling for reliability estimation in high-dimensional non-Gaussian spaces, in: 13th International Conference on Structural Safety & Reliability, ICOSSAR, Shanghai, China, 2022, https://www.researchgate.net/publication/365655375_Quasi-Newton_Hamiltonian_MCMC_Sampling_for_Reliability_Estimation_in_High-dimensional_Non-Gaussian_Spaces.

[81] K.G. Papakonstantinou, E. Eshra, H. Nikbakht, Hamiltonian MCMC based framework for time-variant rare event uncertainty quantification, in: 14th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP, Dublin, Ireland, 2023, <http://www.tara.tcd.ie/handle/2262/103409>.