Uncertainty Management of Safety-Critical Systems: A Solution to the Back-Propagation Problem

Marco de Angelis
Graduate Student, Inst. for Risk and Uncertainty, School of Eng., Univ. of Liverpool, Liverpool, UK

Edoardo Patelli
Lecturer, Inst. for Risk and Uncertainty, School of Eng., Univ. of Liverpool, Liverpool, UK

Michael Beer
Professor, Inst. for Risk and Uncertainty, School of Eng., Univ. of Liverpool, Liverpool, UK

ABSTRACT: In many engineering applications, the assessment of reliability has to be done within a limited amount of information, which does not allow to use exact values for the distributional hyper-parameters. This is achieved defining probability boxes and assessing the reliability computing the failure probability bounds. Probability boxes are often obtained from known probability distribution functions represented by interval hyper-parameters. In the applications, not only it is of interest estimating the failure probability bounds, but it is also required to identify the extreme realizations leading to the estimated bounds. In this paper, we propose a strategy, based on the Kolmogorov-Smirnov test, to identify the parental distribution function that best fit the distribution of extreme realizations, obtained from the min-max propagation. From the results obtained comparing the strategy with a direct search, it has emerged that the proposed method is generally applicable and efficient.

1. INTRODUCTION

In reliability assessment it is of interest computing the effect of epistemic uncertainty on the failure probability, (see e.g. Patelli et al. (2014), Roy and Oberkampf (2010)) and making the least amount of assumptions (see also Beer et al. (2013)). This requires the epistemic uncertainty to be propagated throughout the model and consequently quantified in terms of failure probability intervals. Uncertainty propagation can be performed by means of different strategies, nonetheless a general distinction can be drawn between parametric (de Angelis et al. (2015), Zaffalon (2002)) and non-parametric approaches (Alvarez (2006), Ferson et al. (2002), Kreinovich (1997)). The non-parametric proves to be a more general approach because only bounds of the statistical input quantities are concerned and there is no need to specify any parental probability distributions. However, this approach comes with some limitations, in fact, once the bounds of the interval probability are computed, it is not possible to go back and identify the arguments in the input space corresponding to such bounds.

In this paper we propose a numerical strategy to resolve the issue of back-tracking the failure probability bounds in the input space, using simulation methods and model updating procedure. The interval failure probability is obtained from the Dempster-Shafer (D-S) structure of the output quantity of interest (Dempster (1967)). Any level of the D-S structure identifies a minimum and a maximum of the quantity of interest produced by different searching domains. The expected value of all the minima (maxima) of the D-S structure represents the lower (upper) bound of the expected value of the quantity of interest. To any minimum (maximum) corresponds an argument minimum (maximum) in the input space. Therefore, by gathering all the minima (maxima), and track these minima (maxima) back to the input space, it is possible to construct a cumulative distribution function (CDF) of the corresponding argument minima (maxima). Equivalently, but this time selecting the right combination of extrema, realizations in the input space corresponding to the lower and upper bounds of the failure probability can be identified. These realizations are then gathered to make up a cumulative distribution function for each bound of the failure probability, and subsequently interpreted, for example, by means of model updating techniques. Here, model updating is used to identify the best probability function that fits the two obtained CDFs corresponding to the lower and upper bounds of the failure probability. Eventually, a (parental) probability distribution can be chosen to best fit the resulting CDFs.

2. RELIABILITY ASSESSMENT WITH PROBABILITY BOXES

2.1. *Brief introduction to probability boxes*

Probability boxes (p-boxes) extend the definition of reliability to an interval of possible alternatives, enclosed by a lower and a upper bound. In reliability and risk assessment, p-boxes are invoked to represent what in literature is referred to as uncertainty of Type III, which include both aleatory and epistemic uncertainties. Let $\overline{F}$ and $\underline{F}$ be non-decreasing function from the real line $\mathbb{R}$ into $[0,1]$ and $\underline{F} \leq \overline{F}$ for all $x \in \mathbb{R}$. A p-box is the set of all non-decreasing functions $F : \mathbb{R} \rightarrow [0,1]$, such that $\underline{F}(x) \leq F(x) \leq \overline{F}(x)$.

2.2. *P-box convolution by means of Monte Carlo simulation*

A Dempster-Shafer structure can be seen as the "discrete" equivalent of a p-box and it is key for the reliability assessment of systems. The input-output convolution of p-boxes, in practice, is performed by using Monte Carlo simulation techniques (Kreinovich et al. (1991)). Reliability assessment is, therefore, performed by i) sampling the equivalent D-S structure of the p-boxes, ii) obtaining the output D-S structure, iii) and ultimately estimating the failure probability bounds.

2.3. *Failure probability upper and lower bounds*

Let $\mathscr{G} : \mathbb{R}^n \rightarrow \mathbb{R}$ be the system performance function, $\theta \in \mathbb{R}^n$ be a vector of p-boxes, and Ω_F be the domain of unacceptable states (or failure domain), such that $\Omega_F = \{\theta : \mathscr{G}(\theta) \leq 0\}$. The system performance is evaluated as $g = \mathscr{G}(\theta)$. Each focal element, $\underline{\overline{\theta}}^{\{s\}}$, of the D-S structure is propagated throughout the system, and the corresponding image is obtained as

$$\mathscr{G}(\underline{\overline{\theta}}^{\{s\}}) = [\underline{g},\, \overline{g}]^{\{s\}}; \tag{1}$$

where,

$$\underline{g} = \min_{\theta \in \underline{\overline{\theta}}^{\{s\}}} \mathscr{G}(\theta); \quad \overline{g} = \max_{\theta \in \underline{\overline{\theta}}^{\{s\}}} \mathscr{G}(\theta). \tag{2}$$

The propagation of individual focal elements leads to the failure probability bounds, obtained using the plausibility and belief function as

$$\underline{p_F} = \lim_{N_s \rightarrow \infty} \sum_{\mathscr{G}(\underline{\overline{\theta}}^{\{s\}}) \cap \Omega_F}^{N_s} m(\mathscr{G}(\underline{\overline{\theta}}^{\{s\}})); \tag{3}$$

$$\overline{p_F} = \lim_{N_s \rightarrow \infty} \sum_{\mathscr{G}(\underline{\overline{\theta}}^{\{s\}}) \subseteq \Omega_F}^{N_s} m(\mathscr{G}(\underline{\overline{\theta}}^{\{s\}})); \tag{4}$$

where, m is the mass associated to each focal element of the D-S structure.

2.4. *Reliability assessment with the parametric approach*

The failure probability bounds are obtained by searching among all those distribution functions, which mean values and standard deviations are included in the intervals $\underline{\overline{\mu}}$ and $\underline{\overline{\sigma}}$. The problem can be alternatively formulated as

$$\underline{p_F} = \min_{\underline{\overline{\mu}}, \underline{\overline{\sigma}}} p_F(\mu, \sigma); \quad \overline{p_F} = \max_{\underline{\overline{\mu}}, \underline{\overline{\sigma}}} p_F(\mu, \sigma). \tag{5}$$

Using this approach, the solution is included in the p-box bounding CDFs (see e.g. Figures 1 and 2), and it also belongs to one of the distribution functions defined by the interval hyper-parameters (see e.g. Table 1). In case where the performance function is a black-box, the approach is driven by an optimization procedure that solves the problem of Eq. 5 numerically.

2.5. *Reliability assessment with the non-parametric approach*

Any distribution function contained within the bounding CDFs, even not belonging to any parental distribution model, can be considered for the problem solution. This implies that the failure probability bounds obtained using this approach, are always wider than those obtained from the parametric approach, because a greater set of candidates is searched for. The failure probability bounds are obtained by computing plausibility and belief of the output D-S structure as shown in Eqs. 3 and 4.

One major limitation of the non-parametric approach, despite its efficiency, is the difficulty in identifying the input distribution functions that are responsible for the failure probability bounds. This issue is also known in literature as the tracking (or back-propagation) problem. Here we propose a strategy to tackle this issue.

2.6. *Solution to the tracking problem*

In many applications, it is of interest identifying the distribution functions that lead to the failure probability bounds. Thus, a numerical strategy is needed to characterize the distribution model that best represent the extreme realizations. Here, we propose a strategy, based on the Kolmogorov-Smirnov test, that identifies the function that best-fit the distributions of extrema (Kolmogoroff (1941)). The problem is formulated as

$$\min_{\underline{\mu},\underline{\sigma}} \sup_{x} |F_{N_s}(x) - F(x;\mu,\sigma)|; \qquad (6)$$

where, $\sup_x |F_{N_s}(x) - F(x;\mu,\sigma)|$, is the K-S statistic, which represents a measure of similarity between the CDF obtained with the sample set $\theta^{\{s\}}_{\min(\max)}$ and the distribution function $F(x;\mu,\sigma)$. Solution to the problem is the pair of hyper-parameters (μ^*,σ^*) that minimizes the K-S statistic for each input p-box.

3. NUMERICAL EXAMPLE

The example is formulated to show limitations and advantages of using the two approaches. The following performance function will be considered throughout this section as

$$g = x^2\, y + e^x; \qquad (7)$$

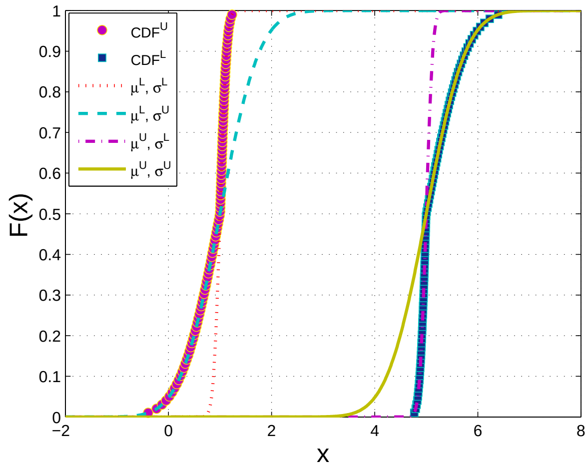

Figure 1: P-box bounding normal CDFs for x

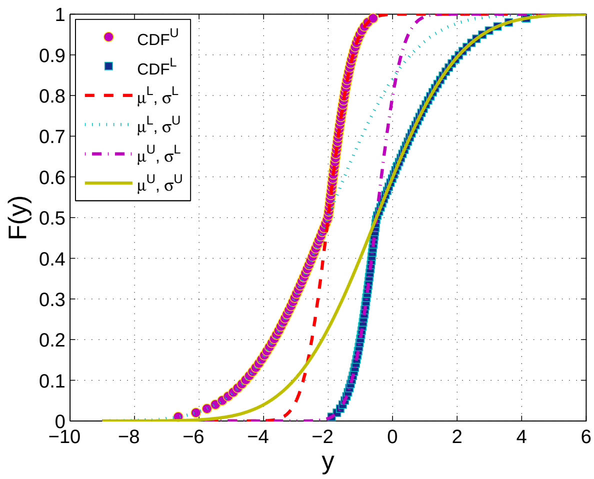

Figure 2: P-box bounding normal CDFs for y

where, x and y are p-boxes obtained from normal distribution functions, which parameters are shown in Tables 1. Let the failure event be defined by the failure region $\Omega_F = \{x,y : \mathscr{G}(x,y) \le 0\}$, then the failure probability is expressed as the interval

$$\underline{\overline{p_F}} = P[g \le 0]. \qquad (8)$$

The p-boxes defined in Table 1 are represented in Figures 1 and 2 in terms of bounding CDFs.

The aim of this example is to identify the failure probability bounds and the corresponding realizations in the input space, i.e. those CDFs that yield the minimum and maximum failure probability.

Depending on how the input space of candidate solutions is searched for, the solution may be signif-

P-box	$\overline{\mu}$	$\overline{\sigma}$	Distribution
x	$[1,\ 5]$	$[0.1,\ 0.6]$	Normal
y	$[-2,\ -0.5]$	$[0.6,\ 2]$	Normal

Table 1: Mean values and standard deviations for the definition of the p-box bounds

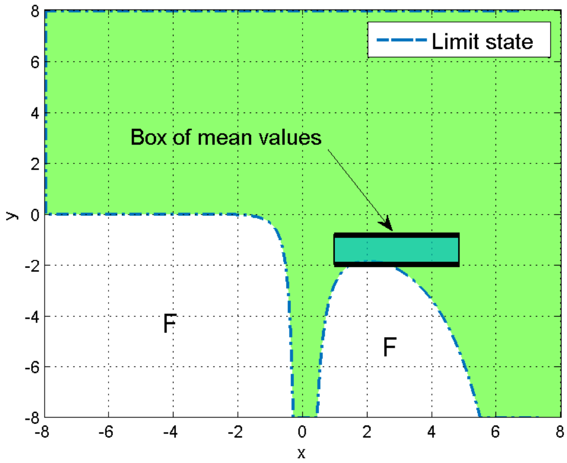

Figure 3: Limit state surface and box of mean values

icantly different. In the next sections two different approaches of searching in the input space of candidate solutions are presented.

3.1. The parametric approach

In this simple case, the optimization can be reduced to a one-dimensional search. In fact, the function of Eq. 3 is monotonically increasing in y, which permits to discharge $\underline{\overline{\mu}}_y$ from the list of candidates, as $\underline{\mu}_y$ and $\overline{\mu}_y$ corresponds to $\overline{p_F}$ and $\underline{p_F}$, respectively. Also, the standard deviations can be taken out of the optimization as only four candidate solutions can be identified, which correspond to the four corners of the domain $\overline{\sigma}$. This becomes more evident if we look to the limit sate surface and the box of candidate mean values, as shown in Figure 3. The optimization is, therefore, reduced to a search along the thick (upper and lower) edges represented in Figure 3. Given the shape of the limit state surface, we expect the maximum failure probability to be located somewhere near the peak of the limit sate surface. Minimum and maximum failure probabilities are obtained on the upper and lower edges of the

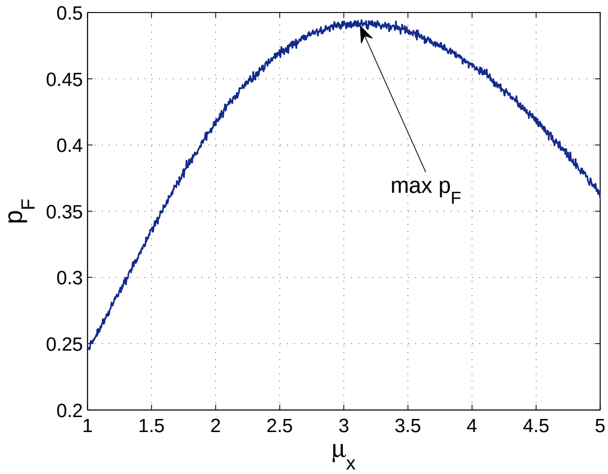

Figure 4: Failure probability values obtained from the optimization along the lower μ_x edge

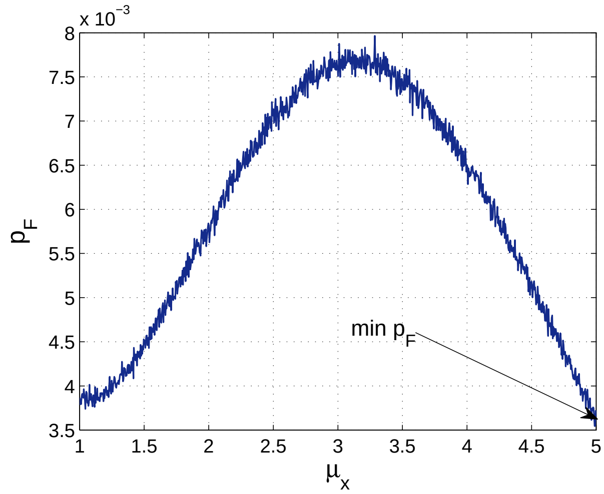

Figure 5: Failure probability values obtained from the optimization along the upper μ_x edge

μ_x domain respectively, populating the space with 1000 realization. On each realization, the failure probability is estimated using MC simulation with 10^5 samples. Results from the edge optimization are shown in Figure 4, where it is shown that the maximum failure probability (unlike the minimum) is attained within the edge, thus not at the corners of the domain. Figure 5 shows that the minimum failure probability is held at the right endpoints of the domain, i.e. for $\mu_x = 5$. The argument optima and corresponding failure probability extrema are reported in Table 2.

P-box	min p_F $5.30\ 10^{-9}$	max p_F 0.493
x	$(\mu_x)_{\cdot} =$ 5.0	$(\mu_x)^{\cdot} =$ 3.49
x	$(\sigma_x)_{\cdot} =$ 0.1	$(\sigma_x)^{\cdot} =$ 0.1
y	$(\mu_y)_{\cdot} =$ -0.5	$(\mu_y)^{\cdot} =$ -2.0
y	$(\sigma_y)_{\cdot} =$ 0.6	$(\sigma_y)^{\cdot} =$ 2.0

Table 2: Failure probability bounds and corresponding extreme normal distributions

3.2. *The non-parametric approach*

The failure probability bounds are obtained by generating a great number of samples, from the p-boxes, and subsequently constructing the associated D-S structures of the response. The procedure for constructing the D-S structure of the response is briefly summarized as,

1. draw a uniform random number, $\alpha^{\{s\}}$, for each p-box, between 0 and 1;

2. get the sample endpoints $\underline{\overline{x}}^{\{s\}}$ using the inverse bounding CDFs, $\underline{\overline{F}}^{-1}$ as;

$$\underline{x}^{\{s\}} = \overline{F}^{-1}\left(\alpha^{\{s\}}\right);\ \overline{x}^{\{s\}} = \underline{F}^{-1}\left(\alpha^{\{s\}}\right); \qquad (9)$$

3. identify minimum, $\underline{r}^{\{s\}}$, and maximum response, $\overline{r}^{\{s\}}$, within the search domain, $\underline{\overline{x}}^{\{s\}}$. This step is also referred to as min-max propagation;

4. repeat the above steps for s from 1 to N_s, i.e. loop over the number of samples N_s;

5. collect samples and corresponding response extrema.

Once the D-S structure of the response is obtained, the failure probability bounds are obtained from the D-S plausibility and belief as

$$\underline{p_F} = \lim_{N_s \to \infty} \frac{1}{N_s} \sum_{s=1}^{N_s} \mathscr{I}[\overline{g}^{\{s\}} < 0]; \qquad (10)$$

$$\overline{p_F} = \lim_{N_s \to \infty} \frac{1}{N_s} \sum_{s=1}^{N_s} \mathscr{I}[\underline{g}^{\{s\}} < 0]; \qquad (11)$$

where, $\mathscr{I} : \mathbb{R} \to \{0, 1\}$ is the indicator function. Most of the attention, in the above procedure, is usually given to the min-max propagation step. In fact, this can be troublesome, especially if the response is the output of a black-box model, thus the propagation is done by invoking global optimization algorithms (such as evolutionary or stochastic algorithms). However, in this case, solution to the propagation task can be found analytically, as the performance function is given explicit mathematical expression.

3.2.1. *The min-max propagation*

The performance function of Eq. 3 is monotonically increasing with respect to y, which is a great advantage as it excludes the presence of relative minima and maxima. Moreover, it implies that, for every value of x, as the variable y decreases/increases so does the performance function. This leads to the following relationships

$$\underline{g} = x_*^2\, \underline{y} + e^{x_*};\quad \overline{g} = x^{*2}\, \overline{y} + e^{x^*}; \qquad (12)$$

where x_* and x^* are yet to be determined. On the other side, the performance function is not monotonic with respect to x. The sign of the first and second derivatives of g, says that the function is monotonically increasing with respect to x only in the portion where $y \in [-1.37, 0]$. Whereas, for $y \in (-\infty, -1.37) \cup (0, \infty)$ the function may have a minimum or maximum. Within the latter portion of domain, the minimum/maximum is identified solving for x the partial derivative $\partial g / \partial x$, and subsequently checking if the obtained value is smaller/greater than the values at the endpoints $\underline{x}$ and $\overline{x}$.

3.2.2. *Solution to the back-propagation problem*

The problem is addressed collecting all those realizations in the input space that correspond to the response extrema. It is, thus, crucial solving the min-max propagation problem by keeping track, back to the input space, of all the minima and maxima. These are also referred to as extreme realizations. The failure probability bounds are computed by means of Eqs. 10 and 11 using 10^5 MC samples.

$$\underline{\overline{p_F}} = [\underline{p_F},\ \overline{p_F}] = [0,\ 0.754].$$

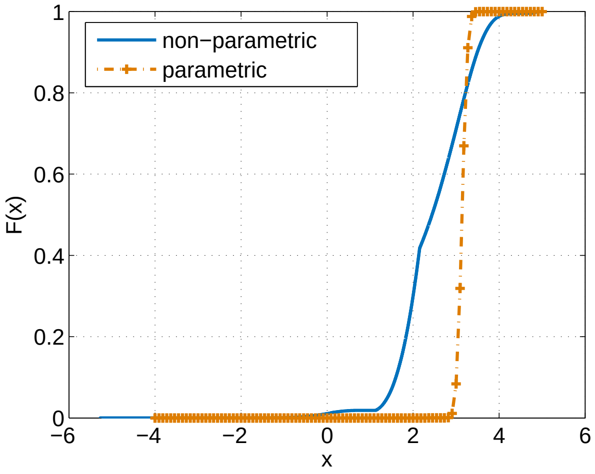

Figure 6: CDFs from the x p-box of extreme realizations leading to the failure probability upper bound

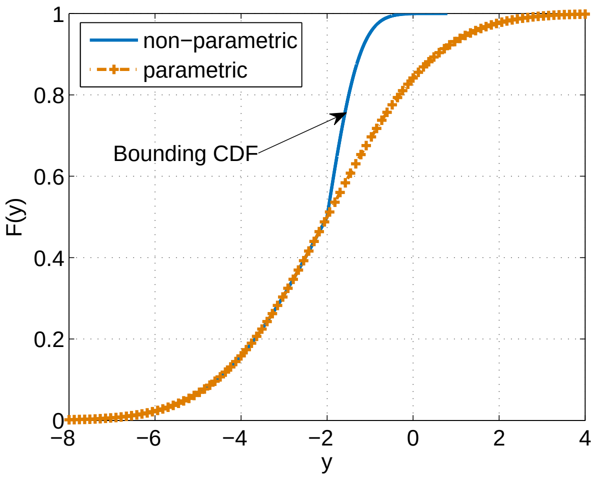

Figure 7: CDFs from the y p-box of extreme realizations leading to the failure probability upper bound

The strategy intends to use the distributions of minima and maxima as probability models corresponding to the upper and lower bounds of the failure probability, respectively. In Figures 6 and 7 the distribution of minima, corresponding to the upper failure probability, are compared with the extreme normal distributions of Table 2 for the maximum failure probability, obtained by solving the optimization problem. Figure 7 shows quite clearly that the normal distribution obtained from the parametric approach, i.e. corresponding to the maximum failure probability, fits quite well the distribution of minima obtained from the non-parametric approach. Figure 6 also shows a good fit, although, this time, it is clear that the normal distribution does not represent the best fit.

P-box	min p_F^* 1.8 10^{-5}		max p_F^* 0.468	
x	$(\mu_x). =$	4.97	$(\mu_x)^{\cdot} =$	2.43
x	$(\sigma_x). =$	0.23	$(\sigma_x)^{\cdot} =$	0.60
y	$(\mu_y). =$	-0.51	$(\mu_y)^{\cdot} =$	-1.98
y	$(\sigma_y). =$	1.10	$(\sigma_y)^{\cdot} =$	1.07

Table 3: Failure probability bounds and corresponding extreme normal distributions obtained with the non-parametric approach

The solution to the back propagation problem can be found by selecting in the space of parental (normal) distribution functions, those hyper-parameters corresponding to the min/max failure probability. Within the non-parametric approach, this can be done searching for the parental distribution functions that provides the best fit to the collected distributions of minima and maxima. Here, the normal distribution that best-fit the extreme realizations is obtained using the Kolmogorov-Smirnov test, by minimizing the statistic (k-s distance) $D_{N_s} = \sup_x |F_{N_s}(x) - F(x)|$. The results from the the k-s distance minimization are shown in Table 3. Figures 8 and 9 show the x and y extreme normal distributions obtained for the failure probability upper bound. It is interesting to see that the two normal distribution functions of extreme values obtained using the two approaches do not differ much. However, the failure probability bounds obtained using the proposed strategy, as it can also be seen in Table 3, are not the optimal ones, as they are enclosed in the ones obtained using the parametric approach,which are shown in Table 2. Note, from Figure 9 that the extreme realizations of p-box y are distributed as the upper CDF as the model is monotonic with respect to this variable.

3.3. Final remarks

From the analysis of the extreme realizations with both parametric and non-parametric approaches we may conclude that

– if the response is monotonic with respect to a

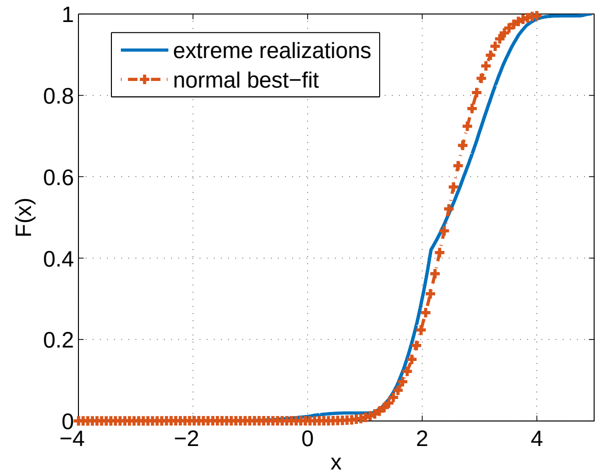

Figure 8: Failure probability upper bound: the normal distribution, from the x p-box, that best-fit the extreme realizations is obtained with a k-s distance of 0.11

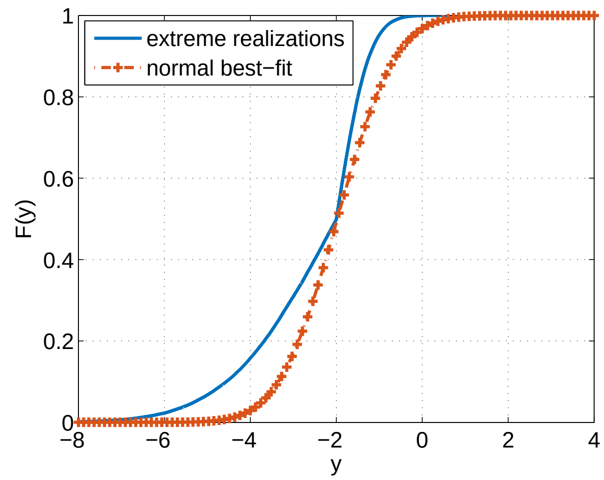

Figure 9: Failure probability upper bound: the normal distribution, from the y p-box, that best-fit the extreme realizations is obtained with a k-s distance of 0.15

p-box, the failure probability bounds are obtained from the bounding CDFs of that p-box,

– if the response is not monotonic with respect to a p-box, the distribution function of the extreme realizations is enclosed in the bounding CDFs of that p-box and may have a complicated form,

– in general, the reconstructed CDF of the extreme realizations is not distributed as the parental model of probability

– if the response is monotonic with respect to all p-boxes and if, for every p-box, the bounding CDFs are made of only two distribution functions (such as in the Beta model), the solution from the two approaches coincides.

4. CONCLUSIONS

In many engineering applications the assessment of reliability requires the consideration of uncertainty in the form of probability boxes. Very often, probability boxes are defined using known probability distribution functions represented by interval hyper-parameters. In these cases, it is of interest not only estimating the bounds on the output statistical quantity of interest, such as the failure probability, but it is also required to identify which extreme realizations led to the estimated bounds. While in some cases it may be sufficient just knowing what the mass function of these realizations is, in other cases it may be necessary to know what is the closest distribution function, from the underlying model of probability. In this paper, we have proposed a strategy, based on the Kolmogorov-Smirnov test, to identify the parental distribution function that is closest to the distribution of extreme realizations. The strategy collects the realizations from the min-max propagation, and search, in the space of feasible hyper-parameters, for the distribution function that best-fit the collected data. From the results obtained comparing the strategy with a direct search, performed by means of the parametric approach, it has emerged that the proposed method works well and shows also to be quite efficient. However, the accuracy of the strategy might not be satisfactory for the lower bound of the failure probability, which is different from the optimal value by orders of magnitude. An improvement in this direction can be sought, for example, putting more emphasis on the tails of the fitted distributions. Also, as one more future research direction, it would be interesting to see how choosing more than one probability model at a time can increase the confidence of the estimation and lead to a better representation of the extreme realizations.

5. REFERENCES

Alvarez, D. A. (2006). “On the calculation of the bounds of probability of events using infinite random sets.” *International journal of approximate reasoning*, 43(3), 241–267.

Beer, M., Ferson, S., and Kreinovich, V. (2013). “Imprecise probabilities in engineering analyses.” *Mechanical systems and signal processing*, 37(1), 4–29.

de Angelis, M., Patelli, E., and Beer, M. (2015). “Advanced line sampling for efficient robust reliability analysis.” *Structural Safety*, 52, 170–182.

Dempster, A. P. (1967). “Upper and lower probabilities induced by a multivalued mapping.” *The annals of mathematical statistics*, 325–339.

Ferson, S., Kreinovich, V., Ginzburg, L., Myers, D. S., and Sentz, K. (2002). *Constructing probability boxes and Dempster-Shafer structures*, Vol. 835. Sandia National Laboratories.

Kolmogoroff, A. (1941). “Confidence limits for an unknown distribution function.” *The Annals of Mathematical Statistics*, 12(4), 461–463.

Kreinovich, V. (1997). “Random sets unify, explain, and aid known uncertainty methods in expert systems.” *Random Sets*, Springer, 321–345.

Kreinovich, V. Y., Bernat, A., Borrett, W., Mariscal, Y., and Villa, E. (1991). “Monte-carlo methods make dempster-shafer formalism feasible.

Patelli, E., Alvarez, D. A., Broggi, M., and Angelis, M. d. (2014). “Uncertainty management in multidisciplinary design of critical safety systems.” *Journal of Aerospace Information Systems*, 1–30.

Roy, C. J. and Oberkampf, W. L. (2010). “A complete framework for verification, validation, and uncertainty quantification in scientific computing.” *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 4–7.

Zaffalon, M. (2002). “The naive credal classifier.” *Journal of statistical planning and inference*, 105(1), 5–21.