Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio
 Published by Research Publishing, Singapore.
doi: 10.3850/978-981-18-8071-1_P675-cd

Failure domain analysis using Sliced-Normal distributions

James Hammond
UQlab, Dept of Aeronautics, Imperial College London, UK. E-mail: j.hammond18@imperial.ac.uk

Luis G. Crespo
NASA Langley Research Center, Hampton, VA, 23681, USA. E-mail: luis.g.crespo@nasa.gov

Francesco Montomoli
UQlab, Dept of Aeronautics, Imperial College London, UK.

Sliced-normal (SN) distributions enable characterization of parameters exhibiting complex dependencies with minimal modeling effort. We leverage the semialgebraic nature of SN distributions to identify the most likely points of failure (MLPs) corresponding to a given failure domain. When this domain is semialgebraic, Sum of Squares (SOS) optimization is used to guarantee that no MLPs are missed within a region of interest. The MLPs not only enable the identification of all the critical points of failure, but also the efficient estimation of failure probabilities using Importance Sampling (IS). The IS density is constructed as a Gaussian Mixture (GM) model with means at the MLPs and covariances equal to the weighted empirical covariance of sample sets drawn in the vicinity of the MLPs.

Keywords: Importance sampling, Reliability analysis, Failure estimation, Parameter dependency, Semidefinite programming, Gaussian mixture.

1. Introduction

The main aim of reliability analysis is to accurately and efficiently quantify the probability of failure of any given system. Consider $\delta \in \mathcal{R}^{n_\delta}$ to be the n_δ uncertain parameters of the system (where $\mathcal{R}$ is the set of real numbers) and $f_\delta(\delta;\theta) : \mathcal{R}^{n_\delta} \to \mathcal{R}^+$ to be a joint density function with shaping hyperparameter θ, having support set $\Delta \subset \mathcal{R}^{n_\delta}$. Consider also n_g individual limit state functions (LSF) $g_j(\delta) = 0 : \mathcal{R}^{n_\delta} \to \mathcal{R}$ for $j = [1, n_g]$. The failure domain associated with an individual LSF is defined as $\mathcal{F}^j = \{\delta \in \Delta : g_j(\delta) \geq 0\}$. The full failure domain is defined as the union of the n_g failure regions

$$\mathcal{F} = \bigcup_{j=1}^{n_g} \mathcal{F}^j. \tag{1}$$

From this, the probability of failure can be calculated as the integral over $\mathcal{F}$ of the joint density function as

$$p_f(\theta) = \int_{\mathcal{F}} f_\delta(\delta;\theta) d\delta. \tag{2}$$

The complement of the failure domain will be called the success domain.

Systems in engineering applications often have a large number of uncertain parameters and comparatively small failure probabilities. Further, the probability of failure in Eq. (2) is a multidimensional integral that cannot be evaluated analytically in general. Therefore, Monte Carlo methods, Bayesian quadrature, or sparse grids approximations are required to approximate it, a task that is often computationally expensive. For instance, accurate estimation of a failure probability of 1×10^{-6} using Monte Carlo sampling requires $N = 1 \times 10^8$ function evaluations (Papaioannou et al., 2016), whose cost might be unacceptably large when $g_j(\delta)$ are computed from numerical simulations. A variety of approaches for the efficient approximation of p_f have been developed using asymptotic approximations such as FORM and SORM (Der Kiureghian, 2005). They yield reasonable results when (i) there is a single LSF, and the corresponding failure domain is closely approximated by a half-space (Hohenbichler et al., 1987) or a quadratic set, and (ii) when there is a single LSF and the failure probability is small. Both of these settings depend heavily on the iden-

tification of the true Most Probable Point (MPP) of failure. This later requirement is far from certain when the LSFs are complex and the uncertain space is high-dimensional. Furthermore, calculation of the MPP requires mapping the joint density to standard normal space. This is not possible when the uncertain parameters exhibit nonlinear parameter dependencies. In contrast, sampling techniques employing variants of the Monte Carlo (MC) simulation relax these requirements at the cost of higher computational expense. To alleviate this expense, importance sampling techniques have been developed which provide a low variance estimate of the failure probability with a reduced number of samples (Tabandeh et al., 2022).

In this work we assume that $f_\delta(\delta;\theta)$ takes the form of a *sliced* distribution. *Sliced* distributions allow characterization of multivariate data exhibiting complex parameter dependencies with minimal modelling effort (Crespo et al., 2019) and have been shown to be versatile and efficient when the number of uncertain parameters is moderate (Crespo et al., 2021). Further, we leverage the semialgebraic nature of *sliced* distributions to define a set of most likely points of failure (MLPs) in a manner analogous to the MPP in FORM/SORM, and construct an importance sampling distribution about these points to efficiently estimate p_f. However, these MLPs, which can be found completely and correctly using semi-infinite programming, are applicable to complex distributions for which no transformation to standard normal space exists.

This paper is organized as follows. In Section 2, *sliced* distributions are introduced. Section 3 introduces the algorithm to search for the complete set of MLPs in the failure domain and Section 4 presents an importance sampling scheme for the accurate estimation of p_f that leverages the MLP set. Finally, Section 5 applies the proposed framework to several examples of increasing level of complexity and realism.

2. *Sliced-Normal* distributions

The developments that follow are based on the *Sliced-Normal* (SN) class of parameter models (Crespo et al., 2019, 2021). SNs have been shown to accurately characterize complex parameter dependencies in moderate dimensions. A brief summary of their underlying structure is presented next. Denote as S^n_{++} the space of symmetric positive definite matrices in $\mathcal{R}^{n\times n}$. Consider the polynomial mapping from physical space $\delta \in \mathcal{R}^{n_\delta}$ to feature space $z \in \mathcal{R}^{n_z}$ given by the function $z = t(\delta;d)$. Here $t(\delta;d) : \mathcal{R}^{n_\delta} \to \mathcal{R}^{n_z}$ with $n_z = \binom{n_p+d}{n_p} - 1$, is the vector of monomials in δ of degree less than or equal to d. The joint density of an SN is

$$f_\delta(\delta;\theta) = \begin{cases} \frac{1}{c(\theta)} \exp\left(-\frac{\phi(z;\theta)}{2}\right) & \text{if } \delta \in \Delta, \\ 0 & \text{otherwise,} \end{cases} \tag{3}$$

where the hyperparameter $\theta = \{\mu, P\}$ comprises the mean $\mu \in \mathcal{R}^{n_z}$ and the precision matrix $P \in S^{n_z}_{++}$,

$$\phi(z;\mu,P) = (z-\mu)^\top P (z-\mu) \tag{4}$$

is a Sum of Squares (SOS) of polynomials in δ of degree $2d$, $c(\theta)$ is the normalization constant, and Δ is the support set.

The maximum likelihood (ML) estimate of an SN entails solving a non-convex optimization program in which μ and P are the decision variables. The *Primal* SN subclass proposed in Crespo et al. (2021) leads to a convex optimization program thereby facilitating its usage in higher dimensions. This subclass is given by

$$\phi(z;\lambda) = (z-\mu^\star)^\top C \,\mathrm{diag}(\lambda)\, C^\top (z-\mu^\star), \tag{5}$$

where $\theta = \lambda \in \mathcal{R}^{n_z}$ with $\lambda \geq 0$ is the hyperparameter of the distribution, $CC^\top$ is a Cholesky decomposition of $P^\star$,

$$\mu^\star = \frac{1}{q}\sum_{i=1}^{q} z^{(i)}, \tag{6}$$

$$P^\star = \left(\frac{1}{q}\sum_{i=1}^{q}\left(z^{(i)}-\mu^\star\right)\left(z^{(i)}-\mu^\star\right)^\top\right)^{-1}, \tag{7}$$

$z^{(i)} = t\left(\delta^{(i)};d\right)$, and $\mathcal{D} = \{\delta^{(i)}\}_{i=1}^q$ is the training data set.

The normalization constant $c(\theta)$ in Eq. (3) can be approximated by

$$c(m,\theta) = \frac{\Delta_V}{m}\sum_{i=1}^{m}\exp\left(-\frac{\phi\left(t(u^{(i)};d);\theta\right)}{2}\right), \tag{8}$$

where $\left\{u^{(i)}\right\}_{i=1}^{m}$ are samples uniformly distributed over Δ, and Δ_V is the volume of Δ.

3. Search for the most likely points of failure (MLPs)

Failure probabilities in reliability analysis are often small, and their accurate estimation is essential. In this article an importance sampling (IS) technique is developed for this purpose. To this end, we first introduce a method to identify the MLPs. These points represent dominant failure modes of the system.

In this section we assume that we have a system with uncertain parameters characterized by a SN with joint density $f_\delta(\delta) : \mathcal{R}^{n_\delta} \to \mathcal{R}^+$, and a set of LSFs $g_j(\delta) : \mathcal{R}^{n_\delta} \to \mathcal{R}$ where the failure domain is defined as in Eq. (1). In this setting we want to identify the elements of the failure domain where the likelihood function has local maxima. As such we are not only interested in the global maximum, which corresponds to the dominant failure mode, but to other failure modes as well.

Key steps of the algorithm used to identify the MLPs, detailed in the Algorithm 1, are explained next:

(1) Find the element of the failure domain where the likelihood is maximal. Call this point MLP_1.
(2) Add a constraint that makes the vicinity of MLP_1 be in the success domain. We assume that this region is a hypersphere centered at MLP_1 of radius r, where r is small.
(3) Solve the optimization program with the sphere exclusion constraint. If the optimum $\delta^\star$, lies on the surface of the sphere, increase r and repeat until the optimum no longer lies on the sphere's surface. This optimum is the next distinct MLP, called MLP_2.
(4) Repeat the process above until the likelihood at the newly discovered MLP is less than $1\,\%$ that of the first MLP. Every time the solution to the optimization program falls on the surface of any of the existing spheres, the corresponding radius of this sphere is increased.

Effectively, once the first MLP has been found, it is contained by a hypersphere whose radius grows sequentially until the optimum jumps. The cycle repeats with each of the spheres being iteratively grown out until no new MLPs are found. The initial radius and the growth rate of the spheres are hyperparameters to be optimized and are set heuristically. Over several test cases, an initial radius corresponding to $5\,\%$ of the support width and a growth rate of $\lambda = 1.25$ reliably obtained the full set of MLPs. Setting either value too high can lead to missing subsequent MLPs in the immediate vicinity of a previously identified MLP. Setting the values too low, whilst guaranteeing identification of all MLPs, leads to a computationally expensive search.

To search for an MLP within a region defined by a set of nonlinear constraints $c(\delta) \leq 0$, the

Algorithm 1 The proposed algorithm to search for the most likely points of failure (MLPs).

Require: $z = t(\delta, d)$
Require: $\phi(z; \mu, P)$ with $P \succ 0$
Require: $g_i(\delta)$ for $i \in [1, n_g]$
Require: $\text{solve}(\phi(z), c(\delta))$
$r, \epsilon, \phi^\star, \delta^\star \leftarrow []$
$\lambda \leftarrow \lambda_0$
$\text{tol} \leftarrow \text{tol}_0$
$k \leftarrow 1$
$c(\delta) = \max(g_i), \forall i \in [1, n_g]$
$\phi_k^\star, \delta_k^\star, \text{feas} \leftarrow \text{solve}(\phi(t(\delta, d)), c(\delta))$
while feas **do**
    $r_k \leftarrow r_0$
    $\epsilon_k \leftarrow 0$
    while $\text{any}(\epsilon) < \text{tol}$ and feas **do**
        $j \leftarrow \text{find}(\min(\epsilon))$
        $r_j \leftarrow \lambda r_j$
        $g_j^{\text{art}}(\delta, \delta_j^\star, r_j) = \|\delta - \delta_j^\star\| - r_j$
        $c \leftarrow [c, g_l^{\text{art}}], \forall i, l \in [1, n_g], [1, k]$
        $\phi_{\text{cnd}}^\star, \delta_{\text{cnd}}^\star, \text{feas} \leftarrow \text{solve}(\phi(z), c(\delta))$
        $\epsilon \leftarrow g_l^{\text{art}}(\delta_{\text{cnd}}^\star, \delta_l^\star, r_l), \forall l \in [1, k]$
    end while
    if feas **then**
        $\phi_{k+1}^\star, \delta_{k+1}^\star \leftarrow \phi_{\text{cnd}}^\star, \delta_{\text{cnd}}^\star$
        $k \leftarrow k + 1$
    end if
end while

following optimization program is used

$$\delta^\star = \underset{\delta\in\Delta}{\operatorname{argmin}} \ \{\phi(t(\delta,d);\mu,P) : c(\delta) \geq 0\}\,, \quad (9)$$

where we recall the definition of $\phi(t(\delta,d);\mu,P)$ given by Eq. (4). This program is equivalent to maximizing the likelihood within the failure domain. In Algorithm 1 this optimization program is denoted by $\text{solve}(\phi(z), c(\delta))$. For the first call of solve (line 6), the region defined by $c(\delta) \geq 0$ (line 5) corresponds to $\max_j g_j(\delta) \geq 0$. In line 14 we define the new region where the next candidate MLP should be searched for as the subset of the failure region outside the artificial spheres centered at the MLPs of radii r_k. As the algorithm progresses, all the points within spheres containing previously found MLPs become infeasible. The program in Eq. (9) can be solved using standard gradient-based optimization algorithms. Note however that such algorithms might fail to converge to the global optimal, thereby rendering underestimations of the failure probability.

However, when the limit state functions are polynomial, the optimization program in Eq. (9) can be solved using semidefinite programming thereby guaranteeing that no local maxima are missed, and therefore that the computed reliability is accurate. The details of this formulation are presented next.

3.1. *Global optimality using semidefinite programming (SDP)*

A general polynomial optimization problem can be stated as

$$\begin{aligned} &\min p(x), \text{ subject to } x \in K \\ &\text{with } K := \{x \in \mathcal{R}^n \mid q_i(x) \geq 0, r_j(x) = 0\}\,, \end{aligned} \quad (10)$$

where p, q_i and r_j are multivariate polynomials, noting that the set K is basic semialgebraic. If we could optimize over the set of polynomials taking non-negative (positive semidefinite) values on a given basic semialgebraic set, then Eq. (10) could be solved globally (Ahmadi, 2018). This can be seen by realizing that the optimal value of Eq. (10) is equivalent to the optimal value of the program

$$\begin{aligned} &\max \gamma, \\ &\text{subject to } p(x) - \gamma \geq 0, \forall x \in K. \end{aligned} \quad (11)$$

The global solution to this program can be tightly bounded using Sum of Squares (SOS) optimization (Packard et al., 2010). If the set of nonlinear constraints $c(\delta) \geq 0$ in Eq. (9) is polynomial (or can be closely approximated by polynomials), the program in Eq. (9) becomes a particular instance of Eq. (10). $\phi(t(\delta,d))$ is the polynomial $p(x)$, which we would like to minimize subject to a set of constraints. Unfortunately, the union of the individual failure domains prescribed by $c(\delta) = \max_j g_j(\delta)$ is not amenable to SDP even when such domains are semialgebraic. This difficulty can be avoided by considering each individual LSF separately.

Therefore, the polynomial optimization program corresponding to the j-th LSF is

$$\begin{aligned} &\min_{\delta\in\Delta} \phi(t(\delta,d)), \text{ subject to } \delta \in K \\ &\text{with } K := \big\{\delta \in \mathcal{R}^{n_\delta} \mid g_j(\delta) \geq 0, \\ &\qquad \|\delta - \delta_l^\star\| - r_l \geq 0,\ l = [1,k]\big\}, \end{aligned} \quad (12)$$

where all inequality constraints are basic semialgebraic. Since K is defined as the intersection of the constraints, it is also basic semialgebraic. Following the redefinition in Eq. (11) we can now search for the largest γ such that the set $\{\delta \in K,\ \phi(t(\delta,d)) - \gamma \geq 0\}$ is empty. Using the Positivstellensatz of Putinar (1993), which is an example of a theorem of the alternative (and is possible over the Positivstellensatz of Stengle (1974) as K is compact), this is equivalent to finding

$$\begin{aligned} &\max \gamma, \text{ such that } \phi(t(\delta,d)) \\ &\quad -\gamma + s_1(\delta) + s_2(\delta) g_j(\delta) \\ &\quad + \sum_{l=1}^{k} s_{l+2}(\delta)(\|\delta - \delta_l^\star\| - r_l) \text{ is SOS}, \end{aligned} \quad (13)$$

where the s_i for $i = 1,\ldots,k+2$ are SOS multipliers. This is now in the form of an SOS program that can readily be solved using standard SDP algorithms and certified using Positivstellensatz certificates (Lacerda and Crespo, 2017). Thus, by running Algorithm 1 once for each $g_j(\delta)$ we can provide a certificate on the global optimality of the detected MLPs. The requirement for using this approach is that the LSFs are polynomial. If this is not the case, standard gradient-based algorithms

can always be used with the advantage that all LSFs can be considered together, but at the risk of missing some MLPs.

4. MLP-based Importance Sampling

Now that a set of MLPs have been identified, we propose an importance sampling algorithm for accurately estimating the failure probability. The importance sampling distribution should locally approximate the target distribution (in this case an SN) in the neighborhood of the MLPs.

A general importance sampling approach utilizes the transformation

$$p_f(\theta) = \int_\Delta \mathcal{I}(\delta \in \mathcal{F}) \frac{f_\delta(\delta)}{f_q(\delta)} f_q(\delta) d\delta, \quad (14)$$

where $\mathcal{I}(\delta \in \mathcal{F})$ is the indicator function which takes value of 1 if $\delta \in \mathcal{F}$ and 0 otherwise. The integrand has been multiplied and divided by an arbitrary distribution $f_q(\delta)$, called the importance sampling distribution, to be chosen by the analyst. The sampling based estimate of this integral is then

$$p_f(\theta) \approx \frac{1}{n_s} \sum_{i=1}^{n_s} \mathcal{I}\left(\delta^{(i)} \in \mathcal{F}\right) \frac{f_\delta(\delta^{(i)})}{f_q(\delta^{(i)})}. \quad (15)$$

A sensible choice of $f_q(\delta)$ can greatly reduce the sampling error introduced in estimating p_f and allow accurate estimation with just a few samples. Several choices of $f_q(\delta)$ reduce Eq. (15) to more recognizable (but not necessarily efficient) approximations, namely:

(1) $f_q(\delta) = f_\delta(\delta)$. In this case the sampling distribution and the target distribution are equivalent and Eq. (15) reduces to a standard Monte Carlo estimate

$$p_f(\theta) \approx \frac{1}{n_s} \sum_{i=1}^{n_s} \mathcal{I}\left(\delta^{(i)} \in \mathcal{F}\right). \quad (16)$$

This estimate is very inefficient for small failure probabilities. In fact, the coefficient of variation is

$$\epsilon_{p_f} = \sqrt{\frac{1 - p_f}{n_s p_f}}, \quad (17)$$

from which it is readily observed that for an accuracy $\epsilon_{p_f} \sim 10\,\%$, when p_f is of the order $O(10^{-k})$, around $O(10^{k+2})$ samples are required. For instance, the estimation of a failure probability of 1×10^{-6} to $10\,\%$ accuracy requires 1×10^8 samples.

(2) $f_q(\delta) \sim \mathcal{U}(\underline{\delta}, \overline{\delta})$. In this case the sampling distribution is uniform over Δ, resulting in the estimation

$$p_f(\theta) \approx \frac{\Delta_V}{n_s} \sum_{i=1}^{n_s} \mathcal{I}\left(\delta^{(i)} \in \mathcal{F}\right) f_\delta(\delta^{(i)}). \quad (18)$$

Sampling from a uniform distribution over the support is very inefficient when the failure region has a much smaller volume.

(3) $f_q(\delta) = f_q^\star(\delta) = \frac{1}{p_f}\mathcal{I}(\delta \in \mathcal{F}) f_\delta(\delta)$. This is the theoretically optimal sampling distribution and leads to a zero variance estimate of p_f (Ang et al., 1992). However, it is not realizable in practice since the normalization constant is itself p_f, which is the ultimate quantity of interest.

Whilst the optimal distribution in (3) is not realizable, it sheds light on the structure that a good IS distribution should inherit. Namely, we want an importance sampling distribution that mimics the target distribution in the subregions of the failure domain having a comparatively large likelihood, while having the smallest probability of success.

The approach proposed uses the MLPs to construct and important sampling distribution. This distribution is given by a Gaussian mixture model having their means at the MLPs and their covariances being equal to the weighted empirical covariance[a] of sample sets drawn in the vicinity of the MLPs. Furthermore, the weights of the mixed Gaussian are made equal to the likelihood at the corresponding MLP.

Means to compute the weighted covariance are detailed next. Consider a uniform random vector

$$X \sim \mathcal{U}(\delta_k^\star - l/2, \delta_k^\star + l/2),\ X \in \mathcal{R}^{n_\delta}, \quad (19)$$

and samples $x^{(i)}$ with $i = 1, \ldots, n_w$ taken from X. The weighted empirical covariance matrix Σ^w, where sample weights are given by the likelihood

[a] Samples of a uniform distribution clustered around an MLP are weighted according to the value of the likelihood for the target distribution therein.

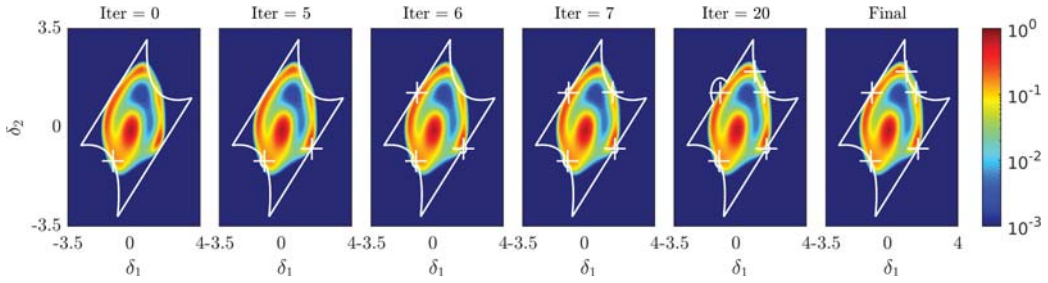

Fig. 1. Upper views of the SN density and contours of the LSF showing the progression of the search for the MLPs (x's) with increasing iteration.

of the SN at the sample, can be calculated from

$$\Sigma^w = \frac{\sum_{i=1}^{n_w} f_\delta(x^{(i)})(x^{(i)} - \mu^w)(x^{(i)} - \mu^w)^T}{\sum_{i=1}^{n_w} f_\delta(x^{(i)})}, \tag{20}$$

where μ^w is the weighted sample mean

$$\mu^w = \frac{\sum_{i=1}^{n_w} f_\delta(x^{(i)}) x^{(i)}}{\sum_{i=1}^{n_w} f_\delta(x^{(i)})}. \tag{21}$$

The importance sampling density is then

$$f_q(\delta) = \sum_{i=1}^{k} f_\delta(\delta_i^\star) \mathcal{N}_i(\delta_i^\star, \Sigma_i^w). \tag{22}$$

Since sampling is needed only from a uniform distribution in the vicinity of each MLP, this method scales well to high dimensions.

5. Numerical Experiments

We now apply the MLP-based importance sampling method to two test cases, one with $n_\delta = 2$ to visualize the method and a second with $n_\delta = 7$ to show its performance in higher dimensions. In both instances we compare the results against sampling uniformly and from the SN itself.

5.1. *Four branch function*

In this example we use an SN density obtained from modelling a Van der Pol oscillator with uncertain input parameters. The failure domain is composed of a modified four branch function giving a true probability of failure of $p_f = 0.021$. This MLP problem can be solved by both SDP and gradient-based optimization as the LSFs are polynomial. Figure 1 shows the five MLPs found

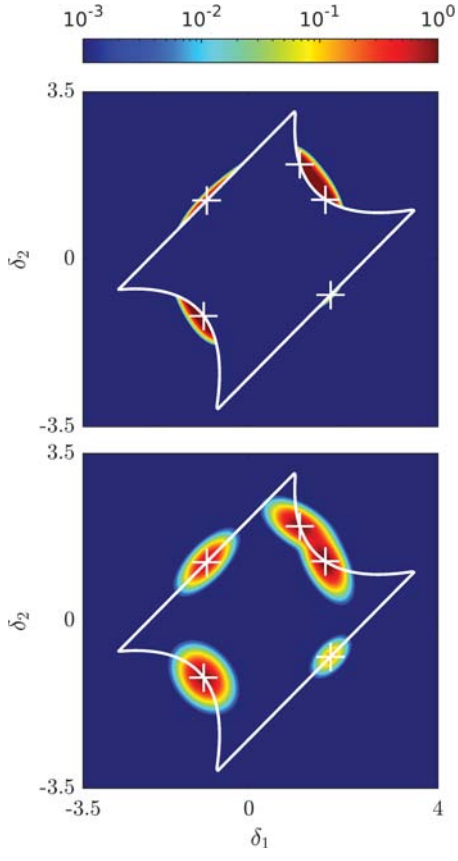

Fig. 2. Upper views of optimal (top) and MLP-based (bottom) importance sampling densities.

in the search. As the search progresses the areas of artificial safety grow around the MLPs, deducting from the failure domain and allowing subsequent MLPs to be found. Optimal points

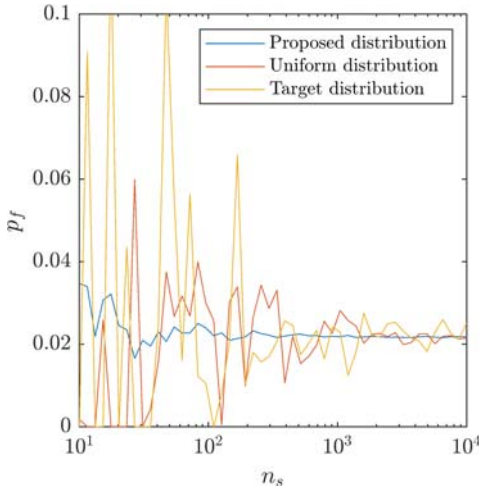

Fig. 3. Convergence of p_f when sampling from the proposed MLP-based IS distribution, a uniform distribution and the target distribution.

may jump between the surfaces of the spheres around the current MLPs, and it is the sphere on which the most recent optimum exists that has its radius increased. This explains why, for example, the sphere around MLP_3 continues to grow after MLP_4 has been found. The result is the detection of 5 distinct MLPs, where two were found in the top right part of the branch function. Visual inspection confirms the correct identification of all points, and the same set of points were obtained using SDP. The entire search took 4s to complete on a single Intel® Core™ i7-8665U processor[b] highlighting the efficiency of the algorithm.

Figure 2 shows the optimal IS density and the constructed MLP-based IS density. Note that the Gaussians are aligned with the target distribution locally. Figure 3 shows the convergence of the three IS techniques as n_s increases. The MLP-based IS density reduces the required n_s for accurate calculation of p_f by an order of magnitude over uniform sampling and two orders of magnitude over sampling directly from the SN. Further, the method is guaranteed to account for modes of failure with small p_f, which a standard Monte-Carlo campaign might miss entirely.

5.2. *Higher dimensional example*

We now consider an example SN in $n_\delta = 7$ dimensions with a failure domain consisting of 15 individual failure functions $g_j(\delta)$. To help visualize

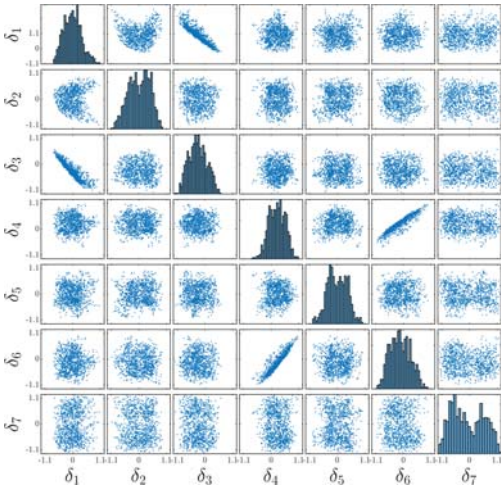

Fig. 4. Samples drawn from SN density with $n_\delta = 7$.

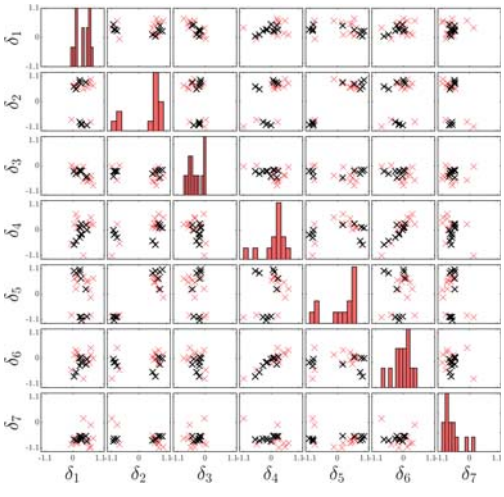

Fig. 5. Sample failures (red x's) and the set of identified MLPs (black x's) for the SN. Diagonal shows histograms of sample failures.

[b]The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

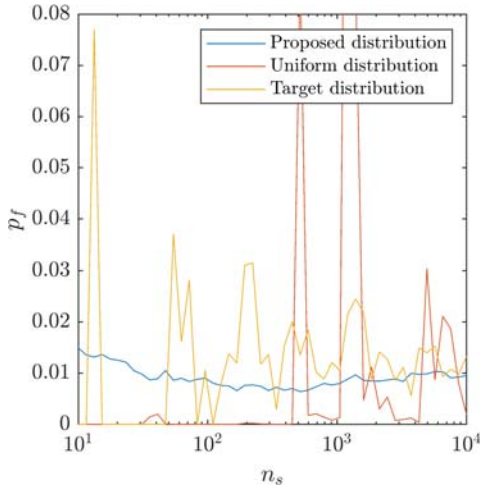

Fig. 6. Convergence of p_f when sampling from the proposed MLP-based IS distribution, a uniform distribution and the target distribution.

the SN, Figure 4 shows empirical marginal densities of samples drawn from it along the diagonal, whereas the off-diagonal plots are projections of the cloud of parameters onto 2-dimensional subspaces. Strong parameter dependencies among δ_1-δ_3, and δ_4 -δ_6, are apparent. Figure 5 shows empirical marginal densities of the samples falling into the failure domain along the diagonal, and the sample failures are plotted on 2-dimensional subspaces alongside the detected MLPs on off-diagonal plots. Note that samples falling into the failure domain are clustered around the MLPs, suggesting the MLP search successfully identified the failure modes. Finally, Figure 6 shows the convergence of p_f for the IS distribution against uniform sampling and sampling from the SN itself. The MLP search took 170s to run on a single Intel® Core™ i7-8665U processor with an additional 6s to calculate p_f from the IS distribution, highlighting the efficiency of the algorithm in higher dimensions

6. Conclusions

This paper presents an IS technique based on discovering a set of most likely points of failure (MLPs). The IS distribution is constructed as a weighted Gaussian mixture with centroids located at the identified MLPs. The method allows the efficient computation of failure probabilities where there are multiple modes in possibly disconnected failure region, and thus where FORM/SORM methods would fail. Whilst there is some overhead cost required to search for the MLPs, only a few samples are required in targeted areas of the support to compute the covariance around each MLP. Consequently, the method performs well in higher dimensional settings and has shown efficient computation up to $n_\delta = 7$.

References

Ahmadi, A. A. (2018). Sum of Squares (SOS) Techniques: An Introduction. Technical report, Princeton ORFE.

Ang, G. L., A. H. Ang, and W. H. Tang (1992). Optimal Importance-Sampling Density Estimator. *Journal of Engineering Mechanics 118*(6), 1146–1163.

Crespo, L. G., B. K. Colbert, S. P. Kenny, and D. P. Giesy (2019). On the quantification of aleatory and epistemic uncertainty using Sliced-Normal distributions. *Systems and Control Letters 134*, 104560.

Crespo, L. G., B. K. Colbert, T. Slager, and S. P. Kenny (2021). Robust Estimation of Sliced-Exponential Distributions. In *60th IEEE Conference on Decision and Control (CDC)*, Austin, TX, pp. 6742–6748. IEEE.

Der Kiureghian, A. (2005). *First- and Second-order reliability methods*, Volume 14. Boca Raton, FL: CRC Press.

Hohenbichler, M., S. Gollwitzer, W. Kruse, and R. Rackwitz (1987). New light on first- and second-order reliability methods. *Structural Safety 4*(4), 267–284.

Lacerda, M. J. and L. G. Crespo (2017). Interval predictor models for data with measurement uncertainty. *Proceedings of the American Control Conference*, 1487–1492.

Packard, A., U. Topcu, P. Seiler, and G. Balas (2010). Help on SOS. *IEEE Control Systems Magazine 30*(4), 18–23.

Papaioannou, I., C. Papadimitriou, and D. Straub (2016). Sequential importance sampling for structural reliability analysis. *Structural Safety 62*, 66–75.

Putinar, M. (1993). Positive Polynomials on Compact Semi-algebraic Sets. *Indiana University Mathematics Journal 42*(3), 969–984.

Stengle, G. (1974). A nullstellensatz and a positivstellensatz in semialgebraic geometry. *Mathematische Annalen 207*(2), 87–97.

Tabandeh, A., G. Jia, and P. Gardoni (2022). A review and assessment of importance sampling methods for reliability analysis. *Structural Safety 97*, 102216.