



Sequential Bayesian optimal experimental design for structural reliability analysis

Christian Agrell^{1,2} · Kristina Rognlien Dahl¹

Received: 1 July 2020 / Accepted: 1 February 2021
© The Author(s) 2021

Abstract

Structural reliability analysis is concerned with estimation of the probability of a critical event taking place, described by $P(g(\mathbf{X}) \leq 0)$ for some n -dimensional random variable \mathbf{X} and some real-valued function g . In many applications the function g is practically unknown, as function evaluation involves time consuming numerical simulation or some other form of experiment that is expensive to perform. The problem we address in this paper is how to optimally design experiments, in a Bayesian decision theoretic fashion, when the goal is to estimate the probability $P(g(\mathbf{X}) \leq 0)$ using a minimal amount of resources. As opposed to existing methods that have been proposed for this purpose, we consider a general structural reliability model given in hierarchical form. We therefore introduce a general formulation of the experimental design problem, where we distinguish between the uncertainty related to the random variable \mathbf{X} and any additional epistemic uncertainty that we want to reduce through experimentation. The effectiveness of a design strategy is evaluated through a measure of residual uncertainty, and efficient approximation of this quantity is crucial if we want to apply algorithms that search for an optimal strategy. The method we propose is based on importance sampling combined with the unscented transform for epistemic uncertainty propagation. We implement this for the myopic (one-step look ahead) alternative, and demonstrate the effectiveness through a series of numerical experiments.

Keywords Optimal experimental design · Structural reliability · Probability of failure · Epistemic and aleatory uncertainty · Unscented transform

1 Introduction

In order to ensure sufficient reliability of engineered systems, such as buildings, ships, offshore structures, aircraft or technological products, uncertainties with respect to the system's capabilities and the system's environment must be accounted for. In probabilistic structural reliability analysis, this is achieved through a probabilistic model of the system

and its environment. A primary objective with such a model is to estimate the probability that the system will fail (e.g. collapse, sink, crash or explode).¹

A probabilistic structural reliability model is commonly defined through a *performance function* (also called a *limit-state function*) $g(\mathbf{X})$ depending on some random variable \mathbf{X} . Here, $g(\mathbf{X}) < 0$ corresponds to system failure, and $g(\mathbf{X}) \geq 0$ corresponds to the system functioning. Typically, \mathbf{X} contains the parameters describing a particular structure, such as the geometry, dimensions and material properties. These quantities may be random, but can be influenced by the designer of the structure. For example, the designer may choose to use a more expensive, but more durable material in order to improve the structural properties of the system. In addition, \mathbf{X} contains the (random) parameters that characterize the sys-

✉ Christian Agrell
chrisagr@math.uio.no; christian.agrell@dnvgl.com
Kristina Rognlien Dahl
kristrd@math.uio.no

¹ Department of Mathematics, University of Oslo, Oslo, Norway

² DNV GL Group Technology and Research, Høvik, Norway

¹ This is rarely interpreted as a frequentist probability. As the model is not the real world, it is common to design models such that the failure probability can be interpreted as a conservative estimate, or as a consistent measure of robustness for comparison with other 'acceptable' systems.

tems environment, such as wind speed, wave height etc., and parameters describing how well the model fits reality (model uncertainties). Given \mathbf{X} and the function $g(\cdot)$, the *probability of failure* is defined as the probability $P(g(\mathbf{X}) < 0)$. Modern engineering requirements for safe design and operation of such systems are usually given as an upper bound on this probability (Madsen et al. 2006).

Hence, for many practical applications, the failure probability computation is an important task. This is often challenging for complex systems, as a computationally feasible stochastic model of the complete system and its environment is not available. To capture this in our modelling framework, we consider additional *epistemic* uncertainties, i.e. uncertainties due to limited data or knowledge that *in principle* can be reduced by gathering more information.

1.1 Epistemic and aleatory uncertainty

The concept of epistemic uncertainty is commonly used in uncertainty quantification (UQ) and in reliability analysis. One often considers two different kinds of uncertainty: Aleatory (stochastic) and epistemic (knowledge-based) uncertainty. Aleatory uncertainty is uncertainty which cannot be removed by collecting more or better information. For instance, the result of throwing a dice is an example of aleatory information, because there is a range of possible outcomes even if we understand the experimental setup. Epistemic uncertainty, on the other hand, is uncertainty which can be affected by collecting more and/or better information. For example, if a quantity or parameter has a definite value, but this value is unknown to us, then the uncertainty considered epistemic. Likewise, uncertainty about the form of a model for a physical phenomenon is epistemic, because more research or experiments could be performed to improve the model.

We note that this characterization of uncertainties will have to depend on the relevant modelling and decision-making context. Given the aleatory example of throwing a dice, one could argue that given sufficient information about initial conditions together with a detailed physics model, it should be possible to predict the outcome (and the uncertainty is therefore epistemic). But based on the modelling context this may not be relevant or a realistic assumption at all. See for instance (Der Kiureghian and Ditlevsen 2009) for a broader discussion.

The following example illustrates that a random variable may contain both epistemic and aleatory uncertainty.

Example 1 Consider two experiments:

Experiment 1 Consider a fair dice that is to be thrown, and denote the outcome \mathbf{A} . Since the distribution of dice

throwing is known ($P(\mathbf{A} = i) = 1/6$ for $i=1, \dots, 6$), the uncertainty in the random variable \mathbf{A} is (purely) aleatory. *Experiment 2* Consider another dice that has been thrown, but where the dice has been covered so that the result is not visible. Now there is uncertainty about the value of the hidden dice. Call this random variable \mathbf{E} . The uncertainty in \mathbf{E} is (purely) epistemic because it could be reduced by gathering more information (removing the cover from the dice).

Assume the (random) quantity of interest is the sum, $\mathbf{S} = \mathbf{A} + \mathbf{E}$, of the result of the two die. If \mathbf{E} is given, then the remaining uncertainty is the aleatory (stochastic) uncertainty in throwing a dice. Without knowing the value of the hidden dice, the uncertainty in the sum \mathbf{S} is both aleatory and epistemic.

In Example 1 we had the option of uncovering the second dice, an *experiment* that would remove all epistemic uncertainty in \mathbf{S} . Generally, we will consider experiments that reduces (but not necessarily completely removes) epistemic uncertainty. For instance, in the context of Example 1, an experiment that would reveal whether \mathbf{E} was an even or odd number, or whether $\mathbf{E} > 1$. Or, the sum \mathbf{S} (but not the value of \mathbf{A}) from a few repeated throws of the aleatory die, from which inference about \mathbf{E} could be made. Reducing epistemic uncertainty usually comes at a cost, where the more informative experiments are more expensive. In this paper we are interested in how to decide on which experiments to perform, where the cost and potential effect of experiments are balanced in an optimal manner.

In the context of structural reliability modelling, the epistemic uncertainty usually comes from one of the two reasons:

1. The function $g(\cdot)$ or the distribution of \mathbf{X} depends on parameters that we do not know the value of.
2. Evaluating $g(\mathbf{x})$ at some single realization \mathbf{x} of \mathbf{X} is *expensive* in terms of money and/or time.

The second part comes from the complex physical nature of failure mechanisms, where experiments are needed to evaluate the function $g(\mathbf{x})$. This includes numerical computer simulations and physical experiments in a laboratory, which are both time consuming and expensive. Hence, due to the limited number of experiments that can be performed in practice, any method for estimating $P(g(\mathbf{X}) < 0)$ that relies on a large number of evaluations of $g(\cdot)$ is practically infeasible. This problem is usually solved by replacing the performance function $g(\cdot)$ with a computationally cheap *surrogate model* or *emulator*², constructed from a small set of experiments.

² The word *emulator* is often used for a surrogate model that can interpolate between noiseless observations coming from a deterministic computer simulation.

When the surrogate model is a stochastic process (viewed as a distribution over functions), we can quantify the added epistemic uncertainty that comes from this simplification.

We will assume that epistemic uncertainty is introduced in a structural reliability model, and that there is a way to reduce this uncertainty by performing experiments.

The problem we address in this paper is how to optimally estimate $P(g(\mathbf{X}) < 0)$ using as little resources as possible. In particular, we want to find an optimal strategy for the scenario where we can perform experiments sequentially, i.e. where each experiment may depend on the preceding ones.

Remark 1 (Why separate between epistemic and aleatory uncertainties?) Note that if there is no epistemic uncertainty in our model, then there is no incentive for performing experiments to collect more information, since the uncertainty cannot be reduced no matter what experiment we do. Hence, for our problem formulation to make sense, it is crucial to know that there is epistemic uncertainty present. In Example 1 we can consider the conditional probability $P(\mathbf{S}|\mathbf{E} = \mathbf{e})$, which for any fixed realization \mathbf{e} is a property of the aleatory uncertainty alone. When we do not know the value of \mathbf{E} , the quantity $P(\mathbf{S}|\mathbf{E})$ becomes a random variable of purely epistemic uncertainty. We are going to treat the failure probability $P(g(\mathbf{X}) < 0)$ in this way, by conditioning on epistemic information, so that we can study the potential effect of experiments.

Furthermore, Der Kiureghian and Ditlevsen (2009) show that by not separating between these two types of uncertainty in risk and reliability assessment, one may either over- or underestimate the failure probability by a significant magnitude (depending on the problem at hand), and conclude that distinguishing between aleatory and epistemic uncertainty in risk assessment is important. This is also supported by the examples we present in Sect. 6.

1.2 Hierarchical modelling

The scenario where $g(\cdot)$ is replaced by a surrogate model created from a finite set of observations $\{g(\mathbf{x}_i)\}_{i=1}^n$ has already been studied extensively (Bect et al. 2012; Echard et al. 2011; Bichon et al. 2008; Sun et al. 2017; Jian et al. 2017; Perrin 2016; Schueremans and Gemert 2005). The most common approach is to approximate $g(\cdot)$ using a Gaussian process, and make use of the convenient fact that a surrogate model given by the posterior predictive distribution of the Gaussian process has a closed form solution. However, structural reliability models are often hierarchical, and the reason why $g(\cdot)$ is expensive comes from one or more expensive *sub-components*³. An example is shown in Fig. 1, where

³ For instance, $g(\mathbf{x})$ is often a function of a structures capacity and the effect of loads acting on the structure, where each of which are determined from separate types of experiments.

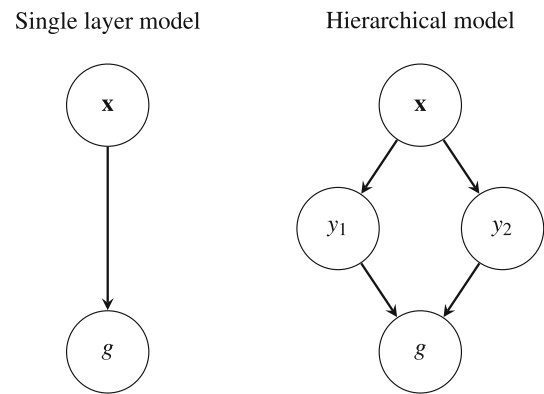


Fig. 1 Left: Single layer model. Right: Example of an hierarchical (2 layer) model where $g(\mathbf{x}) = g(y_1(\mathbf{x}), y_2(\mathbf{x}))$

$g(\mathbf{x}) = g(y_1(\mathbf{x}), y_2(\mathbf{x}))$. Assume here that $\mathbf{x} \in \mathbb{R}^m$, then the index set of the Gaussian process approximation of $g(\mathbf{x})$ is m -dimensional. Naturally, the number of experiments needed is highly dependent on m . If $g(\mathbf{x})$ is expensive, then this must be because one (or more) of the functions, $y_1(\mathbf{x})$, $y_2(\mathbf{x})$ or $g(y_1, y_2)$ is expensive. Very often, the effective domains⁴ of these functions have dimensionality much smaller than m , so fitting a Gaussian process to observations of $g(\mathbf{x})$ is not very efficient. There is also some practical inconvenience here, which is that some of the expensive sub-components (for instance load models) may be applicable in different structural reliability models, so there is a potential for re-use if we create a surrogate model for, say $y_1(\mathbf{x})$, instead of $g(\mathbf{x})$. Kzyurova et al. (2018) also consider a similar scenario and give some examples, for the 2-layer case where each component is replaced by a Gaussian process emulator.

In this paper we will work with hierarchical models (not necessarily with the structure illustrated in Fig. 1), where we assume that some of the intermediate variables are stochastic processes with epistemic (potentially reducible) uncertainty. Note that this also covers the case where we just introduce additional epistemic variables into the model. Actually, in the approximate numerical solution we propose in this paper, these two problems become equivalent. Moreover, as Gaussianity generally is lost in the hierarchical setting, we will only make assumptions on existence of second order moments of the stochastic processes used as surrogates. We will present a general formulation of the problem of finding an optimal strategy for performing experiments based on Bellman’s principle of optimality, and discuss some alternative routes for solving such problems. For the myopic (one step look-ahead) strategy, we propose an efficient numerical procedure, based on finite-dimensional approximation of the

⁴ If for instance $y_1(\mathbf{x}) : \mathbb{R}^m \rightarrow \mathbb{R}$ depends only on x_1, \dots, x_n for $n \leq m$, the effective domain of y_1 is n -dimensional.

stochastic processes and uncertainty propagation using the unscented transform.

1.3 Structure and main contributions of the paper

The structure of the remaining part of the paper is as follows: Through Sects. 2 and 3 we develop the Bayesian optimal experimental design problem for a general structural reliability model. We introduce a framework for separation of aleatory and epistemic uncertainties using conditional expectations, from which we can express any type of experiment associated with a structural reliability problem. For the purpose of estimating a failure probability, we consider three alternative optimization objectives, and in Sect. 3 we discuss how the experimental design problem may be tackled using dynamic programming and the one-step lookahead approximation. Optimization problems of this form will involve evaluation of a *measure of residual uncertainty*, and in Sect. 4 we present an approach for approximating this quantity. We implement this in Sect. 5 to develop an efficient numerical procedure for the one-step lookahead case, which we illustrate through a series of examples in Sect. 6. Finally, our concluding remarks are given in Sect. 7, and some supporting material used throughout the paper is included in the Appendices.

2 Problem formulation

Given a probabilistic surrogate of a structural reliability model, we are interested in how to optimally improve the model for failure probability estimation, given a fixed experimental budget. More generally, given a structural reliability model with epistemic uncertainty (e.g. as introduced when using a surrogate), and a set of possible experiments that can be performed, we want to select the experiments in an optimal manner. The choice of experiment is called a *decision*, $d \in \mathbb{D}$ where \mathbb{D} is a space of feasible decisions. Note that this set may include different kinds of decisions, such as performing computer experiments, lab experiments or performing physical measurements in the field.

In the following subsections we present a rigorous formulation of the Bayesian optimal experimental design problem for structural reliability analysis. Here we will need a way to express uncertainty about the *performance function* used in structural reliability models, and a way to model uncertainty about future outcomes of potential experiments that can be made. For this purpose we will define a model (ξ, δ) , where

- ξ is a stochastic representation of the *performance function* $g(\mathbf{x})$ evaluated at some fixed input \mathbf{x} .

- $\delta(d)$ is a predictive model of experimental outcomes given a decision d . In other words, δ models the data generating process of potential experiments.

We will consistently write \mathbf{X} as a random variable with values in $\mathbb{X} \subseteq \mathbb{R}^m$, and let \mathbf{x} be a deterministic realization. ξ and δ are stochastic processes, indexed over inputs \mathbf{x} and decisions d , respectively. In structural reliability analysis, we are interested in the random variable $g(\mathbf{X})$, and likewise we will consider $\xi(\mathbf{X})$, but now where $\xi(\mathbf{x})$ is also random for any fixed \mathbf{x} . Here, for notational convenience, we suppress the $\omega \in \Omega$ when referring to the random variable $\mathbf{X} : \Omega \rightarrow \mathbb{X}$ or the stochastic process $\xi(\mathbf{x}) : \mathbb{X} \times \Omega \rightarrow \mathbb{R}$. That is, we define the notation $\xi(\mathbf{X}) := \xi(\mathbf{X}(\omega), \omega)$ to describe the random variable $\xi(\mathbf{X}) : \Omega \rightarrow \mathbb{R}$.

Remark 2 Note that $\xi(\mathbf{x})$ is a stochastic representation of the performance function $g(\cdot)$. When making decisions d , we aim to reduce the uncertainty in $\xi(\mathbf{X})$, where also the input \mathbf{X} is random. Hence, the process $\delta(d)$ is linked with $\xi(\mathbf{X})$ through its reduction of uncertainty (see Sect. 2.3).

As the purpose of performing experiments will be to provide information about ξ , note that ξ and δ are generally not independent. A detailed description of how (ξ, δ) is constructed is provided in the following subsections.

2.1 Structural reliability analysis

Let $\mathbb{X} \subseteq \mathbb{R}^m$, and let \mathbf{X} be a random variable on the probability space (Ω, \mathcal{F}, P) with values in \mathbb{X} and $g : \mathbb{X} \rightarrow \mathbb{R}$ a measurable function. We call g the *performance function* or *limit state*, with the associated failure set

$$F_g = \{\mathbf{x} \in \mathbb{X} \mid g(\mathbf{x}) \leq 0\}.$$

In structural reliability analysis, we are interested in estimating the *failure probability*, which we here denote $\bar{\alpha}$. It is defined as

$$\bar{\alpha}(g) = P(F_g) = E[\mathbf{1}(g(\mathbf{X}) \leq 0)], \quad (1)$$

where $E[\cdot]$ denotes the expectation with respect to P and $\mathbf{1}(\cdot)$ is the indicator function.

In most real-world cases it is difficult to derive an analytical expression for the failure probability. To overcome this, several approximation and simulation methods have been suggested, see e.g. Madsen et al. (2006) or Huang et al. (2017). Two traditional methods are the first- and second-order reliability method (FORM/SORM), where the failure boundary is approximated at a specific point using a Taylor expansion up to the first and second order, respectively. Different sampling procedures have also been developed, which often make use of intermediate results obtained from

FORM/SORM. Other relevant techniques involve the construction of environmental contours and the estimation of buffered failure probabilities as in (Dahl and Huseby 2019). In this paper, our focus is different from these methods in the sense that we are mainly interested in how to estimate the failure probability as well as possible, given a limited experimental budget. To do so, we need to separate between different kinds of uncertainty in our model.

2.2 Separating epistemic and aleatory uncertainties

Ideally, the uncertainty related to the random variable $g(\mathbf{X})$ in (1) is *aleatory*, in the sense that that it relates to inherent variability of the physical phenomenon that is being modelled, but in reality we must also include *epistemic* uncertainty due to lack of information or knowledge. For instance, assume that $g(\mathbf{x}, \mathbf{e})$ depends on the aleatory variable \mathbf{x} and some fixed but unknown parameter \mathbf{e} . Assume further that \mathbf{X} is the aleatory random variable representing variability in \mathbf{x} , \mathbf{E} is the epistemic random variable representing our belief about \mathbf{e} , and that \mathbf{X} and \mathbf{E} are independent with laws P_x and P_e . It is then relevant to view the failure probability as a random quantity with epistemic uncertainty, $\alpha(\mathbf{E}) = \int \mathbf{1}(g(\mathbf{x}, \mathbf{E}) \leq 0) P_x(d\mathbf{x})$. For engineering applications, one would then typically be interested in some specified upper percentile values of $\alpha(\mathbf{E})$, i.e. ensuring that the epistemic uncertainty is under control.

In the following, we will assume that we have a performance function $\xi(\cdot)$ that depends on a strictly aleatory random variable \mathbf{X} , and some other random quantity with epistemic uncertainty. We will need to formulate this with a bit of generality, in order to cover the different ways epistemic uncertainty can be introduced in a structural reliability model.

As in Sect. 2.1 we will work with (Ω, \mathcal{F}, P) as the global probability space, capturing all forms of uncertainty. We then let \mathcal{A} and \mathcal{E} be two sub σ -algebras representing, respectively, aleatory and epistemic information i.e.,

$$\mathcal{A}, \mathcal{E} \subseteq \mathcal{F}. \tag{2}$$

Though all uncertainty in our model is assumed to be either epistemic or aleatory, Example 1 illustrates that random variables may contain both aleatory and epistemic information.

We will assume that \mathbf{X} is \mathcal{A} -measurable. Furthermore, for any $\mathbf{x} \in \mathbb{X}$ we assume that $\xi(\mathbf{x})$ is \mathcal{E} -measurable. That is, $\xi : \mathbb{X} \times \Omega \rightarrow \mathbb{R}$ is a stochastic process indexed by $\mathbf{x} \in \mathbb{X}$ (this is also called a random field), and $\xi(\mathbf{X})$ is a real-valued random variable. We will write $\xi(\cdot)$ instead of $g(\cdot)$ whenever epistemic uncertainty has been introduced, as for instance in the canonical case where a deterministic performance function $g(\cdot)$ is approximated with a probabilistic surrogate $\xi(\cdot)$.

We can now define the failure probability with epistemic uncertainty as the \mathcal{E} -measurable random variable

$$\alpha(\xi) = E [\mathbf{1}(\xi(\mathbf{X}) \leq 0) | \mathcal{E}]. \tag{3}$$

Note that (3) coincides with (1) in the case where the performance function is not affected by epistemic uncertainty, and in general as $\bar{\alpha}(\xi) = E [\alpha(\xi)]$ because

$$\begin{aligned} E [\alpha(\xi)] &= E [E [\mathbf{1}(\xi(\mathbf{X}) \leq 0) | \mathcal{E}]] \\ &= E [\mathbf{1}(\xi(\mathbf{X}) \leq 0)] \\ &= \bar{\alpha}(\xi), \end{aligned} \tag{4}$$

where the second equality uses the double expectation property.

In the following we will just write α or $\bar{\alpha}$ without the dependency on ξ when there is no risk of confusion.

Example 2 Assume ξ is a deterministic function of the aleatory random variable \mathbf{X} and epistemic random variable \mathbf{E} , both defined on (Ω, \mathcal{F}, P) . Then $\mathcal{A} = \sigma(\mathbf{X})$ and $\mathcal{E} = \sigma(\mathbf{E})$, i.e., the σ -algebras generated by the random variables \mathbf{X} and \mathbf{E} , respectively.

Note that the converse of Example 2 also holds true, as we can always view ξ as a deterministic function applied to two random variables \mathbf{X} and \mathbf{E} . That is, where $\xi(\mathbf{x}, \mathbf{e})$ is a deterministic function for \mathbf{x} and \mathbf{e} fixed, and we can write the stochastic process $\xi(\mathbf{x}, \omega)$ as $\xi(\mathbf{x}, \mathbf{E})$. It is sometimes useful to think of ξ in this way. In particular, the numerical approximation we propose later in this paper is based on obtaining a finite-dimensional approximation of \mathbf{E} .

Example 3 Let g be given as in the hierarchical model in Fig. 1, and \mathbf{X} a random variable defined on some measure space $(\Omega_x, \mathcal{F}_x, P_x)$. Assume that y_1 and y_2 are expensive to evaluate, so we replace them with surrogate models in the form of two stochastic processes \tilde{y}_1 and \tilde{y}_2 defined on another measure space $(\Omega_y, \mathcal{F}_y, P_y)$. Note that we assume that both \tilde{y}_1 and \tilde{y}_2 are defined on the same measure space. Then, the measure space for the experimental design problem is given by $(\Omega, \mathcal{F}, P) = (\Omega_x \times \Omega_y, \mathcal{F}_x \otimes \mathcal{F}_y, P_x \times P_y)$, $\mathcal{A} = \mathcal{F}_x$ and $\mathcal{E} = \mathcal{F}_y$ (up to isomorphism), and we would write $\xi(\mathbf{x}) = g(\tilde{y}_1(\mathbf{x}), \tilde{y}_2(\mathbf{x}))$.

2.3 Decisions, outcomes and experiments

We are interested in the case where the epistemic uncertainty in α can be reduced by running experiments. For instance, in Example 2 the epistemic variable \mathbf{E} could be a fixed but unknown parameter, and maybe additional measurements could be performed to reduce the uncertainty in \mathbf{E} . Or in Example 3, additional experiments could be performed to infer the values of y_1 or y_2 at some given input \mathbf{x}' , in order to reduce uncertainty in the surrogate models \tilde{y}_1 and \tilde{y}_2 .

These are examples of possible *decisions* we could make to reduce epistemic uncertainty. We will let \mathbb{D} denote the set of all possible decisions, and \mathbb{O} the set of all possible *outcomes*. For any decision $d \in \mathbb{D}$, the corresponding outcome is uncertain a priori, and in order to evaluate the potential impact of a decision we will need to specify (possibly subjectively) a distribution representing the possible outcomes. We will let $\delta(d)$ denote the random outcome of a decision $d \in \mathbb{D}$ with values in \mathbb{O} . For any realization $o \in \mathbb{O}$ of $\delta(d)$, we will refer to the pair (d, o) as an *experiment*.

In our modelling framework, we will assume that $\xi(\mathbf{x})$ as defined in Sect. 2.2 is provided together with (Ω, \mathcal{F}, P) and the sub σ -algebras \mathcal{A} and \mathcal{E} , and that a decision process $\delta(d)$ is given where $\delta(d)$ is \mathcal{E} -measurable for any $d \in \mathbb{D}$. Table 1 gives an overview of the notation we have introduced so far, in order to define the problem of optimal experimental design for structural reliability analysis.

Example 4 Continuing from Example 3, assume that noise perturbed observations of y_1 can be made. Let $d(\mathbf{x}) = \{\text{observe } y_1(\mathbf{x})\}$, and define \mathbb{D} as the union of such events for all \mathbf{x} . If we assume that observations come with additive noise, $o(\mathbf{x}) = y_1(\mathbf{x}) + \epsilon(\mathbf{x})$, for some specified noise process ϵ , then we can let $\delta(d(\mathbf{x})) = \tilde{y}_1(\mathbf{x}) + \epsilon(\mathbf{x})$. In a similar fashion, \mathbb{D} and $\delta(d)$ could be extended to include observations of y_2 as well.

We will note that the noise-free alternative to Example 4, i.e. the case where $\epsilon \equiv 0$, is a common scenario when dealing with deterministic computer simulations. Another related scenario that is also of relevance here, is that of multifidelity modelling (Fernandez et al. 2017), in which case inaccurate estimates of $y_1(\mathbf{x})$ could be available at the same time, but at a lower cost.

2.4 Sequential model updating

Now, having defined a random variable \mathbf{X} and the two processes $\{\xi(\mathbf{x})\}_{\mathbf{x} \in \mathbb{X}}$ and $\{\delta(d)\}_{d \in \mathbb{D}}$, we want to perform a sequence of experiments, $(d_0, o_0), (d_1, o_1), \dots$, and update ξ and δ accordingly.

We let $I_k := \{(d_0, o_0), \dots, (d_{k-1}, o_{k-1})\}$ denote the information or history up to the k th experiment, and define \mathcal{E}_k as the σ -algebra generated by \mathcal{E} and I_k . Hence, \mathcal{E}_k is all the information regarding epistemic quantities that is available after k experiments. We introduce the notation $P_k(\cdot)$ and $E_k[\cdot]$ to denote the conditional distribution $P(\cdot | \mathcal{E}_k)$ and conditional expectation $E[\cdot | \mathcal{E}_k]$ given the updated information \mathcal{E}_k . For convenience we define $I_0 = \emptyset$, so that we can use the index $k = 0$ with these definitions for the scenario before any experiment has been made. We will write ξ_k and δ_k as the updated processes $\xi|I_k$ and $\delta|I_k$ corresponding to P_k . Per definition,

$$(\xi_{k+1}, \delta_{k+1}) = (\xi_k, \delta_k) | d_k, o_k = (\xi_0, \delta_0) | I_k, d_k, o_k.$$

In the following example, we show how this sequential update can be done via Bayes' theorem.

Example 5 Let $k \in \mathbb{N}$. Assume (ξ, δ) admits a joint probability density at any finite subset of $\mathbb{X} \times \mathbb{D}$ with respect to P_k , which we write $p_k(\xi, \delta)$ for short. E.g. $p_k(\xi)$ means

$$P_k \left(\left(\xi(\mathbf{x}^{(1)}), \dots, \xi(\mathbf{x}^{(n)}) \right) = \left(\xi^{(1)}, \dots, \xi^{(n)} \right) \right)$$

for some $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)} \in \mathbb{X}$ and $\xi^{(1)}, \dots, \xi^{(n)} \in \mathbb{R}$. Then $p_k(\xi) = p_0(\xi_k)$, $p_k(\delta) = p_0(\delta_k)$, and the update of the probabilities is done by using Bayes' theorem:

$$p_{k+1}(\xi) = p_k(\xi|d_k, o_k) = \frac{p_k(o_k|\xi, d_k)p_k(\xi)}{p_k(o_k|d_k)},$$

$$p_{k+1}(\delta) = p_k(\delta|d_k, o_k) = \frac{p_k(o_k|d_k, \delta)p_k(\delta)}{p_k(o_k|d_k)},$$
(5)

where $p_k(\cdot)$ is the relevant density with respect to P_k .

Example 6 For a specific problem there will typically be simpler ways of updating the model than the generic formulation given in the previous example. Continuing again from Examples 3 and 4, assume $\delta(d) = \delta(\mathbf{x}, \tilde{y}_1, \tilde{y}_2)$ corresponds to observing $\tilde{y}_1(\mathbf{x}) + \epsilon_1(\mathbf{x})$ or $\tilde{y}_2(\mathbf{x}) + \epsilon_2(\mathbf{x})$. Then \tilde{y}_1 and \tilde{y}_2 can be updated directly, and we let $\xi|I_k = g(\tilde{y}_1|I_k, \tilde{y}_2|I_k)$ and $\delta|I_k = \delta(\mathbf{x}, \tilde{y}_1|I_k, \tilde{y}_2|I_k)$.

In fact, if \tilde{y}_1 and \tilde{y}_2 and the noise terms ϵ_1 and ϵ_2 are all Gaussian processes, then $\tilde{y}_1|I_k$ and $\tilde{y}_2|I_k$ are also Gaussian and closed form representations are available (see Appendix A). Note that in this case the model update could include updating the Gaussian process hyperparameters as well.

2.5 Optimization objective

Following the formulation of Bect et al. (2012, 2019), a strategy for uncertainty reduction starts with a *measure of residual uncertainty* for the quantity of interest after k experiments. This is a functional

$$H_k = \mathcal{H}(P_k)$$
(6)

of the conditional distribution P_k . In this paper we will consider three specific alternatives for H_k .

Assume k experiments have been performed, resulting in the updated probabilistic model (ξ_k, δ_k) . The updated failure probability according to (3) can then be defined as

$$\alpha_k = \alpha(\xi_k) = E_k[\mathbf{1}(\xi(\mathbf{X}) \leq 0)], \quad \bar{\alpha}_k = E[\alpha_k].$$
(7)

Table 1 Overview of the framework for the optimal experimental design problem for structural reliability analysis

Symbol	Description	Type
\mathbf{X}	Parameters describing structure and environment	\mathbb{R}^m -valued random variable
\mathbf{x}	Deterministic realization of \mathbf{X}	values in \mathbb{R}^m
$g(\mathbf{x})$	Performance function of structure	real-valued
F_g	Failure set of $g(\cdot)$	subset of \mathbb{R}^m
$\bar{\alpha}(g)$	The failure probability, $P(F_g)$	values in $[0, 1]$
$\xi(\mathbf{x})$	Stochastic approximation of $g(\cdot)$	real-valued stochastic process
$\alpha(\xi)$	The failure probability with epistemic uncertainty	values in $[0, 1]$
d	Decision	contained in set of decisions \mathbb{D}
o	Outcome of experiment	contained in set of outcomes \mathbb{O}
(d, o)	Summary of an experiment	contained in $\mathbb{D} \times \mathbb{O}$
$\delta(d)$	Model of experiment outcomes	\mathbb{O} -valued stochastic process
\mathbf{E}	Parameters for epistemic uncertainty, independent of \mathbf{X}	random variable
\mathcal{A}	Aleatory information	σ -algebra
\mathcal{E}	Epistemic information	σ -algebra
(ξ, δ)	The model	$(\mathbb{R} \times \mathbb{O})$ -valued

As we are interested in reducing uncertainty in α , a natural optimization objective is to minimize $\text{Var}(\alpha_k) = E[(\alpha_k - \bar{\alpha}_k)^2]$. However, computation of $\text{Var}(\alpha_k)$ can be problematic in practice. Most of the proposed methods for design of experiments in (non-hierarchical) structural reliability models therefore make use of alternative *heuristic* optimization objectives. That is, some alternative function $H_k(\cdot)$ that is easier to compute than $\text{Var}(\alpha_k)$, and where the design that minimizes $H_k(\cdot)$ hopefully also performs well with respect to $\text{Var}(\alpha_k)$.

Bect et al. (2012) present a few such criteria, some of which will also be considered in this paper. Let

$$\begin{aligned} p_k(\mathbf{X}) &= P_k(\xi(\mathbf{X}) \leq 0), \\ \gamma_k(\mathbf{X}) &= p_k(\mathbf{X})(1 - p_k(\mathbf{X})). \end{aligned} \tag{8}$$

Observe that

$$\begin{aligned} \text{Var}(\mathbf{1}(\xi_k(\mathbf{x}) \leq 0)) &= E[(\mathbf{1}(\xi_k(\mathbf{x}) \leq 0))^2] - \\ &= E[\mathbf{1}(\xi_k(\mathbf{x}) \leq 0)]^2 \\ &= E[(\mathbf{1}(\xi_k(\mathbf{x}) \leq 0))] - p_k(\mathbf{x})^2 \\ &= p_k(\mathbf{x}) - p_k(\mathbf{x})^2 \\ &= \gamma_k(\mathbf{x}), \end{aligned} \tag{9}$$

and also that $\gamma_k(\mathbf{x})/2$ is the probability that two i.i.d. samples from $\xi_k(\mathbf{x})$ have the same sign. Hence, γ_k provides a measure of how accurate $\xi_k(\mathbf{x})$ is around the critical value $\xi_k = 0$. We will introduce two measures of residual uncertainty based on taking the expectation of γ_k with respect the distribution of \mathbf{X} , which we denote $P_{\mathbb{X}}$. In total, we will consider the following three alternatives for H_k :

$$\begin{aligned} H_{1,k} &= E_k[(\alpha - \bar{\alpha})^2], \\ H_{2,k} &= \int_{\mathbb{X}} \gamma_k dP_{\mathbb{X}} = E[\gamma_k], \\ H_{3,k} &= \left(\int_{\mathbb{X}} \sqrt{\gamma_k} dP_{\mathbb{X}} \right)^2 = E[\sqrt{\gamma_k}]^2. \end{aligned} \tag{10}$$

Here $H_{2,k}$ and $H_{3,k}$ can also be motivated by realizing that they serve as upper bounds on $H_{1,k}$. In fact, $H_{1,k} \leq H_{3,k} \leq H_{2,k}$ (see Proposition 3 in Bect et al. 2012).

For optimal design of experiments we will consider loss functions given by the above measures of residual uncertainty, potentially in combination with an additional penalty term that represents the cost of performing a given experiment. In the Bayesian decision-theoretic framework, given such a loss function depending on a *policy* for selecting experiments π , we can evaluate the policy by looking n -steps ahead. For instance, a relevant loss function for minimizing uncertainty in α after n additional experiments, following after the current experiment k , could be given as $J_k(\pi) = E_k[H_{1,k+n}]$ where \mathcal{E}_{k+n} corresponds to following the policy π . The additional notation introduced with respect to the measure of residual uncertainty and sequential model updating is summarized in Table 2.

3 Modelling information and experimental design

In this section, we introduce the experimental design framework and explain how the development of information is modelled in this context. In the following, let $k =$

Table 2 Overview of the framework for the optimal experimental design problem for structural reliability analysis with sequential model updating

Symbol	Description	Type
I_k	Information up to k 'th experiment	Sequence of decisions and outcomes
\mathcal{E}_k	Information given ξ and I_k	σ -algebra
P_k	Conditional probability given ξ_k	Values in $[0, 1]$
ξ_k	Update of ξ given I_k	Stochastic process indexed by k
δ_k	Update of δ given I_k	Stochastic process indexed by k
$H_k = \mathcal{H}(P_k)$	Measure of residual uncertainty	Functional from space of probability distributions to \mathbb{R}
α_k	Updated epistemic failure probability	Values in $[0, 1]$
$\tilde{\alpha}_k$	Updated expected failure probability	Values in $[0, 1]$

$0, 1, \dots, K - 1$ be the *experiment index* which keeps track of the number of performed experiments.

3.1 The dynamic programming formulation

Huan and Marzouk (2016) introduce a general framework for sequential optimal experimental design: Let the *state*⁵ of the system after experiment $k - 1$ be denoted by s_k . The input (decided by the experimental designers) to experiment k is denoted by d_k . We want to determine a *policy*

$$\pi := (\pi_0, \pi_1, \dots, \pi_{K-1})$$

where $d_k = \pi_k(s_k)$. That is, given the current state of the system, the policy is a function which tells the experimental designer the input to the next experiment.

From each experiment, we get *observations* o_k . These observations may include measurement noise and modelling errors. Associated to each experiment, we have a *stage reward* $R_k(s_k, o_k, d_k)$. The stage reward reflects the cost of doing the experiment (measured in e.g. money or time) plus any additional benefits or penalties of doing the experiment (measured in the same unit). Furthermore, we have a *terminal reward* $R_K(s_K)$ only depending on the final state of the system.

In order to model the development of the system of experiments, we have the *system dynamics*:

$$s_{k+1} = \mathcal{V}(s_k, d_k, o_k)$$

where $\mathcal{V}(\cdot)$ is some function specifying the transition from a current state to a new state based on the performed experiment. The optimal experimental design problem can then be formulated as follows:

⁵ In (Huan and Marzouk 2016) the *state* is written as $s_k = (s_k^{(b)}, s_k^{(p)})$, where $s_k^{(b)}$ denotes the *uncertainty state* and $s_k^{(p)}$ denotes the *physical state* that describes any additional deterministic decision-relevant variables. Herein we will not write s_k specifically in this form.

Maximize

$$E_{o_0, \dots, o_{K-1}} \left[\sum_{k=0}^{K-1} R_k(s_k, o_k, \pi_k(s_k)) + R_K(s_K) \right] \tag{11}$$

such that $s_{k+1} = \mathcal{V}(s_k, d_k, o_k)$,

and the maximization is done over all policies π that do not look into the future (in the sense that information about future results of experiments are used in current policy making). That is, when deciding policy π_k , only what is known up to experiment $k - 1$ can be used. Another way of saying this is that the policy π should be *adapted* to the filtration generated by the processes $\{s_k\}$, $\{o_k\}$ and $\{d_k\}$.

To adapt this framework to the experimental design problem for structural reliability analysis, we write

$$s_k = (\xi_k, \delta_k, I_k), \quad d_k = \pi_k(s_k), \quad o_k = \delta_k(d_k), \tag{12}$$

and where the dynamics $s_{k+1} = \mathcal{V}(s_k, d_k, o_k)$ is given by updating ξ_k, δ_k and I_k with respect to the experiment (d_k, o_k) as described in Sect. 2.4.

Remark 3 Note that the expectation in (11) is with respect to future outcomes o_0, \dots, o_{K-1} which a priori are uncertain, and where each outcome o_k depends on the previous outcomes o_0, \dots, o_{k-1} . An equivalent formulation can be given in terms of conditional expectations. Let each reward be defined by backwards induction:

$$R_k = \max_d E_k [R_{k+1} | d_k = d], \quad k = K - 1, \dots, 0,$$

where $R_K = R_K(s_K)$ only depends on the final state of the system. Then, the policy defined by selecting for each k the decision

$$\begin{aligned} d_k^* &= \arg \max_{d \in \mathbb{D}} E_k [R_{k+1} | d_k = d] \\ &= \arg \max_{d \in \mathbb{D}} E_k [\max_{d \in \mathbb{D}} E_{k+1} \dots E_K R_K | d_k = d] \end{aligned}$$

is optimal. This corresponds with the formulation used by Bect et al. (2012).

Problem (11) is a dynamic programming problem. Though theoretically optimal, such problems are known for suffering from the so-called *curse of dimensionality*. That is, possible sequences of design and observation realizations grow exponentially with the dimension of the state space. According to Defourny et al. (2011), the curse of dimensionality implies that dynamic programming can only be solved numerically for state spaces embedded in \mathbb{R}^d with $d \leq 10$. Therefore, such problems can often only be solved approximately via *approximate dynamic programming*, see (Huan and Marzouk 2016). Note also that this type of formulation is based on a *Markovianity* assumption, i.e., that there is no memory in the dynamics of the system. This assumption is necessary in order to perform the simplification to only having dependency on the current state of the system in Remark 3. If the system is not Markovian, in the sense that the decision at any time depends not only on the current state of the system, but also on some of the previous states, we cannot solve the experimental design problem by backwards induction. The reason for this is that the Bellman equation, which backwards induction is based on, does not hold in this case. In such cases, the experimental design problem can for instance be solved via the maximum principle, see e.g. Dahl et al. (2016) for an example of systems with memory in continuous time.

Remark 4 An alternative solution method to dynamic programming for problem (11) is to use a scenario tree based approach, see Defourny et al. (2011). Scenario tree based approaches are not sensitive to curse of dimensionality based on the state space, but based on the number of experiments. Hence, a scenario based approach can be attempted whenever there are few experiments (less than or equal 10), but potentially a large dimensional state space. If the number of experiments is large (greater than 10), but the state space dimension is small (less than or equal 10), dynamic programming is a viable solution method. If both the state space dimension and the number of experiments is large, one can try approximate dynamic programming (see Huan and Marzouk (2016)) or a one-step lookahead (myopic)⁶ formulation as an alternative to the dynamic programming one. In Sect. 3.2, we consider such a one-step lookahead formulation.

Note that problem (11) is maximization problem of a reward, but can trivially be transformed to a minimization problem with some loss function $L_k = -R_k$ instead. For the application considered in this paper, we are interested in minimization problems associated with the residual uncertainty described in Sect. 2.5.

⁶ Some authors, for instance Huan and Marzouk (2016) and Bect et al. (2012), remark that all strategies which consider fewer terms in the summation are *myopic*. Other authors use *myopic* only in the case where no future decisions are taken into account, i.e. the horizon is zero. When we say myopic in this paper, we mean one-step lookahead.

Example 7 Let $\lambda(d_k)$ denote the cost of decision d_k . A relevant set of loss functions could then be: $L_k(s_k, d_k, o_k) = 0$ for $k < K$ and $L_K = H_K \cdot \sum_{k < K} \lambda(d_k)$, where $H_K = H_{1,k}, H_{2,k}$ or $H_{3,k}$ as described in Sect. 2.5. Or, letting $L_k(s_k, d_k, o_k) = \eta^k \lambda(d_k) H_k$ for $k < K$ where η is some discount factor, $\eta \in (0, 1)$, would produce a similar but more greedy policy. Another relevant alternative is to define $L_K = \sum_{k < k^*} \lambda(d_k)$ as the sum of costs up to the iteration k^* where some target level, $H_k < H^*$ for $k > k^*$, has been reached.

3.2 The one-step lookahead formulation

As mentioned in Sect. 3.1, the dynamic programming formulation suffers from the curse of dimensionality. An approximation to the dynamic programming formulation which mends this problem, is the *myopic formulation* or *one-step lookahead*. This corresponds to truncating the dynamic programming sum in (11) and only looking at one time-step ahead.

In this section, we define the the *one-step lookahead optimal decision* $d \in \mathbb{D}$ at step k as the minimizer of the following function

$$J_{i,k}(d) = \lambda(d) E_{k,d} [H_{i,k+1}] \text{ for } i = 1, 2, \text{ or } 3. \tag{13}$$

Here $H_{i,k}$ are the measures of residual uncertainty defined in Sect. 2.5, and $E_{k,d}$ represents the conditional expectation with respect to \mathcal{E}_k with $d_k = d$. Hence, $E_{k,d} [H_{i,k+1}]$ represents how desirable decision d is for reducing the expected remaining uncertainty in α at experiment $k + 1$, if the next experiment is performed with input d . We let $\lambda(d)$ be a deterministic function representing the cost associated with decision d , and we will refer to a function $J_{i,k}(d)$ as the *acquisition* function for myopic (one-step lookahead) design. Other ways of introducing additional rewards or penalties associated with an experiment are of course also possible. In fact, there is no particular reason why we write (13) as a *product* of cost and the measure of residual uncertainty, besides emphasizing that $J_{i,k}(d)$ should be a function of these two terms.

Remark 5 We have assumed here that a total number K of experiments that are to be performed, where we want to perform each experiment optimally. But in practice it is relevant to consider stopping before the K th experiment, when some objective has been reached, or when the potential gain of new experiments diminishes. Section 5.3 we introduce a criterion for stopping when the variance in the failure probability is sufficiently low.

4 Approximating the measure of residual uncertainty

Assume k experiments have been performed, resulting in the updated probabilistic model (ξ_k, δ_k) . A simple method for estimating the measures of residual uncertainty described in Sect. 2.5, is by a double-loop Monte Carlo simulation: Let $N_1, N_2 \in \mathbb{N}$ and let $h_{i,j}^{(k)} = \mathbf{1}(\xi_{k,j}(\mathbf{x}_i) \leq 0)$, where $\mathbf{x}_1, \dots, \mathbf{x}_{N_1}$ are N_1 i.i.d. samples of \mathbf{X} and $\xi_{k,1}(\mathbf{x}_i), \dots, \xi_{k,N_2}(\mathbf{x}_i)$ are N_2 i.i.d. performance functions sampled from ξ_k and evaluated at each \mathbf{x}_i . Then $H_{1,k}$ can be obtained as the sample variance of the N_2 samples of the form $\hat{\alpha}_{k,j} = \frac{1}{N_1} \sum_i h_{i,j}^{(k)}$. Similarly, $H_{2,k}$ and $H_{3,k}$ can be estimated from $\hat{\rho}_k(\mathbf{x}_i) = \frac{1}{N_2} \sum_j h_{i,j}^{(k)}$.

This approach is problematic for several reasons. First of all, $\hat{\alpha}_{k,j}$ is an unbiased estimator of the failure probability $\alpha_{k,j} = \alpha(\xi_{k,j})$ corresponding to the deterministic performance function $\xi_{k,j}$. When $\alpha_{k,j}$ is small, the variance of this estimator is $\text{var}(\hat{\alpha}_{k,j}) = \alpha_{k,j}(1 - \alpha_{k,j})/N_1 \approx \alpha_{k,j}/N_1$. If we want to achieve an accuracy, of say $\sqrt{\text{var}(\hat{\alpha}_{k,j})} < 0.1\alpha_{k,j}$, and $\alpha_{k,j} = 10^{-m}$, then the number of samples required would be approximately $N_1 = 10^{m+2}$. The failure probabilities considered in structural reliability analysis can typically be in the range from 10^{-6} to 10^{-2} .

When N_1 is large, it can also be a practical challenge to obtain the samples $\xi_{k,j}(\mathbf{x}_1), \dots, \xi_{k,j}(\mathbf{x}_{N_1})$ simultaneously for a fixed j . Moreover, the total number of samples needed to evaluate the measures of residual uncertainty $H_{i,k}$ is $N_1 N_2$, and we are interested in optimization over $H_{i,k}$ that will require multiple simulations of this kind.

In this section we present a procedure for efficient approximation of the measures of residual uncertainty. We will start by introducing a finite-dimensional approximation of $\xi_k(\mathbf{x})$, given as a deterministic function $\hat{\xi}_k(\mathbf{x}, \mathbf{E})$ depending on \mathbf{x} and a finite-dimensional \mathcal{E}_k -measurable random variable \mathbf{E} . Then, in Sect. 4.2 we consider how the mean and variance, $E[f(\mathbf{E})]$ and $\text{var}(f(\mathbf{E}))$, can be approximated for any \mathcal{E}_k -measurable function $f(\mathbf{e})$ using the unscented transform. In Sects. 4.3 and 4.4 we present an importance sampling scheme for the case where $f(\mathbf{e})$ is defined in terms of an expectation over \mathbf{X} . Finally, in Sect. 4.5 we consider the case where $f(\mathbf{e}) = \alpha(\hat{\xi}_k(\mathbf{X}, \mathbf{e}))$, which provides the approximations $\hat{\alpha}_k = f(\mathbf{E})$ and $\hat{H}_{1,k} = \text{var}(f(\mathbf{E}))$, and where approximations of $H_{2,k}$ and $H_{3,k}$ are obtained in a similar manner.

In summary, this kind of approximation which we will refer to as UT-MCIS from now on, makes use of the unscented transform (UT) for epistemic uncertainty propagation and Monte Carlo simulation with importance sampling (MCIS) for aleatory uncertainty propagation. The motivation behind this specific setup is that a technique such as MCIS is needed to obtain low variance estimates of $\alpha(\hat{\xi}_k(\mathbf{X}, \mathbf{e}))$, which will typically be a small number. The sampling scheme

we propose is also designed to be efficient in the case where subsequent estimates corresponding to perturbations of $\alpha(\hat{\xi}_k(\mathbf{X}, \mathbf{e}))$ are needed, which is relevant for estimation of e.g. $\alpha(\hat{\xi}_{k+1}(\mathbf{X}, \mathbf{e}))$ or $\alpha(\hat{\xi}_k(\mathbf{X}, \mathbf{e}'))$ for some $\mathbf{e}' \neq \mathbf{e}$ if $\alpha(\hat{\xi}_k(\mathbf{X}, \mathbf{e}))$ has already been estimated. As for epistemic uncertainty propagation, when $\alpha(\hat{\xi}_k(\mathbf{x}, \mathbf{E}))$ is viewed as an \mathcal{E}_k -measurable random variable, the UT alternative which is both simpler and more efficient seems like a viable alternative, in particular for the purpose of optimization with respect to future decisions.

4.1 The finite-dimensional approximation of ξ_k

In our framework, we have defined ξ_k as a \mathcal{E}_k -measurable stochastic process indexed by $\mathbf{x} \in \mathbb{X}$ (often called a random field), and we view ξ_k as a distribution over some (generally infinite-dimensional) space of functions. The special case where $\xi_k = \xi_k(\mathbf{x}, \mathbf{E})$ for some finite-dimensional \mathcal{E}_k -measurable random variable \mathbf{E} can be very useful for simulation. That is, if samples \mathbf{e}_j of \mathbf{E} can be generated efficiently, then random functions $\xi_{k,j}(\mathbf{x}) = \xi_k(\mathbf{x}, \mathbf{e}_j)$ can be sampled as well. As long as ξ_k is square integrable, such a representation of ξ_k is always available from the Karhunen-Loève transform:

$$\xi_k(\mathbf{x}) - E[\xi_k(\mathbf{x})] = \sum_{i=1}^{\infty} E_i \phi_i(\mathbf{x}),$$

where the functions ϕ_i are deterministic and E_i are uncorrelated random variables with zero mean. The canonical ordering of the terms $E_i \phi_i(\mathbf{x})$ also provides a suitable method for approximating $\xi_k(\mathbf{x})$, by truncating the sum at some finite $i = M$, and we could then let $\mathbf{E} = (E_1, \dots, E_M)$ (see for instance Wang 2008).

But obtaining the Karhunen-Loève transform can also be challenging. Because of this, we present an extremely simple approximation, that just relies on computation of the first two moments of ξ_k . We let \mathbf{E} be a 1-dimensional random variable with $E[\mathbf{E}] = 0$ and $E[\mathbf{E}^2] = 1$, and define

$$\hat{\xi}_k(\mathbf{x}) = E[\xi_k(\mathbf{x})] + \mathbf{E} \sqrt{\text{var}(\xi_k(\mathbf{x}))}. \tag{14}$$

This is indeed a very crude approximation, as essentially we assume that the values of ξ_k at any set of inputs \mathbf{x} are fully correlated. But for probabilistic surrogates used in structural reliability models, this is actually not that unreasonable, and as it turns out, for the examples we consider in Sect. 6 it seems sufficient.

Remark 6 Note that to update the approximate model $\hat{\xi}_k(\mathbf{x})$ in (14) given some new experiment (d_k, o_k) , we only need to update the mean and variance functions. This is in line




Fig. 2 Illustration of the finite-dimensional approximation (14)

with the numerically efficient Bayes linear approach (Goldstein and Wooff 2007), where random variables are specified only through the first two moments, and where the Bayesian updating given some experiment corresponds to computation of an adjusted mean and covariance. An application of the Bayes linear theory to sequential optimal design of experiments can be found in (Jones et al. 2018).

We note also that in the case where Gaussian processes are used as surrogate models, the classical and linear Bayesian approaches are computationally equivalent. Moreover, in the following section we will introduce the unscented transform for approximation of the updated/adjusted moments, and as a consequence the complete prior probability specification of \mathbf{E} becomes less relevant.

In the case where we are dealing with a hierarchical model, it might not be convenient to compute $E[\xi_k(\mathbf{x})]$ and $\text{var}(\xi_k(\mathbf{x}))$. If $\xi_k(\mathbf{x}) = g(\mathbf{Y}_k(\mathbf{x}))$ where $\mathbf{Y}_k(\mathbf{x})$ is a stochastic process with values in \mathbb{R}^n for any $\mathbf{x} \in \mathbb{X}$, we would instead approximate \mathbf{Y}_k with

$$\hat{\mathbf{Y}}_k = E[\mathbf{Y}_k] + L\mathbf{E}, \tag{15}$$

where \mathbf{E} is n -dimensional with $E[\mathbf{E}] = 0$, $E[\mathbf{E}\mathbf{E}^T] = I$, and the matrix L satisfies $LL^T = (\mathbf{Y}_k - E[\mathbf{Y}_k])(\mathbf{Y}_k - E[\mathbf{Y}_k])^T$. The approximation of ξ_k is then obtained as $\hat{\xi}_k(\mathbf{x}) = g(\hat{\mathbf{Y}}_k(\mathbf{x}))$. The same goes for the scenario with more than two layers in the hierarchy, for instance $\xi_k(\mathbf{x}) = g(\mathbf{Z}_k(\mathbf{Y}_k(\mathbf{x})))$, where we would approximate both $\mathbf{Z}_k(\mathbf{y})$ and $\mathbf{Y}_k(\mathbf{x})$. In any case, we end up with a finite-dimensional random variable \mathbf{E} , and we can define the approximation $\hat{\xi}_k(\mathbf{x}, \mathbf{E})$.

4.2 The unscented transform for epistemic uncertainty propagation

The unscented transform (UT) is a very efficient method for approximating the mean and covariance of a random variable after nonlinear transformation. UT is commonly applied in the context of Kalman filtering, and it is based on the general idea that *it is easier to approximate a probability distribution than an arbitrary nonlinear transformation* (Uhlmann 1995; Julier and Uhlmann 2004). Intuitively, given any finite-dimensional random variable \mathbf{E} we may define a set of *weighted sigma-points* $\{(v_i, \mathbf{e}_i)\}$, such that if $\{(v_i, \mathbf{e}_i)\}$

was considered as a discrete probability distribution, then its mean and covariance would coincide with \mathbf{E} . For any nonlinear transformation $\mathbf{Y} = f(\mathbf{E})$, if \mathbf{E} was discrete we could compute the mean and covariance of \mathbf{Y} exactly. The UT approximation is the result of such computation, where we make use of a small set of weighted points $\{(v_i, \mathbf{e}_i)\}$.

Specifically, let \mathbf{E} be a finite-dimensional random variable with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. A set of sigma-points for \mathbf{E} is a set of weighted samples $\{(v_1, \mathbf{e}_1), \dots, (v_n, \mathbf{e}_n)\}$ such that

$$\boldsymbol{\mu} = \sum_{i=1}^n v_i \mathbf{e}_i, \quad \boldsymbol{\Sigma} = \sum_{i=1}^n v_i (\mathbf{e}_i - \boldsymbol{\mu})(\mathbf{e}_i - \boldsymbol{\mu})^T. \tag{16}$$

If $\mathbf{y} = f(\mathbf{e})$ is any (generally nonlinear) transformation, the UT approximation of the mean and covariance of $\mathbf{Y} = f(\mathbf{E})$ are then obtained as

$$\begin{aligned} \hat{E}[\mathbf{Y}] &= \sum_{i=1}^n v_i \mathbf{y}_i, \\ \widehat{\text{Cov}}[\mathbf{Y}] &= \sum_{i=1}^n v_i (\mathbf{y}_i - \hat{E}[\mathbf{Y}])(\mathbf{y}_i - \hat{E}[\mathbf{Y}])^T, \end{aligned} \tag{17}$$

where $\mathbf{y}_i = f(\mathbf{e}_i)$.

Naturally, the selection of appropriate sigma-points is essential for UT to be successful. It is important to note that, although we may view the sigma-points as *weighted samples*, v_i and \mathbf{e}_i are fixed or given by some deterministic procedure. Moreover, the definition of sigma-points given in (16) does not require that the weights are nonnegative and sum to one. Although this conflicts with the intuition of approximating \mathbf{E} with a discrete random variable, the unscented transform still makes sense as a procedure for approximating statistics after nonlinear transformation.

Since the introduction of UT to Kalman filters in the 1990's, many different alternatives to sigma-point selection have been proposed (Menegaz et al. 2015). These mostly focus on applications where \mathbf{E} follows a multivariate Gaussian distribution, but we do not see this as a restriction since we will assume that \mathbf{E} can be represented as a transformation $\mathbf{E} = \mathcal{T}^{-1}(\mathbf{U})$ of a multivariate Gaussian variable \mathbf{U} . For the applications considered in this paper, we will let $\{(v_i, \mathbf{u}_i)\}$ denote a set of sigma-points that are appropriate for the multivariate standard normal $\mathbf{U} \sim \mathcal{N}(0, I)$ where $\text{dim}(\mathbf{U}) = \text{dim}(\mathbf{E})$. If \mathcal{T} is the corresponding isoprobabilistic transformation, i.e. $\mathcal{T}(\mathbf{E}) \sim \mathcal{N}(0, I)$ (see Appendix B.1), we will use $\{(v_i, \mathcal{T}^{-1}(\mathbf{u}_i))\}$ as a set of sigma-points for \mathbf{E} . Equivalently, we could also view this as taking the UT approximation of \mathbf{U} under a different transformation given by $f \circ \mathcal{T}$. For the numerical examples we present in this paper, we have made use of the the method developed by Merwe (2004), which produces a set of $n = 2 \cdot \text{dim}(\mathbf{E}) + 1$

points \mathbf{e}_i with corresponding weights⁷. Determining sigma-points with this procedure is quite straightforward, and the details are given in Appendix C. We note again that for any structural reliability model, as long as we do not change dimensionality of \mathbf{E} , determining the sigma-points is a one-time computation, and any subsequent UT approximation of $\mathbf{Y} = f(\mathbf{E})$, for some nonlinear transformation $f(\cdot)$, is computationally very efficient.

Remark 7 Note that it is not necessary that the sigma points used in the approximation of the mean and covariance in (17) are the same. In fact, the method presented in Appendix C makes use of two different sets of weights for these approximations. As this is not of any relevance for the remaining part of this paper, we will keep writing $\{v_i, \mathbf{e}_i\}$ as a single set of sigma-points to simplify the notation.

4.3 Generating samples in \mathbb{X}

In order to estimate the measures of residual uncertainty, we will need a set of samples of \mathbf{X} . We will generate a finite set of 3-tuples $\{(\mathbf{x}_i, w_i, \hat{\eta}_i)\}$, where $\{(\mathbf{x}_i, w_i)\}$ are weighted samples in \mathbb{X} suitable for obtaining importance sampling estimates of failure probabilities, and $\hat{\eta}_i$ is a number describing how influential a given sample (\mathbf{x}_i, w_i) is expected to be in such an estimate. In other words, $\{\mathbf{x}_i\}$ should be constructed to "cover the relevant regions in \mathbb{X} ", and for estimation we will only make use of a subset of $\{(\mathbf{x}_i, w_i)\}$. The relevant subset will be determined from the *measure of insignificance* $|\hat{\eta}_i|$, where we will only consider samples (\mathbf{x}_i, w_i) where $|\hat{\eta}_i|$ is below some threshold. We start by describing how the weighted samples $\{(\mathbf{x}_i, w_i)\}$ are generated.

4.3.1 Importance sampling

The general idea behind importance sampling is that if we select some random variable $Q \geq 0$ with law P_Q , such that $E_{P_X}[Q] = 1$ and $Q \neq 0$ P_X -almost surely, then

$$E_{P_X}[f(\mathbf{X})] = E_{P_Q}[f(\mathbf{X})/Q], \tag{18}$$

for any \mathcal{A} -measurable function $f(\mathbf{x})$. This is often useful for estimation, for instance when sampling from P_X is difficult, and in the case where we can find a Q such that estimates with respect to the right hand side of (18) are better (have lower variance) than estimating $E_{P_X}[f(\mathbf{X})]$ directly.

⁷ Other alternatives for sigma-point selection could also be applied, potentially with better performance. The method by Merwe (2004) depends on a set of parameters, and it could also be relevant to refine or learn the appropriate parameter values as in (Turner and Rasmussen 2010). However, in our current implementation we have only considered the fixed set of sigma-points given in Appendix C.

In the case where \mathbf{X} admits a probability density p_X , we can let q_X be any density function such that $q_X(\mathbf{x}) > 0$ whenever $p_X(\mathbf{x}) > 0$. Let $\mathbf{x}_1, \dots, \mathbf{x}_N$ be i.i.d. samples generated according to q_X , and define $w_i = p_X(\mathbf{x}_i)/q_X(\mathbf{x}_i)$. The importance sampling estimate of $E_{P_X}[f(\mathbf{X})]$ with respect to the proposal density q_X is then obtained as

$$E_{P_X}[f(\mathbf{X})] = E_{P_Q} \left[f(\mathbf{X}) \frac{p_X(\mathbf{X})}{q_X(\mathbf{X})} \right] \approx \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_i) w_i. \tag{19}$$

We now assume that the stochastic limit state can be written as $\xi_k(\mathbf{x}, \mathbf{e})$ for some finite-dimensional random variable \mathbf{E} , and for any deterministic performance function $\xi_k(\mathbf{x}, \mathbf{e})$ we will write $\alpha_k(\mathbf{e}) = \alpha(\xi_k(\mathbf{X}, \mathbf{e}))$ as the corresponding failure probability. An importance sampling estimate of $\alpha_k(\mathbf{e})$ is then given by (19) with $f(\mathbf{x}) = \mathbf{1}(\xi_k(\mathbf{x}, \mathbf{e}) \leq 0)$, that is

$$\hat{\alpha}_k(\mathbf{e}) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(\xi_k(\mathbf{x}_i, \mathbf{e}) \leq 0) w_i. \tag{20}$$

In order to obtain a good estimate of $\alpha_k(\mathbf{e})$, we would like the proposal distribution q_X to produce samples such that there is an even balance between the samples where $\xi_k(\mathbf{x}, \mathbf{e}) \leq 0$ and $\xi_k(\mathbf{x}, \mathbf{e}) > 0$, where at the same time p_X is as large as possible. One way to achieve this is to generate samples in the vicinity of points on the surface $\xi_k(\mathbf{x}, \mathbf{e}) = 0$ with (locally) maximal density. A point with this property is called a *design point*⁸ or *most probable failure point* in the structural reliability literature. We will let q_X represent a mixture of distributions, centered around different design points that are appropriate for different values of \mathbf{e} . The full details are given in Appendix B, where we also describe a simpler alternative than can be used in the case where design point searching is difficult or not appropriate.

4.3.2 The measure of insignificance $|\eta_i|$

Assume $\{(\mathbf{x}_i, w_i)\}$ is a set of samples capable of providing a satisfactory estimate of $\alpha_k(\mathbf{e})$, and we now want to estimate $\alpha_k(\mathbf{e}')$ for some new value \mathbf{e}' . If we know that the sign of $\xi_k(\mathbf{x}_i, \mathbf{e})$ and $\xi_k(\mathbf{x}_i, \mathbf{e}')$ will coincide for many of the samples \mathbf{x}_i , then the estimate of $\alpha_k(\mathbf{e}')$ can be obtained more efficiently by not computing all the terms in the sum (20). This is typically the case when \mathbf{e} and \mathbf{e}' are both sampled from \mathbf{E} . It is also true in the case where we want to estimate $\alpha_{k+1}(\mathbf{e}')$ given some new experiment (d_k, o_k) , if we assume that updating with respect to (d_k, o_k) has local effect (i.e. there are always regions in \mathbb{X} where $\xi_{k+1}(\mathbf{x}) \approx \xi_k(\mathbf{x})$), or if

⁸ The most common definition of a design point is that it is the point on the limit state surface with maximal density after transformation to the standard normal space. See Appendix B.1

the experiment is carried out to reduce the uncertainty in the level set $\xi_k = 0$ (which is what we intend to do).

In other words, we consider some perturbation of the performance function $\xi_k(\mathbf{x}, \mathbf{e})$, and we are interested in identifying the samples \mathbf{x}_i where $\mathbf{1}(\xi_k(\mathbf{x}_i, \mathbf{e}) \leq 0)$ does not change under the perturbation. For this purpose we define the function

$$\eta(\mathbf{x}, \xi) = E[\xi(\mathbf{x})] / \sqrt{\text{var}(\xi(\mathbf{x}))}, \tag{21}$$

and let $\eta_i = \eta(\mathbf{x}_i, \xi_k)$ be defined with respect to the relevant process ξ_k . Here η_i describes how uncertain $\xi_k(\mathbf{x}_i)$ is around the critical value $\xi_k = 0$, in the sense that if $|\eta_i|$ is small (close to zero) then $\xi_k(\mathbf{x}_i) > 0$ and $\xi_k(\mathbf{x}_i) \leq 0$ may both be probable outcomes. Conversely, if $|\eta_i|$ is large then either $P(\xi_k(\mathbf{x}_i) \leq 0) \approx 0$ or $P(\xi_k(\mathbf{x}_i) \leq 0) \approx 1$, and the input \mathbf{x}_i is *insignificant* as it is unnecessary to keep track of changes in $\mathbf{1}(\xi_k(\mathbf{x}_i) \leq 0)$. We will use η_i to *prune* the sample set $\{(\mathbf{x}_i, w_i)\}$, by only considering the samples where $|\eta_i|$ is below a given threshold τ . Although this is an intuitive idea, we may also justify the definition of η and selection of a threshold τ more formally by making use of the following proposition.

Proposition 41 *Given any process $\xi(\mathbf{x})$, let $\eta(\mathbf{x}) = \eta(\mathbf{x}, \xi)$ be defined as in (21) and let $\tau > \sqrt{2}$. Assume $\xi^{(1)}$ and $\xi^{(2)}$ are two i.i.d. random samples from $\xi(\mathbf{x})$. Then,*

$$P\left(\mathbf{1}(\xi^{(1)} \leq 0) \neq \mathbf{1}(\xi^{(2)} \leq 0) \mid |\eta| \geq \tau\right) \leq \frac{2}{\tau^2} \left(1 - \frac{1}{\tau^2}\right). \tag{22}$$

Proof Let $p = P(\xi(\mathbf{x}) \leq 0)$ and $\gamma(p) = p(1 - p)$ for short (note also that this is (8) for $\xi = \xi_k$), and observe that $P(\mathbf{1}(\xi^{(1)} \leq 0) \neq \mathbf{1}(\xi^{(2)} \leq 0)) = 2\gamma(p)$. Assume first that $\eta > 0$. Then $E[\xi] > 0$ and by Chebyshev’s one-sided inequality we get

$$\eta = \tau \Rightarrow p \leq \frac{\text{var}(\xi(\mathbf{x}))}{(\text{var}(\xi(\mathbf{x})) + E[\xi(\mathbf{x})]^2)} \leq \frac{1}{\tau^2},$$

and as $\tau > \sqrt{2}$ we also get $p \leq 1/2$. Since $\gamma(p)$ is increasing for $p \in [0, 1/2]$, we must have $\gamma(p) \leq \gamma(1/\tau^2)$.

Conversely, if $-\tau = \eta < 0$ then $p \geq 1 - 1/\tau^2 \geq 1/2$, and as $\gamma(p)$ is decreasing for $p \in [1/2, 1]$ we have that $\gamma(p) \leq \gamma(1 - 1/\tau^2) = \gamma(1/\tau^2)$. Hence, combining both cases we get $|\eta| = \tau \Rightarrow \gamma(p) \leq \gamma(1/\tau^2)$, and (22) is proved by observing that $\gamma(1/(\tau + \varepsilon)^2) \leq \gamma(1/\tau^2)$ for any $\varepsilon > 0$. \square

Although Proposition 41 holds in general, tighter (and probably more realistic) bounds can be obtained by making assumptions on the form of $\xi(\mathbf{x})$. For instance, in the case

where $\xi(\mathbf{x})$ is Gaussian we obtain

$$P\left(\mathbf{1}(\xi^{(1)} \leq 0) \neq \mathbf{1}(\xi^{(2)} \leq 0) \mid |\eta| \geq \tau\right) \leq 2\Phi(\tau)\Phi(-\tau), \tag{23}$$

where $\Phi(\cdot)$ is the standard normal CDF.

We will make use of $\hat{\eta}_i$ obtained as the UT approximation of η_i . That is, $\hat{\eta}_i$ is in general obtained from the finite-dimensional approximation described in Sect. 4.1, combined with the UT approximation (17) with $\mathbf{Y} = \hat{\xi}_k(\mathbf{x}, \mathbf{E})$.

4.4 Importance sampling estimates with pruning

Let $\{(\mathbf{x}_i, w_i, \hat{\eta}_i) \mid i \in \mathcal{I}\}$, $\mathcal{I} = \{1, \dots, N_0\}$ be a set of samples generated as described in Sect. 4.3. Given some fixed threshold $\tau > 0$, we define the subset of *pruned samples* as the ones corresponding to the index set $\mathcal{I}_\tau = \{i \in \mathcal{I} \mid \hat{\eta}_i < \tau\}$, and define $\bar{\mathcal{I}}_\tau = \mathcal{I} \setminus \mathcal{I}_\tau$. If $f(\mathbf{x})$ is some \mathcal{A} -measurable function where we know a priori the value of $f_i = f(\mathbf{x}_i)$ for all $i \in \bar{\mathcal{I}}_\tau$, then we can immediately compute

$$\bar{h} = \frac{1}{N_0} \sum_{i \in \bar{\mathcal{I}}_\tau} f_i w_i, \tag{24}$$

and the importance sampling estimate of the expectation of $f(\mathbf{X})$ becomes

$$\hat{E}[f(\mathbf{X})] = \bar{h} + \frac{1}{N_0} \sum_{i \in \mathcal{I}_\tau} f(\mathbf{x}_i) w_i. \tag{25}$$

If we let

$$s_{\bar{h}} = \frac{1}{N_0} \sum_{i \in \bar{\mathcal{I}}_\tau} (f_i w_i - \hat{E}[f(\mathbf{X})])^2, \tag{26}$$

then an unbiased estimate of the sample variance is given as

$$\widehat{\text{var}}(\hat{E}[f(\mathbf{X})]) = \frac{s_{\bar{h}}}{N_0 - 1} + \frac{1}{N_0(N_0 - 1)} \sum_{i \in \mathcal{I}_\tau} (f(\mathbf{x}_i) w_i - \hat{E}[f(\mathbf{X})])^2, \tag{27}$$

which shows the general idea with this *pruning*, namely that low variance estimates of $E[f(\mathbf{X})]$ can be obtained with a small number of evaluations $f(\mathbf{x}_i)$, assuming that the subset \mathcal{I}_τ is small compared to \mathcal{I} (and that the assumed values f_i are correct).

One drawback with this procedure is that we do not have control over the number of pruned samples, which still might be very large. In order to set an upper bound on the number of evaluations $f(\mathbf{x}_i)$, we let $\mathcal{I}_\tau^n \subseteq \mathcal{I}_\tau$ contain the first n elements of \mathcal{I}_τ (or some other subset, as long as the elements of

$\{\mathbf{x}_i \mid i \in \mathcal{I}_\tau^n\}$ remain independent). An importance sampling estimate of $E[f(\mathbf{X})]$ using only samples from \mathcal{I}_τ^n is given as

$$\widehat{E}[f(\mathbf{X})] = \bar{h} + \bar{r}, \quad \bar{r} = \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}_\tau^n} f(\mathbf{x}_i)w_i, \tag{28}$$

where $N_\tau = |\mathcal{I}_\tau|$, and we may estimate the sample variance as

$$\widehat{\text{var}}(\widehat{E}[f(\mathbf{X})]) = \frac{1}{N_0 - 1} (s_{\bar{h}} - \bar{h}^2) + \frac{N_\tau}{nN_0 - N_\tau} \left(-\bar{r}^2 + \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}_\tau^n} (f(\mathbf{x}_i)w_i)^2 \right). \tag{29}$$

Obtaining consistency results is easy under the ideal assumption that $n(N_0 - N_\tau)/N_\tau$ is an integer, and the formulas in (28)–(29) comes as a consequence of the following result.

Proposition 42 *Assume $n(N_0 - N_\tau)/N_\tau \in \mathbb{N}$. Then (28) is an unbiased estimate of $E[f(\mathbf{X})]$ and (29) is an unbiased estimate of the sample variance.*

Proof Let $\bar{\mathcal{I}}_\tau^n$ be a set of $n(N_0 - N_\tau)/N_\tau$ elements selected uniformly random from $\bar{\mathcal{I}}_\tau$ and define $\mathcal{I}^n = \mathcal{I}_\tau^n \cup \bar{\mathcal{I}}_\tau^n$. Then $\{\mathbf{x}_i \mid i \in \mathcal{I}^n\}$ is a set of size $|\mathcal{I}^n| = nN_0/N_\tau$, containing i.i.d. samples from the proposal distribution with density $q(\mathbf{x})$. To show consistency we replace each sample \mathbf{x}_i with i.i.d. random variables \mathbf{X}_i distributed according to q . We then define $\hat{\mu} = \hat{\mu}_1 + \hat{\mu}_2$ where

$$\hat{\mu}_1 = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \mathbf{1}(\eta(\mathbf{X}_i) \geq \tau) f(\mathbf{X}_i)w(\mathbf{X}_i),$$

$$\hat{\mu}_2 = \frac{1}{|\mathcal{I}^n|} \sum_{i \in \mathcal{I}^n} \mathbf{1}(\eta(\mathbf{X}_i) < \tau) f(\mathbf{X}_i)w(\mathbf{X}_i),$$

and where $w(\mathbf{x}) = p(\mathbf{x})/q(\mathbf{x})$, and we can observe that $\hat{\mu} = \widehat{E}[f(\mathbf{X})]$ when $\mathbf{X}_i = \mathbf{x}_i$.

To show that $\widehat{E}[f(\mathbf{X})]$ is unbiased it is enough to observe that $E_q[\hat{\mu}] = E_q[\mathbf{1}(\eta(\mathbf{X}) \geq \tau) f(\mathbf{X})w(\mathbf{X})] + E_q[\mathbf{1}(\eta(\mathbf{X}) < \tau) f(\mathbf{X})w(\mathbf{X})] = E_q[f(\mathbf{X})w(\mathbf{X})] = E[f(\mathbf{X})]$.

As for the variance, we first observe that $\text{var}(\hat{\mu}) = \text{var}(\hat{\mu}_1) + \text{var}(\hat{\mu}_2)$ where $\text{var}(\hat{\mu}_1) = \text{var}(\mathbf{1}(\eta(\mathbf{X}) \geq \tau) f(\mathbf{X})w(\mathbf{X})/|\mathcal{I}|)$ and $\text{var}(\hat{\mu}_2) = \text{var}(\mathbf{1}(\eta(\mathbf{X}) < \tau) f(\mathbf{X})w(\mathbf{X})/|\mathcal{I}^n|)$. Replacing $\text{var}(\hat{\mu}_1)$ and $\text{var}(\hat{\mu}_2)$ with unbiased sample variances using the samples $\mathbf{X}_i = \mathbf{x}_i$ we obtain

$$\widehat{\text{var}}(\hat{\mu}_1) = \frac{1}{|\mathcal{I}|(|\mathcal{I}|-1)} \sum_{i \in \mathcal{I}} (\mathbf{1}(\eta(\mathbf{x}_i) \geq \tau) f(\mathbf{x}_i)w(\mathbf{x}_i) - \bar{h})^2$$

$$= \frac{1}{|\mathcal{I}|-1} \left(-\bar{h}^2 + \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} (\mathbf{1}(\eta(\mathbf{x}_i) \geq \tau) f(\mathbf{x}_i)w(\mathbf{x}_i))^2 \right)$$

$$= \frac{1}{|\mathcal{I}|-1} (-\bar{h}^2 + s_{\bar{h}}),$$

and similarly

$$\widehat{\text{var}}(\hat{\mu}_2) = \frac{1}{|\mathcal{I}^n|-1} \left(-\bar{r}^2 + \frac{1}{|\mathcal{I}^n|} \sum_{i \in \mathcal{I}^n} (f(\mathbf{x}_i)w(\mathbf{x}_i))^2 \right),$$

where we have used that \bar{h} and \bar{r} are unbiased estimates of $E_q[\hat{\mu}_1]$ and $E_q[\hat{\mu}_2]$, respectively. The expression in (29) is then obtained as $\widehat{\text{var}}(\hat{\mu}_1) + \widehat{\text{var}}(\hat{\mu}_2)$ using that $|\mathcal{I}| = N_0$ and $|\mathcal{I}^n| = nN_0/N_\tau$. \square

4.5 The UT-MCIS approximation of $H_{1,k}, H_{2,k}$ and $H_{3,k}$

Using the tools introduced in the preceding subsections, we now present how the measures of residual uncertainty, $H_{1,k}, H_{2,k}$ and $H_{3,k}$, can be approximated using Monte Carlo simulation with importance sampling (MCIS) combined with the unscented transform (UT) for epistemic uncertainty propagation.

We first let $\hat{\xi}_k(\mathbf{x}, \mathbf{E})$ be the finite-dimensional approximation introduced in Sect. 4.1, with the corresponding failure probability $\hat{\alpha}_k(\mathbf{E}) = \alpha(\hat{\xi}_k(\mathbf{x}, \mathbf{E}))$. We then let $\{(\mathbf{x}_i, w_i, \hat{\eta}_i) \mid i \in \mathcal{I}\}$, $\mathcal{I} = \{1, \dots, N_0\}$ be a set of samples generated as described in Sect. 4.3, where $\hat{\eta}_i$ is obtained using the UT approximation of $\hat{\xi}_k(\mathbf{x}_i, \mathbf{E})$. We will make use of importance sampling estimates as introduced in Sect. 4.4, where $\mathcal{I}_\tau = \{i \in \mathcal{I} \mid \hat{\eta}_i < \tau\}$, and estimation is based on a small subset $\{(\mathbf{x}_i, w_i, \hat{\eta}_i) \mid i \in \mathcal{I}_\tau^n\}$ where $\mathcal{I}_\tau^n \subset \mathcal{I}_\tau$ and $|\mathcal{I}_\tau^n| = n < N_\tau = |\mathcal{I}_\tau|$.

4.5.1 Approximating $H_{1,k}$

Let $f_i = \mathbf{1}(\hat{\eta}_i \leq 0)$ for $i \in \bar{\mathcal{I}}_\tau$ and compute \bar{h}_1 as in (24). We will let $\{v_j, \mathbf{e}_j \mid j = 1, \dots, M\}$ denote the set of sigma-points as introduced in Sect. 4.2.

For any fixed \mathbf{e}_j , the corresponding importance sampling estimate of the failure probability $\hat{\alpha}_k(\mathbf{e}_j)$ is obtained as

$$\hat{\alpha}_k^j = \bar{h}_1 + \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}_\tau^n} \mathbf{1}(\hat{\xi}_k(\mathbf{x}_i, \mathbf{e}_j) \leq 0) w_i, \tag{30}$$

and we let $\hat{H}_{1,k}$ be given by the UT approximation

$$\widehat{E}[\hat{\alpha}_k] = \sum_{j=1}^M v_j \hat{\alpha}_k^j, \tag{31}$$

$$\hat{H}_{1,k} = \widehat{\text{var}}[\hat{\alpha}_k] = \sum_{j=1}^M v_j (\hat{\alpha}_k^j - \widehat{E}[\hat{\alpha}_k])^2.$$

4.5.2 Approximating $H_{2,k}$ and $H_{3,k}$

Both $H_{2,k}$ and $H_{3,k}$ are defined through the function $\gamma_k(\mathbf{x})$, which represents the uncertainty in the sign of $\xi_k(\mathbf{x})$. We will approximate $\gamma_k(\mathbf{x}_i)$ with the following function

$$\hat{\gamma}_k^i = \Phi(\hat{\eta}_i)\Phi(-\hat{\eta}_i), \tag{32}$$

where $\Phi(\cdot)$ is the standard normal CDF. There are two ways of interpreting this approximation. First of all, $\hat{\gamma}_{k,i}$ corresponds to the case where $\hat{\xi}_k(\mathbf{x}_i, \mathbf{E})$ is Gaussian, which may or may not be an appropriate assumption. Alternatively, we can think of $\gamma_k(\mathbf{x})$ as a measure of uncertainty in $\mathbf{1}(\xi_k(\mathbf{x}) \leq 0)$, and any $\gamma(\mathbf{x}) \propto -|\eta(\mathbf{x})| = -|E[\xi_k(\mathbf{x})]|/\sqrt{\text{var}(\xi_k(\mathbf{x}))}$ is reasonable. In this scenario it is natural to consider $\gamma = s(\eta)s(-\eta)$ for some sigmoid function $s(\cdot)$, and the function $\Phi(\cdot)$ in (32) is one such alternative.

For a single approximation of $H_{2,k}$ or $H_{3,k}$ it is really not necessary to split the importance sampling estimate as in (24)–(28), but we will present it in this form as it will be convenient when we consider strategies for optimization. Given $\hat{\gamma}_k^i$ as in (32), we approximate $H_{2,k}$ and $H_{3,k}$ by

$$\begin{aligned} \hat{H}_{2,k} &= \bar{h}_2 + \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}_\tau^n} \hat{\gamma}_k^i w_i, \\ \hat{H}_{3,k} &= \left(\bar{h}_3 + \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}_\tau^n} \sqrt{\hat{\gamma}_k^i} w_i \right)^2, \end{aligned} \tag{33}$$

where we let $\bar{h}_2 = \bar{h}_3 = 0$. Alternatively, if the intention is to use $H_{2,k}$ and $H_{3,k}$ as upper bounds on $H_{1,k}$, we could let $\bar{h}_2 = \frac{1}{N_0} \Phi(\tau)\Phi(-\tau) \sum w_i$, $\bar{h}_3 = \frac{1}{N_0} \sqrt{\Phi(\tau)\Phi(-\tau)} \sum w_i$ where the sums are over $i \in \mathcal{I}_\tau$.

5 Numerical procedure for one-step lookahead optimization

In the one-step lookahead case, the optimal decision d_k at each time step k is found by solving the following optimization problem

$$d_k = \arg \min_{d \in \mathbb{D}} J_{i,k}(d) \text{ for } i = 1, 2, \text{ or } 3, \tag{34}$$

where $J_{i,k}(d)$ is the relevant acquisition function as defined in (13). We propose a procedure where we make use of a UT-MCIS approximation of $J_{i,k}(d)$ to find an approximate solution to (34). This will build on the approximation of $H_{i,k}$ introduced in Sect. 4, but where we now also make use of the predictive model δ to approximate expectations with respect to future values of $H_{i,k+1}$.

In Sects. 5.1 and 5.2 we present how the UT-MCIS approximation of $J_{i,k}(d)$ is obtained, and in Sect. 5.3 we propose a criterion for determining when the sequence of experiments should be stopped. The final algorithm is summarized in Sect. 5.4

5.1 The probabilistic model $(\hat{\xi}_k, \hat{\delta}_k)$

Starting with some probabilistic model (ξ_k, δ_k) , recall that ξ_k represents uncertainty about the performance of the system under consideration, and δ_k represents uncertainty with respect to outcomes of certain decisions. We have already discussed how to obtain a finite-dimensional approximation of ξ_k , and likewise, this will also be needed for δ_k .

Assuming δ_k is square integrable, we will make use of the same type of finite-dimensional approximation as the one introduced for ξ_k in Sect. 4.1. In this way, we end up with two finite-dimensional \mathcal{E}_k -measurable random variables \mathbf{E}^ξ and \mathbf{E}^δ , which in turn determine the approximations $\hat{\xi}_k(\mathbf{x}, \mathbf{E}^\xi)$ and $\hat{\delta}_k(d, \mathbf{E}^\delta)$, where both $\hat{\xi}_k(\mathbf{x}, \mathbf{e})$ and $\hat{\delta}_k(d, \mathbf{e})$ are deterministic functions for \mathbf{e} fixed. Here \mathbf{E}^ξ and \mathbf{E}^δ are generally not independent.

Remark 8 Note that if $\delta(d)$ is a function of some of the uncertain sub-components of ξ , then we might already have a finite-dimensional approximation of δ available.

Consider for instance the model in Example 4 and the discussion in the end of Sect. 4.1. In this case, $\hat{\xi}$ is obtained as a function of the finite-dimensional approximation $\hat{y}_1(\mathbf{x}, \mathbf{E})$ of a sub-component $\tilde{y}_1(\mathbf{x})$, and $\delta(d)$ is given as $\delta(d(\mathbf{x})) = \tilde{y}_1(\mathbf{x}) + \epsilon(\mathbf{x})$. Hence, all we need is to find a finite-dimensional representation of the noise $\epsilon(\mathbf{x})$. But observational noise such as $\epsilon(\mathbf{x})$ is often described as a function of \mathbf{x} and some 1-dimensional random variable, in which case no additional approximation will be needed.

We will let $(\hat{\xi}_k, \hat{\delta}_k)$ denote the finite-dimensional approximation of (ξ_k, δ_k) corresponding to a finite-dimensional random variable $\mathbf{E} = (\mathbf{E}^\xi, \mathbf{E}^\delta)$, and where $(\hat{\xi}_0, \hat{\delta}_0)$ is the initial model that is used as input for determining the first decision d_1 .

Remark 9 In the canonical case where a surrogate $\tilde{y}(\mathbf{x})$ is used to represent some unknown function $y(\mathbf{x})$, an initial set of experiments is often performed to establish $\tilde{y}(\mathbf{x})$ before any sequential strategy is started. For instance, in the case where evaluation of $y(\mathbf{x})$ means running deterministic computer code, it is normal to set up a space-filling initial design using e.g. Latin Hypercube Sampling.

When $\tilde{y}(\mathbf{x})$ is a Gaussian process model as described in Appendix A, specific mean and covariance functions may also be selected based on knowledge or assumptions about

the phenomenon that is being modelled by $y(\mathbf{x})$. For estimation of failure probabilities it is also convenient to make use of conservative prior mean values. That is, prior to any experiment $\tilde{y}(\mathbf{x})$ will correspond to a value associated with poor structural performance (small ξ), such that $\alpha(\xi)$ will be biased towards higher failure probabilities in the absence of experimental evidence. This reasonable from a safety perspective, and also numerically as larger failure probabilities are easier to estimate.

5.2 Acquisition function approximation

To find an approximate solution to the optimization problem (34), we will replace the acquisition function $J_{i,k}(d)$ with an approximation $\hat{J}_{i,k}(d)$. Recall that $J_{i,k}(d)$ as defined in (13) is a function of $E_{k,d}[H_{i,k+1}]$, where $E_{k,d}$ is the conditional expectation with respect to \mathcal{E}_k with $d_k = d$. In Sect. 4 we introduced an approximation $H_{i,k}$, and we will make use of the same idea to approximate $E_{k,d}[H_{i,k+1}]$.

Assume k experiments have been performed, giving rise to the model (ξ_k, δ_k) and the approximation $(\hat{\xi}_k, \hat{\delta}_k)$. If we consider the k th decision $d_k = d$, then $H_{i,k+1}$ is a priori a $\delta_k(d)$ -measurable random variable. That is, $H_{i,k+1}$ is a function of $\delta_k(d)$, and we are interested in the expectation $E_{k,d}[H_{i,k+1}] = E[H_{i,k+1}(\delta_k(d))]$. To approximate this quantity, we can make use of $(\hat{\xi}_k, \hat{\delta}_k)$ in the place of (ξ_k, δ_k) , in which case $H_{i,k+1}$ becomes a function of \mathbf{E} and we can approximate its expectation using UT.

The approximate acquisition functions are then given as

$$\hat{J}_{i,k}(d) = \lambda(d)\hat{E}_{k,d}[\hat{H}_{i,k+1}], \tag{35}$$

where $\hat{E}_{k,d}[\hat{H}_{i,k+1}]$ is obtained as follows:

5.2.1 Generating samples of $\hat{\xi}_{k+1}$

Let $\{(v_j^\xi, \mathbf{e}_j^\xi) \mid j = 1, \dots, M^\xi\}$ and $\{(v_m^\delta, \mathbf{e}_m^\delta) \mid m = 1, \dots, M^\delta\}$ denote sigma-points as introduced in Sect. 4.2 for \mathbf{E}^ξ and \mathbf{E}^δ , respectively. We then let $\{(\mathbf{x}_i, w_i, \hat{\eta}_i) \mid i \in \mathcal{I}\}$, $\mathcal{I} = \{1, \dots, N_0\}$ be a set of samples generated as described in Sect. 4.3, where $\hat{\eta}_i$ is obtained using the UT approximation of $\hat{\xi}_k(\mathbf{x}_i, \mathbf{E}^\xi)$. As for the approximation of $H_{i,k}$ discussed in Sect. 4.5, we let $\mathcal{I}_\tau = \{(\mathbf{x}_i, w_i, \hat{\eta}_i) \mid \hat{\eta}_i < \tau\}$ and define the subset $\mathcal{I}_\tau^n \subseteq \mathcal{I}_\tau$ of size n .

The approximations of $E_{k,d}[H_{i,k+1}]$ for $i = 1, 2$ and 3 will all be based on samples of $\hat{\xi}_{k+1}$ of the form

$$\hat{\xi}_{k+1}^{m,i,j} = \hat{\xi}_{k+1}(\mathbf{x}, \mathbf{e}_j^\xi, d, \mathbf{e}_m^\delta), \tag{36}$$

where $\hat{\xi}_{k+1}(\mathbf{x}, \mathbf{e}_j^\xi, d, \mathbf{e}_m^\delta)$ is the finite-dimensional approximation of $\xi_k \mid d_k = d, o_k = \hat{\delta}(\mathbf{e}_m^\delta)$ evaluated at $(\mathbf{x}, \mathbf{e}_j^\xi)$. The scalar $\hat{\xi}_{k+1}^{m,i,j}$ is computed for all $j = 1, \dots, M^\xi$, $m = 1, \dots, M^\delta$ and $i \in \mathcal{I}_\tau^n$. As in Sect. 4.5 we set $\bar{h}_2 = \bar{h}_3 = 0$ and compute \bar{h}_1 as in (24) with $f_i = \mathbf{1}(\hat{\eta}_i \leq 0)$ for $i \notin \mathcal{I}_\tau$.

5.2.2 The UT-MCIS approximation of $E_{k,d}[H_{1,k+1}]$

The approximation $\hat{E}_{k,d}[\hat{H}_{1,k+1}]$ is just a weighted sum of the terms in (36), but for clarity we present it in the following three steps

MCIS of $\alpha(\hat{\xi}_{k+1})$:

$$\hat{\alpha}_{k+1}^{m,j} = \bar{h}_1 + \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}_\tau^n} \mathbf{1}(\hat{\xi}_{k+1}^{m,i,j} \leq 0) w_i, \tag{37}$$

UT of $H_{1,k+1}$:

$$\hat{H}_{1,k+1}^m = \sum_{j=1}^{M^\xi} v_j^\xi (\hat{\alpha}_{k+1}^{m,j})^2 - \left(\sum_{j=1}^{M^\xi} v_j^\xi \hat{\alpha}_{k+1}^{m,j} \right)^2, \tag{38}$$

UT of $E_{k,d}[H_{1,k+1}]$:

$$\hat{E}_{k,d}[\hat{H}_{1,k+1}] = \sum_{m=1}^{M^\delta} v_m^\delta \hat{H}_{1,k+1}^m. \tag{39}$$

5.2.3 The UT-MCIS approximation of $E_{k,d}[H_{2,k+1}]$ and $E_{k,d}[H_{3,k+1}]$

The weighted sums that gives the approximations of $E_{k,d}[H_{2,k+1}]$ and $E_{k,d}[H_{3,k+1}]$ can be obtained as follows

$$\text{UT of } E[\hat{\xi}_{k+1}(\mathbf{x}_i)]: \hat{\mu}_{k+1}^{i,m} = \sum_{j=1}^{M^\xi} v_j^\xi \hat{\xi}_{k+1}^{m,i,j}, \tag{40}$$

UT of $\text{var}[\hat{\xi}_{k+1}(\mathbf{x}_i)]:$

$$(\hat{\sigma}_{k+1}^{i,m})^2 = \sum_{j=1}^{M^\xi} v_j^\xi (\hat{\xi}_{k+1}^{m,i,j} - \hat{\mu}_{k+1}^{i,m})^2, \tag{41}$$

Using Φ to approximate $\hat{\gamma}_{k+1}(\xi_i)$:

$$\hat{\gamma}_{k+1}^{i,m} = \Phi(\hat{\eta}_{k+1}^{i,m})\Phi(-\hat{\eta}_{k+1}^{i,m}), \hat{\eta}_{k+1}^{i,m} = \hat{\mu}_{k+1}^{i,m}/\hat{\sigma}_{k+1}^{i,m} \tag{42}$$

MCIS of $H_{2,k+1}$:

$$\hat{H}_{2,k+1}^m = \bar{h}_2 + \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}_\tau^n} \hat{\gamma}_{k+1}^{i,m} w_i, \tag{43}$$

MCIS of $H_{3,k+1}$:

$$\sqrt{\hat{H}_{3,k+1}^m} = \bar{h}_3 + \frac{N_\tau}{nN_0} \sum_{i \in \mathcal{I}^n} \sqrt{\hat{\gamma}_{k+1}^{i,m}} w_i, \tag{44}$$

and where $\hat{E}_{k,d}[\hat{H}_{2,k+1}]$ and $\hat{E}_{k,d}[\hat{H}_{3,k+1}]$ are obtained with the same formula as for $\hat{E}_{k,d}[\hat{H}_{1,k+1}]$ in (39).

Remark 10 The number of model updates and function evaluations needed to generate the set $\{\hat{\xi}_{k+1}^{m,i,j}\}$ are M^δ and $nM^\xi M^\delta$. We can view this as a discretization of the system dynamics, where there are only M^δ possible future scenarios corresponding to the decision $d_k = d$, which are given by the model updates $\xi_k \rightarrow \xi_{k+1}(\mathbf{e}_m^\delta) = \xi_k | d_k = d, o_k = \hat{\delta}_k(\mathbf{e}_m^\delta)$. The samples in (36) are the ones needed for approximating the measure of residual uncertainty corresponding to $\xi_{k+1}(\mathbf{e}_m^\delta)$ for each $m = 1, \dots, M^\delta$.

Moreover, although the approximations $\hat{E}_{k,d}[\hat{H}_{i,k+1}]$ are presented as weighted sums of the $nM^\xi M^\delta$ terms $\hat{\xi}_{k+1}^{m,i,j}$, this can also be obtained from a sequence of nested loops for a more memory efficient implementation. See for instance the schematic illustration in Fig. 3.

5.3 Stopping criterion

For design strategies that make use of heuristic acquisition functions, it can be challenging to determine an appropriate stopping criterion. Here, we have considered the approximation $\hat{H}_{1,k}$ which has a natural interpretation. Hence, even if we make use of a criteria such as $\hat{H}_{2,k}$ or $\hat{H}_{3,k}$ to determine the next optimal decision, it makes sense to use $\hat{H}_{1,k}$ as an indicator of when the potential uncertainty reduction from future experiments is diminishing.

We will let $\hat{E}[\hat{\alpha}_k]$ and $\hat{H}_{1,k}$ be given as in (31), and define

$$\hat{V}_k = \frac{\sqrt{\hat{H}_{1,k}}}{\hat{E}[\hat{\alpha}_k]}. \tag{45}$$

Then \hat{V}_k is the UT-MCIS approximation of the *coefficient of variation* of the failure probability α_k with respect to epistemic uncertainty. We will let $\hat{V}_k \leq V_{\max}$ for some threshold V_{\max} serve as a criterion for stopping the iteration procedure, in the case where a predefined maximum number of iterations K_{\max} has not already been reached.

Remark 11 The coefficient of variation is often used as a numerical criterion for convergence in Monte Carlo simulation. In structural reliability analysis, a coefficient of variation below 0.05 is often used as an acceptable level for failure probability estimation.

Note also that the criterion $\hat{V}_k \leq V_{\max}$ for arbitrary $V_{\max} \geq 0$ implicitly assumes that the epistemic uncertainty can be reduced to zero in the limit. If this is not the case, one might

instead consider stopping when \hat{V}_k is no longer decreasing. A different stopping criterion is also considered in Sect. 6.4.

5.4 Algorithm

The complete procedure for myopic/one-step lookahead optimization is summarized in Algorithm 1. Note that for simplicity the number of MCIS samples N_0 and n are specified as input, but one may also consider deciding these using (28) and (29). Using a standard technique in Monte Carlo simulation, one could keep increasing N_0 and n until the coefficient of variation (std/mean) of the relevant estimator is sufficiently small.

Algorithm 1: One-step lookahead optimization

input: Model and sigma-points: $(\hat{\xi}_0, \hat{\delta}_0)$ and $\{(v_j^\xi, \mathbf{e}_j^\xi)\}, \{(v_j^\delta, \mathbf{e}_j^\delta)\}$.
 Number of samples for UT-MCIS and threshold: $N_0, n \in \mathbb{N}$ and $\tau > 0$,
 Max number of iterations and convergence criteria: K_{\max} and V_{\max} .

for $k = 0$ **to** $K_{\max} - 1$ **do**

- (1) Generate samples $\{(\mathbf{x}_i, w_i, \hat{\eta}_i)\}$ as described in Section 4.3 and compute $\hat{h}_1 = \frac{1}{N_0} \sum_{|\hat{\eta}_i| \geq \tau} \mathbf{1}(\hat{\eta}_i \leq 0) w_i$.
- (2) Compute \hat{V}_k as in (45)
- if** $\hat{V}_k > V_{\max}$ **then**
 - (3) Compute the set $\{\hat{\xi}_{k+1}^{m,i,j}\}$ as in (36) and define the function $\hat{J}_{i,k}(d)$ as in (35) (for $i = 1, 2$, or 3 depending on the acquisition function of choice)
 - (4) Find the optimal decision: $d_k = \arg \min_{d \in \mathbb{D}} \hat{J}_{i,k}(d)$
 - (5) Make decision d_k and obtain (d_k, o_k)
 - (6) Update model $(\hat{\xi}_{k+1}, \hat{\delta}_{k+1}) = (\hat{\xi}_k, \hat{\delta}_k) | (d_k, o_k)$
- else**
 - Break. Convergence has been reached before K_{\max} iterations.

6 Numerical experiments

Here we present a few numerical experiments using the algorithm for one-step lookahead optimal design presented in Sect. 5.4. Four experiments are presented, each with its own objective:

- (1) Section 6.1: A toy example in 1d for conceptual illustration of the sequential design procedure.
- (2) Section 6.2: A hierarchical model with multiple 'expensive' sub-components.
- (3) Section 6.3: A non-hierarchical benchmark problem for comparison against alternative strategies.




Fig. 3 Illustration of how $\widehat{E}_{k,d}[\widehat{H}_{1,k+1}]$ is obtained using UT over epistemic uncertainties. Here $\widehat{H}_{1,k+1}(\mathbf{e}_m^\delta)$ for $m = 1, \dots, M^\delta$ is obtained from the MCIS estimates of $\alpha(\widehat{\xi}_{k+1}(\mathbf{e}_m^\xi))$

(4) Section 6.4 A model that is more in resemblance of a realistic application in structural reliability analysis, where we introduce different types of decisions by considering both probabilistic function approximation and Bayesian inference of model parameters through measurements with noise.

All numerical experiments have been performed using Algorithm 1 with the parameters $\tau = 3$, $N_0 = 10^4$, $n = 10^3$ and $V_{\max} = 0.05$. This choice of V_{\max} corresponds to a 5% coefficient of variation on the estimated failure probability, and $\tau = 3$ should give a reasonable coverage for importance sampling (from Proposition 41 the probability of misclassification is less than 0.2 in the extreme case (Chebyshev) and less than 0.003 under the Gaussianity assumption). The number of samples, N_0 and n was chosen to make evaluation of the acquisition function reasonably cheap, and the choice $N_0 = 10^4$, $n = 10^3$ worked well in all of our experiments. Note that for final estimates of the failure probability, after an optimal decision has been found, a larger number of samples may be used for increased accuracy. The probabilistic surrogate models used in the examples are all Gaussian process (GP) models with Matérn 5/2 covariance. A short summary of the relevant Gaussian process theory is given in Appendix A.

6.1 Example 1: Illustrative 1d example

To illustrate the one-step lookahead procedure, we present a simple 1d example similar to the one given in (Bect et al. 2012), where we aim to emulate the limit state function

$$g(x) = 1 - ((0.4x - 0.3)^2 + \exp(-11.534|x|^{1.95}) + \exp(-5(x - 0.8)^2)). \tag{46}$$

We assume that $g(x)$ can be evaluated at any $x \in \mathbb{R}$ without error, but that function evaluations are expensive. We will let $\xi(x)$ be the probabilistic surrogate in the form of a Gaussian process, where we use a prior mean $\mu(x) = -0.5$ together with a Matérn 5/2 covariance function with fixed kernel variance $\sigma_c^2 = 0.1$ and length scale $l = 0.5$.

We assume that X follows Normal distribution with mean $\mu_X = -0.5$ and standard deviation $\sigma_X = 0.2$, and our goal is to estimate $\alpha(g) = P(g(X) \leq 0)$ using only a small number of evaluations of $g(\cdot)$. The set of decisions is therefore $\mathbb{D} = \cup_x \{\text{evaluate } g(x)\}$ with respective outcomes $o(x) = g(x)$, and a predictive model for outcomes given as $\delta(x) = \xi(x)$.

Using a large number of samples of $g(X)$ we estimate $\alpha(g) \approx 0.0234$, and we will consider this as the 'true' failure probability for comparison.

We initiate ξ by evaluating $g(x)$ at $x = \mu_X$. For subsequent function evaluations, we minimize the expected variance in the failure probability. I.e. we minimize the acquisition function $J_{1,k}$ given in (13) with $\lambda \equiv 1$. For comparison we also evaluate $J_{2,k}$ and $J_{3,k}$, and in this example it seems that all three acquisition functions would perform equally well. Figure 4 shows ξ_k and the corresponding three acquisition functions for the first few experiments, and Fig. 5 shows how $\alpha(\xi_k)$ evolves before converging after $k = 3$ iterations.

6.2 Example 2: A 3 layer hierarchical model with 7d input

In this example we consider the structural reliability benchmark problem given as problem RP38 in (Rozsas and Slobbe 2019). Here, $\mathbf{x} = (x_1, \dots, x_7) \in \mathbb{X} = \mathbb{R}^7$, and the limit state function $g(\mathbf{x})$ can be written in terms of intermediate variables as follows:

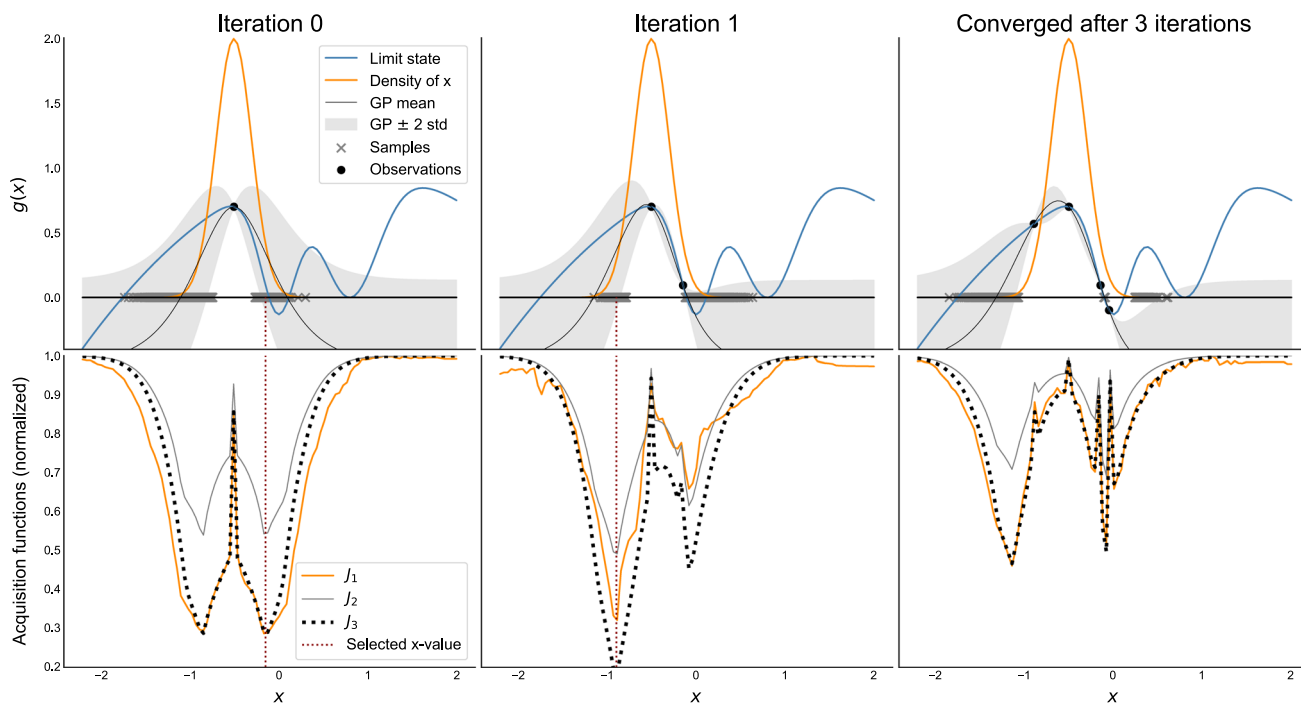


Fig. 4 (Example 1) The top row shows the true limit state function $g(x)$, the probability density of X , and the mean ± 2 standard deviations of the GP ξ_k for $k = 0, 1$ and 3 . The samples indicated with \times on

the x -axis are used in the importance sampling estimates of $J_{1,k}$, $J_{2,k}$ and $J_{3,k}$ that are shown in the bottom row

$$\begin{aligned}
 y_1(\mathbf{x}) &= \frac{x_1 x_2^3}{2c_4 x_3^3}, & y_2(\mathbf{x}) &= \frac{x_4^2}{c_2}, \\
 y_3(\mathbf{x}) &= -4x_5 x_6 x_7^2 + x_4(x_6 + 4x_5 + 2x_6 x_7), \\
 y_4(\mathbf{x}) &= x_4 x_5(x_4 + x_6 + 2x_6 x_7), \\
 z_1(\mathbf{y}) &= \frac{c_4 y_1 y_2}{c_3}, & g(\mathbf{y}, z_1) &= 1 - \frac{c_2 c_3 z_1 + c_4 y_1 y_3}{c_1 y_4},
 \end{aligned}$$

where c_1, c_2, c_3 and c_4 are constants: $c_1 = 15.59 \cdot 10^4$, $c_2 = 6 \cdot 10^4$, $c_3 = 2 \cdot 10^5$, $c_4 = 1 \cdot 10^6$.

Figure 6 shows a graphical representation of how $g(\mathbf{x})$ depends on the intermediate variables z_1, y_1, y_3 and y_4 . We will assume that the functions $y_2(\mathbf{x})$ and $z_1(\mathbf{y})$ will require probabilistic surrogates, where $y_2(\mathbf{x})$ and $z_1(\mathbf{y})$ can be evaluated without error for any input \mathbf{x} and \mathbf{y} . We will also assume that there is no difference in the cost associated with evaluating y_2 or z_1 , and our goal is to estimate the failure probability $\alpha(g)$ while keeping the total number of function evaluations of $y_2(\mathbf{x})$ and $z_1(\mathbf{y})$ as small as possible. Note that the effective domain of y_2 is 1-dimensional and the effective domain of z_1 is 2-dimensional. Hence, using surrogates for y_2 and z_1 should be much more efficient than building a single surrogate for g using samples $g(\mathbf{x}_i)$.

As for the random variable $\mathbf{X} = (X_1, \dots, X_7)$, we assume that all X_i 's are independent and normally distributed, $X_i \sim \mathcal{N}(\mu_i, \sigma_i)$, with means $\mu_1 = 350$, $\mu_2 = 50.8$, $\mu_3 = 3.81$,

$\mu_4 = 173$, $\mu_5 = 9.38$, $\mu_6 = 33.1$, $\mu_7 = 0.036$, and standard deviation $\sigma_i = 0.1\mu_i$. The 'true' failure probability we aim to estimate is $\alpha(g) \approx 8.1 \cdot 10^{-3}$.

Assuming y_2 and z_1 are expensive to evaluate, we introduce two Matérn 5/2 GP surrogates, \tilde{y}_2 and \tilde{z}_1 . The initial kernel parameters are $(\sigma_c^2 = 0.03, l = 20)$ and $(\sigma_c^2 = 2, l = [0.5, 0.5])$ for \tilde{y}_2 and \tilde{z}_1 , respectively. These parameters may be updated by maximum likelihood estimation, but not until a few observations (resp. 2 and 5) have been made. We know that large values of y_2 or z_1 will result in poor structural performance (small $g(\mathbf{x})$), so we initiate the GP models with conservative prior means of $\mu(\mathbf{x}) = 1$ for \tilde{y}_2 and $\mu(\mathbf{y}) = 5$ for \tilde{z}_1 . Both models are initially updated with one observation each, $\tilde{y}_2(\mu_4) = y_2^0$ and $\tilde{z}_1(y_1^0, y_2^0) = z_1^0$ for $y_1^0 = y_1(\mu_1, \mu_2, \mu_3)$, $y_2^0 = y_2(\mu_4)$ and $z_1^0 = z_1(y_1^0, y_2^0)$.

In this example, we would then define $\xi(\mathbf{x}) = g(y_1, \tilde{z}_1, \tilde{y}_2, y_3, y_4)$. With respect to \tilde{z}_1 , there is a set of possible decisions for uncertainty reduction, namely $\mathbb{D} = \cup_{y_1, y_2} \{\text{evaluate } z_1(y_1, y_2)\}$, with a corresponding set of observations $\mathbb{O} = \cup_{y_1, y_2} \{z_1(y_1, y_2)\}$, and a predictive model $\delta(y_1, y_2) = \tilde{z}_1(y_1, y_2)$. Similarly, we obtain a set of decisions, outcomes and a predictive model for \tilde{y}_2 , and we can update \mathbb{D}, \mathbb{O} and $\delta(d)$ accordingly.

Convergence was reached at iteration $k = 10$, after 2 additional evaluations of y_2 and 8 additional evaluations of z_1 . Figure 7 shows the updated surrogate models, $\tilde{y}_2|I_k$ and

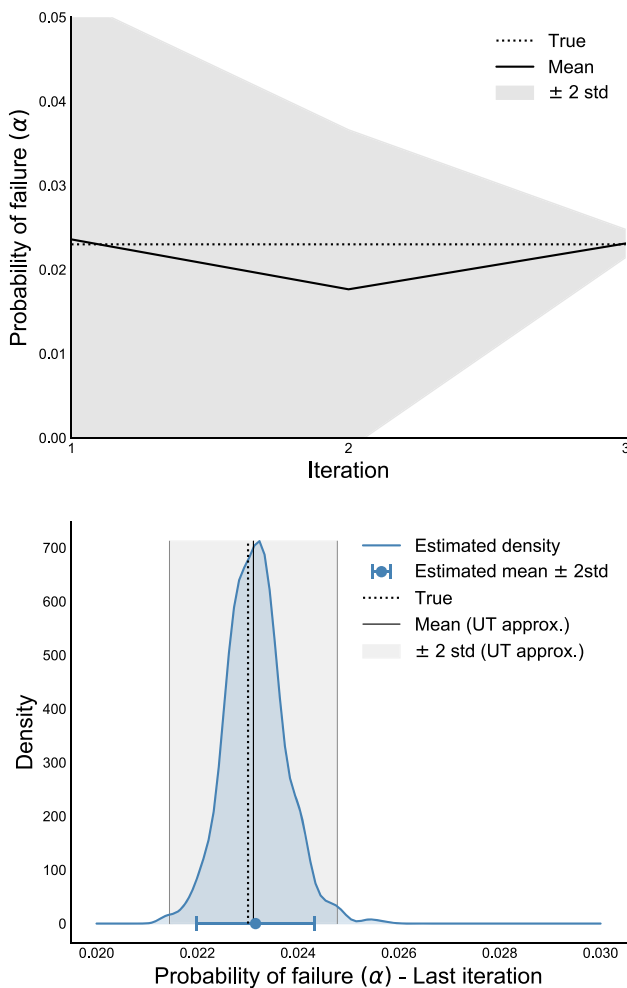


Fig. 5 (Example 1) Top: Mean \pm 2 standard deviations of α_k after k iterations, as computed using the approximation described in Sect. 4. Bottom: The distribution of α_k at the final iteration $k = 3$, estimated from a double-loop Monte Carlo (i.e. by sampling from $\alpha(\xi_k)$ without any approximation)




Fig. 6 (Example 2) Hierarchical representation of $g(\mathbf{x})$. We assume that the intermediate variables $y_2(\mathbf{x})$ and $z_1(\mathbf{y})$ are expensive to evaluate

$\tilde{z}_1|I_k$ for $k = 10$, and Fig. 8 shows how $\alpha(\xi_k)$ evolves with each iteration. At each iteration, the next experiment was decided by minimizing the acquisition function $J_{3,k}$ with respect to updating each of the two surrogate models.

6.3 Example 3: The 4 branch system

Here we consider the 'four branch system', a classical 2D benchmark problem given by the limit state

$$g(\mathbf{x}) = \min \left\{ \begin{array}{l} 3 + 0.1(x_1 - x_2)^2 - (x_1 + x_2)\sqrt{2}; \\ 3 + 0.1(x_1 - x_2)^2 + (x_1 + x_2)\sqrt{2}; \\ (x_1 - x_2) + 6\sqrt{2}; \\ (x_2 - x_1) + 6\sqrt{2} \end{array} \right\}, \quad (47)$$

and where x_1 and x_2 are independent standard normal variables. In this example we will not write (47) as an hierarchical model, in order to compare our method with other alternatives that are tailored to non-hierarchical setting. We therefore let $\xi(\mathbf{x})$ be a Gaussian process surrogate of $g(\mathbf{x})$, constructed from observations $(\mathbf{x}_i, g(\mathbf{x}_i))$. For the initial 'conservative' Gaussian process we select a prior mean of -1 , a Matérn 5/2 kernel with parameters of $(\sigma_c = 1, l = 3)$, and condition on the initial observation $(\mathbf{0}, g(\mathbf{0}))$.

According to Huang et al. (2017), the method called AK-MCS developed by Echard et al. (2011) is considered a typical and mature approach, and should therefore be a suitable candidate for comparison. In addition, Echard et al. (2011) also provide the results from using a number of other alternatives proposed in Schueremans and Gemert (2005). Table 3 gives a summary of the results from Echard et al. (2011), together with the those obtained using the approach presented in this paper.

Our results in Table 3 are obtained using Algorithm 1 with three different stopping criteria, $V_{max} = 0.1$, $V_{max} = 0.05$ and $V_{max} = 0.025$. Instead of point estimates we provide prediction intervals, which in this example contain the 'true' failure probability obtained with Monte Carlo in each scenario. From a practical perspective, even the estimates obtained using only 35 evaluations ($V_{max} = 0.1$) of (47) seems acceptable. If we were to use the mean + 2 standard deviations as a conservative estimate, the relative error with respect to the 'true' failure probability is still less than 3%. After an additional 30 iterations, this number drops to 0.65%. Hence, our approach performs well with respect to the alternatives considered in (Echard et al. 2011; Schueremans and Gemert 2005). It should also be noted that the Directional Sampling alternative in Table 3 is a method that is especially suitable for the specific 'radial' type of limit state surfaces as considered here, and at this level of performance is not expected in general.

Optimization was performed using the approximate acquisition function $\hat{J}_{3,k}$, and Fig. 9 shows how the sequence of

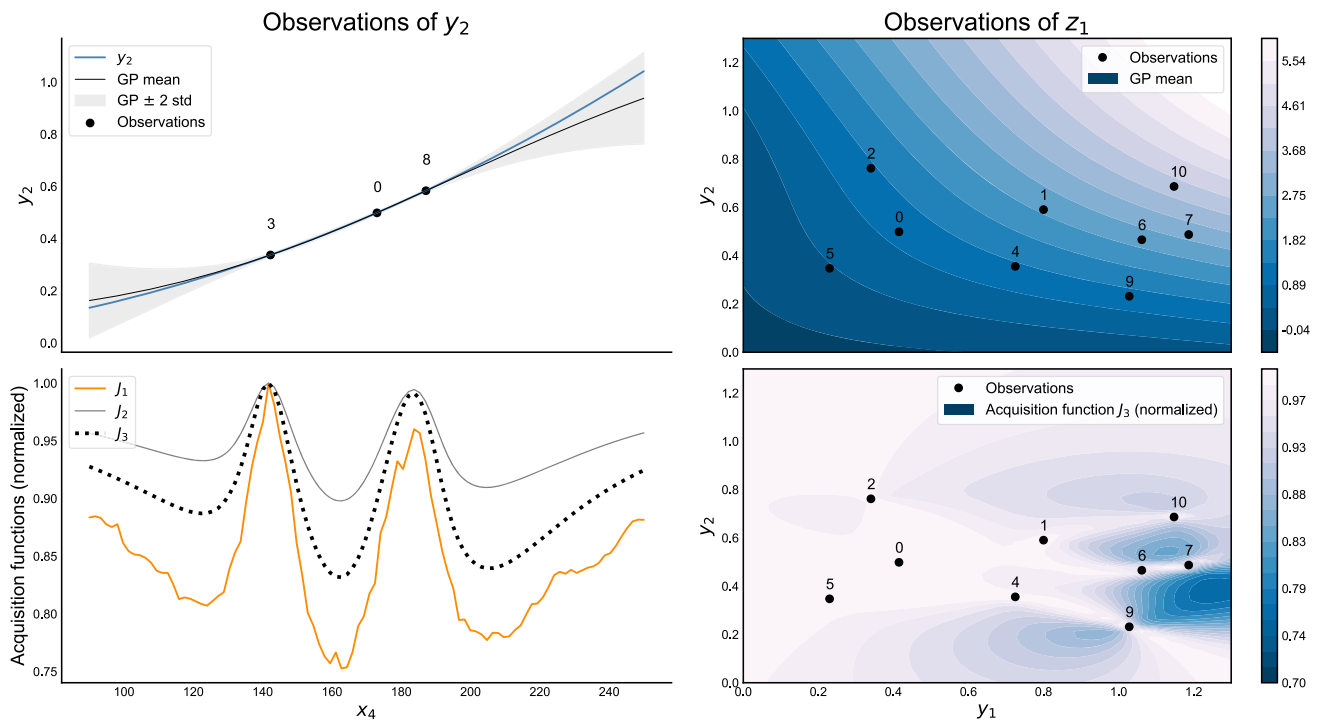


Fig. 7 (Example 2) The top row shows the GP models \tilde{y}_2 and \tilde{z}_1 with respect to P_k for $k = 10$. The number above each observations is the iteration index k , and convergence was obtained after 2 evaluations of y_2 and 8 evaluations of z_1 . The final acquisition functions are shown in the bottom row

Table 3 (Example 3) Table 2 from Echard et al. (2011), where we have appended the method from this paper (UT-MCIS) in the bottom row. The reported failure probabilities (\hat{p}_f) are the estimated mean ± 2 standard deviations of $\alpha(\xi_k)$ for $k = 35$ (stopped at $\hat{V}_k \leq 0.1$), $k = 48$ (stopped at $\hat{V}_k \leq 0.05$), and $k = 65$ (stopped at $\hat{V}_k \leq 0.025$)

Method	N_{call}	$\hat{p}_f \times 10^3$
Monte Carlo	10^6	4.416
AK-MCS+U	126	4.416
AK-MCS+EFF	124	4.412
Directional Sampling (DS)	52	4.5
DS + Response Surface	1745	5.0
DS + Spline	145	2.4
DS + Neural Network	165	4.1
Importance Sampling (IS)	1469	4.9
DS + Response Surface	1375	4.5
IS + Spline	428	4.5
IS + Neural Network	52	5.7
UT-MCIS ($V_{\text{max}} = 2.5\%$)	65	(4.347–4.444)
UT-MCIS ($V_{\text{max}} = 5\%$)	48	(4.288–4.470)
UT-MCIS ($V_{\text{max}} = 10\%$)	35	(4.163–4.547)

observations are located with respect to the failure set $g = 0$. The resulting sequence of failure probabilities after each iteration is illustrated in Fig. 10.

6.4 Example 4: Corroded pipeline example

To give an example of a scenario where there are different types of experiments, we consider a probabilistic model which is recommended for engineering assessment of offshore pipelines with corrosion (DNV GL 2017). The failure mode under consideration is where a pipeline bursts, when the pipeline’s ability to withstand the high internal pressure has been reduced as a consequence of corrosion.

6.4.1 The structural reliability model

Figure 11 shows a graphical representation of the structural reliability model. Here, a steel pipeline is characterised by the outer diameter (D [mm]), the wall thickness (t [mm]) and the ultimate tensile strength (s [MPa]). In this example we let $D = 800$, $t \sim \mathcal{N}(\mu = 20, \text{cov} = 0.03)$, and $s \sim \mathcal{N}(\mu = 545, \text{cov} = 0.06)$, where cov is the coefficient of variation (standard deviation / mean).

The pipeline contains a rectangular shaped defect with a given depth (d [mm]) and length (l [mm]), where $l \sim \mathcal{N}(\mu = 200, \sigma^2 = 1.49)$ and where d will be inferred from observations. Given a pipeline (D, t, s) with a defect (d, l), we can determine the pipeline’s pressure resistance capacity (the maximum differential pressure the pipeline can withstand before bursting). We let p_{FE} [MPa] denote the capacity coming from a Finite Element simulation of the physical phe-

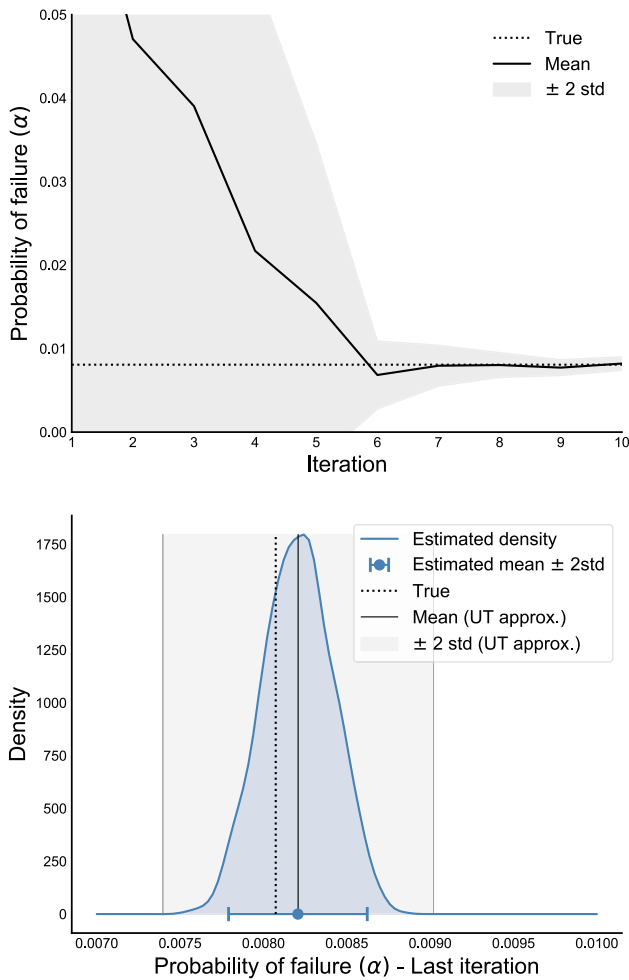


Fig. 8 (Example 2) Top: Mean \pm 2 standard deviations of α_k after k iterations, as computed using the approximation described in Sect. 4. Bottom: The distribution of α_k at the final iteration $k = 10$, estimated from a double-loop Monte Carlo (i.e. by sampling from $\alpha(\xi_k)$ without any approximation)

nomenon. From the theoretical capacity p_{FE} , we model the true pipeline capacity as $p_c = X_m \cdot p_{FE}$, where X_m is the **model discrepancy**, $X_m \sim \mathcal{N}(\mu_m, \sigma_m^2)$. For simplicity we have assumed that X_m does not depend on the type of pipeline and defect, and we will also assume that $\sigma_m = 0.1$, where only the mean μ_m will be inferred from observations of the form p_c/p_{FE} . Finally, the pressure **load** (in MPa) is modelled as a Gumbel distribution with mean 15.75 and standard deviation 0.4725. The limit state representing the transition to failure is then given as $g = p_c - p_d$.

6.4.2 Different types of decisions

We consider the following three types of decisions

1. *Defect measurement* We assume that unbiased measurements of the relative depth d/t can be obtained.




Fig. 9 (Example 3) The limit state (47) together with the expected failure surface $E[\xi_{65}]$ and the 65 observations collected before convergence at $\hat{V}_{65} < 0.025$. The proposal distribution $q_{\mathbf{X}}$ used for importance sampling is a mixture of Gaussian random variables centered at the four design points (x) as described in Appendix B. The pruned samples shown in the figure are mostly located around $E[\xi_{65}] = 0$ and in other regions where the level set $\xi_{65} = 0$ is uncertain

The measurements come with additive Gaussian noise, $\epsilon \sim \mathcal{N}(0, \sigma_{d/t}^2)$, and we will assume that three types of inspection are available, corresponding to $\sigma_{d/t} = 0.02, 0.04$ and 0.08 .

2. *Computer experiment* Evaluate p_{FE} at some deterministic input (D, t, s, d, l) .
3. *Lab experiment* Obtain one observation of X_m .

In order to generate synthetic data for this experiment, we assume that the true defect depth is $d = 0.3t = 6$ mm and that $\mu_m = 1.0$. Instead of running a full Finite Element simulation to obtain p_{FE} , we will make use of the simplified capacity equation in (DNV GL 2017), in which case

$$p_{FE} = 1.05 \frac{2ts}{D-t} \frac{1-d/t}{1-d/t/Q}, \quad Q = \sqrt{1 + 0.31 \frac{l^2}{Dt}}$$

6.4.3 Results

To define the initial model ξ_0 we need a prior specification over the epistemic quantities d, μ_m and p_{FE} . We let d be a priori normal with mean 0.5 and standard deviation 0.15, and μ_m normal with mean 1.0 and standard deviation 0.1. Consequently, the posteriors of d and μ_m (and also X_m) given any number of observations are all normal. The function p_{FE} is replaced by a GP surrogate with prior mean $\mu = -10$ and $\sigma_c = 10, l = [1, 1, 1, 1]$ Matérn 5/2 parameters, which we initiate using a single observations at the expected value of the input.

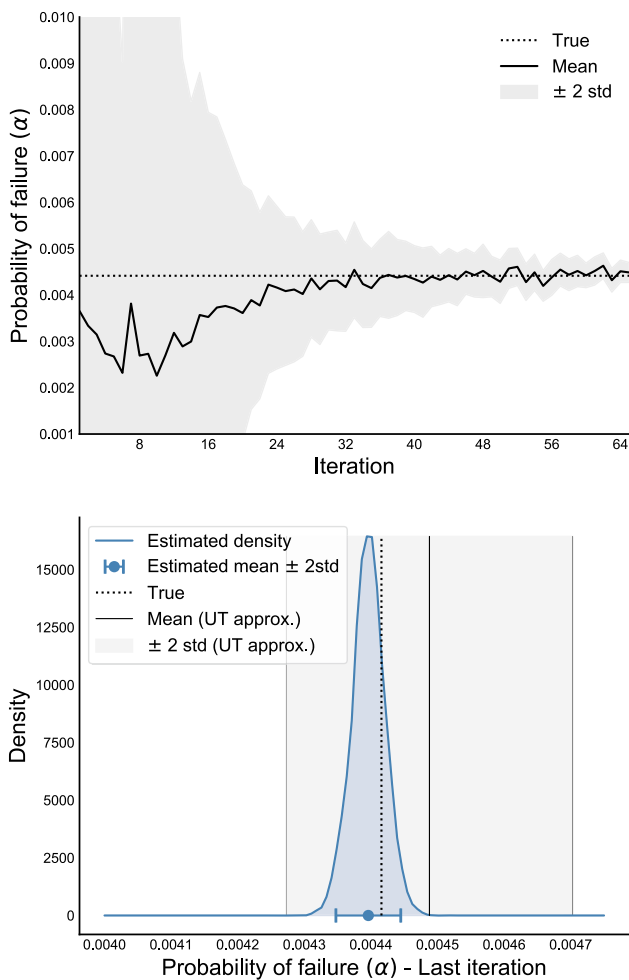


Fig. 10 (Example 3) Top: Mean ± 2 standard deviations of α_k after k iterations, as computed using the approximation described in Sect. 4. Bottom: The distribution of α_k at the final iteration $k = 65$, estimated from a double-loop Monte Carlo (i.e. by sampling from $\alpha(\xi_k)$ without any approximation)




Fig. 11 (Example 4) Graphical representation of the corroded pipeline structural reliability model. The shaded nodes d , p_{FE} and μ_m have associated epistemic uncertainty that can be reduced through experiments

We assume that the computer experiments are cheap compared to the lab experiments, and that the direct measurements of d/t is most expensive. To reflect these varying costs, we specify the acquisition function

$$\hat{J}_{i,k}(d) = c(d) \frac{\hat{E}_{k,d}[\hat{H}_{i,k+1}]}{\hat{H}_{i,k}}, \tag{48}$$

where $c(d)$ is the cost of a given decision. (Note that in (48) the variable d refers to a *decision*, but for the remaining part of this example d will only refer to the *defect depth*). In (48) we have normalized the expected future measure of residual uncertainty with the current, which gives an estimate of the expected improvement given a certain decision. The numerical values representing difference in costs is given by $c = 1$ for computer experiments, $c = 1.1$ for lab experiments, and $c = 1.11, 1.12, 1.13$ for measurements of d/t with accuracy $\sigma_{d/t} = 0.08, 0.04$ and 0.02 , respectively.

In structural reliability analysis, the objective is not always to obtain an estimate of the failure probability that is as accurate as possible. A relevant problem in practice is to determine whether a structure satisfies some prescribed target reliability level α_{target} . In this example, we aim to either confirm that the failure probability is less than the target $\alpha_{target} = 10^{-3}$ (in which case we can continue operations as normal), or to detect with confidence that the target is exceeded (and we have to intervene). For this purpose we intend to stop the iterative procedure if the difference between the expected and target failure probability is at least 4 standard deviations. In addition to the standard stopping criterion for convergence (45), we therefore introduce the stopping criterion

$$|\hat{E}[\hat{\alpha}_k] - \alpha_{target}| < 4\sqrt{\hat{H}_{1,k}}. \tag{49}$$

Figure 12 shows how the UT-MCIS approximation of the failure probability evolves throughout 100 iterations. We have made use of $i = 3$ in (48) as we found the corresponding acquisition surface for p_{FE} smoother than the alternative $i = 1$, and hence easier to minimize numerically. The stopping criterion (49) is reached after $k = 25$ iterations, and Fig. 13 shows the corresponding posteriors of the relative defect depth d/t and the model discrepancy X_m .

Throughout the examples in this paper we have initiated GP surrogate models using a single observation at the expected input. A different approach that is often found in practical applications is to initiate the GP surrogate with a space-filling design. A very common alternative is to make use of a Latin Hypercube sample (LHS), of size no more than $10 \times$ the input dimension (although the appropriate number of samples naturally depends on how nonlinear the response is expected to be, see e.g. Loepky et al. 2009).

Table 4 shows a summary of the results from running this example with and without an initial design consisting of 10

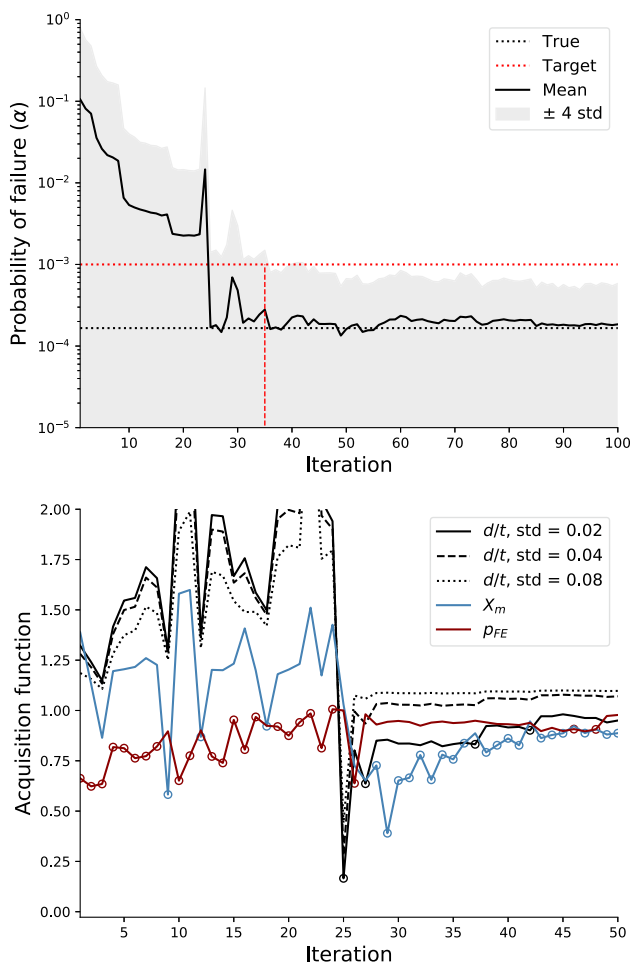


Fig. 12 (Example 4) Top: Mean ± 4 standard deviations of α_k after k iterations, as computed using the approximation described in Sect. 4. The stopping criterion (49) is reached after 25 iterations. Bottom: The acquisition functions (48) for each type of experiment during the first 50 iterations

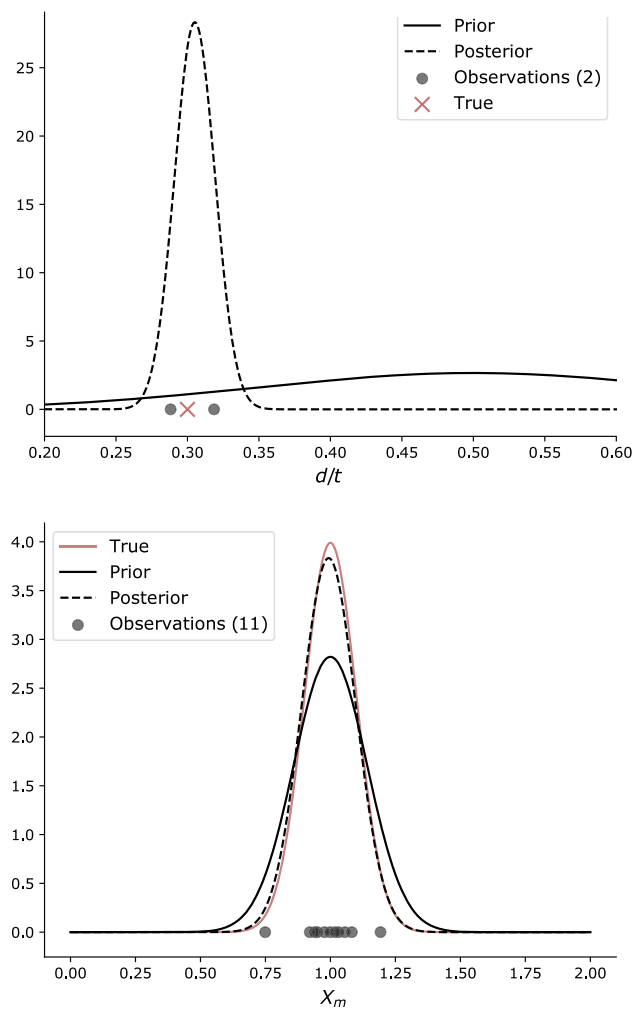


Fig. 13 (Example 4) The posterior distributions of d/t and X_m when the stopping criterion (49) is reached at $k = 25$

Table 4 (Example 4) Averages over 100 runs, using 1 versus 10 initial observations of p_{FE}

Initial Design	Stop at Target	Cov of α_k (\hat{V}_k)	Number of observations		
			p_{FE}	X_m	d/t
$E[X]$	Yes	1.39	23 + 1	10	2
	No	0.63	46 + 1	47	7
LHS 10	Yes	1.37	12 + 10	8	2
	No	0.90	45 + 10	48	7

LHS samples. For this example it does not seem to make any significant difference, but we see why the stopping criterion (49) is useful, as on average we can conclude that the failure probability is below the target value after around 30-40 iterations.

We leave this numerical experiment with an important remark, which is that specifying an appropriate cost in (48) can be difficult. If for instance the cost related to a measurement of d/t is set very high, then the decision to measure d/t will never be taken. In this example, it is not possible to reach the stopping criterion given in (49) without at least one such measurement, and hence, the one-step lookahead strategy will keep requesting measurements of X_m and evaluations of p_{FE} indefinitely, accumulating a potentially infinite cost. This is indeed a drawback of the one-step lookahead strategy. Note that a full dynamic programming implementation is typically not feasible in practice as it is too computationally expensive, and this may also be the case for implementations looking only a few more steps ahead. An idea for dealing with this is to study the problem via reinforcement learning instead. This is a work in progress.

7 Concluding remarks

We have presented a general formulation of the Bayesian optimal experimental design problem based on separation of aleatory randomness associated with a physical system, and the epistemic uncertainty that we wish to reduce through experimentation. The effectiveness of a design strategy is evaluated through a *measure of residual uncertainty*, and efficient approximation of this quantity is crucial if we want to apply algorithms that search for an optimal strategy. We make use of a pruned importance sampling scheme for subsequent estimation of (typically small) failure probabilities for a given epistemic realization, combined with the unscented transform epistemic uncertainty propagation. In our numerical experiments, we made use of a rather naive implementation of the unscented transform, in the sense that the number of sigma-points is very low, and that these are determined a priori with a deterministic procedure. Since the alternative by Merwe and Wan (2003) produced satisfactory results in all of our numerical examples, no further consideration was made with respect to alternative methods for sigma-point selection. From applications to Kalman filtering, it has been observed that this version of the unscented transform has a tendency to over-estimate the variance, which is something we notice also in our experiments.

For the application we consider in this paper, we emphasize that the unscented transform is used as a *proxy* for the measure of residual uncertainty to be used in optimization, as a numerically efficient alternative that should be proportional to the true objective. Hence, we view the unscented transform as a tool to find the best decision or strategy, where we get the possibility of exploring many decisions approximately rather than a few exactly. Once an optimal strategy is found, we estimate the corresponding measure of residual uncertainty using a pure Monte Carlo alternative which is exact in the limit. We note that for global optimization of acquisition functions, we have used a combination of random sampling and gradient based local optimization. With this procedure, an optimization objective given by $H_{3,k}$ (and also $H_{2,k}$) is generally more suitable than $H_{1,k}$, as it is less susceptible to noise coming from Monte Carlo estimation (see for instance Fig. 7). On the other hand, $H_{1,k}$ has a natural interpretation (the variance of the failure probability), and is therefore a better measure for evaluating convergence, or for early stopping as discussed in Sect. 6.4.

In Example 4 (Sect. 6.4), we briefly discussed the common alternative of applying a space-filling design, and we observed that starting with an initial LHS design did not make any significant difference, when the remaining design was determined using the one-step lookahead strategy. Similarly, we may compare the one-step lookahead strategy to a naive LHS design, to investigate how useful it is to apply this strategy at all. The number of experiments needed to

converge at $V_{\max} = 0.05$ in Examples 1, 2 and 3 in Sect. 6 was 4, 10 and 48. If we instead were to use a (maximin) LHS design over the set of inputs with non-negligible probability density, the expected number of experiments needed to reach $V_{\max} \leq 0.05$ is 30, 150 and 400. Hence, in these examples, the number of experiments is reduced by roughly a factor of 10 by applying the one-step lookahead strategy instead of a space-filling design.

Although we focus on the estimation of a *failure probability* in this paper, many of the main ideas we present should also be applicable for other estimation objectives using models where a hierarchical structure can be utilized. For instance, when α_k is some other quantity of interest depending on the random variable $g(\mathbf{X})$, not necessarily given by an indicator function as in (1). In general, the problem we consider in this paper is estimating the volume of the excursion set $\{\mathbf{x} \in \mathbb{X} | g(\mathbf{x}) \leq 0\}$, under some specified measure on $\mathbb{X} \subseteq \mathbb{R}^m$. For the specific applications considered in this paper, we have assumed that an isoprobabilistic transformation of \mathbf{X} to a standard normal variable is available, which is often the case in structural reliability models. We make use of this assumption only to apply some well known techniques for failure probability estimation, but note that other alternatives, for instance the one presented in Appendix B.3, can be used instead.

There are several ways to improve the methodology presented in this paper. For instance, other alternatives of the unscented transform could be applied, see for instance Menegaz et al. (2015), or the parameters determining the set of sigma-points used in this paper could be optimized as in (Turner and Rasmussen 2010).

As seen in Sect. 6.4, the one-step lookahead/myopic strategy, can make it impossible to reach the stopping criterion of the algorithm. As mentioned, a way to avoid this problem is by looking at the whole dynamic programming formulation (11). However, this formulation suffers from the curse of dimensionality. Since the myopic formulation corresponds to truncating the sum in the dynamic programming formulation (11) to only one term, it is of interest to study methods where more terms of the sum are included (multi-step look ahead). How much better do the estimations get by including an extra term, and how much does the computation time increase? Is it possible to determine an optimal choice of truncation where we weigh accuracy and computation time against one another? Different ways of finding approximate solutions to the complete dynamic programming problem has been the focus of much research within areas such as operations research, optimal control and reinforcement learning, and trying out some of these alternatives is certainly interesting avenue for further research.

Another interesting topic worth investigating is how the numerical examples in this paper compare to the case where we estimate the *buffered failure probability* instead of the

classical failure probability. Buffered failure probabilities were introduced by Rockafellar and Royset (2010) as an alternative to classical failure probabilities in order to take into account the tail distribution of the performance function. See Dahl and Huseby (2019) for an application of this concept to structural reliability analysis.

One may also discuss whether using heuristic optimization objectives chosen to approximate the variance is reasonable. By essentially focusing on minimizing the variance of the failure probability, we say that all deviations from the true value is equally bad. In reality, overestimating the failure probability can be costly, but is not nearly as problematic as underestimating the failure probability. Because of this, the variance may not be the most appropriate measure of risk. It would be interesting to also derive heuristic optimization objectives based on approximating other risk measures.

These questions are of interest, but beyond the scope of the current paper, and the topics are left for future research.

Acknowledgements This work has been supported by Grant 276282 from the Norwegian Research Council and DNV GL Group Technology and Research (Christian Agrell), and by Project 29989 from the Research Council of Norway as part of the SCROLLER project (Kristina Rognlien Dahl). We would also like to thank the reviewers for helpful comments and feedback.

Funding Open access funding provided by University of Oslo (incl Oslo University Hospital).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendices

A Gaussian process surrogate models

Here we briefly review the Gaussian process (GP) surrogate model in its canonical form, for Bayesian nonparametric function estimation. For a broader overview of the relevant theory see e.g. Rasmussen and Williams (2006). For applications related to uncertainty quantification (UQ) dealing with deterministic computer simulations, Kennedy and O’Hagan (2001) is a classical reference.

Let $f : \mathbb{X} \rightarrow \mathbb{R}$ denote a function that we want to estimate, and assume that a set of k observations $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_k, \mathbf{y}_k)$

have been made. For instance, evaluating $f(\mathbf{x})$ could correspond to running a deterministic (and time consuming) computer simulation, in which case noiseless observations, $\mathbf{y}_i = f(\mathbf{x}_i)$, can be obtained. Alternatively, $f(\mathbf{x}_i)$ could correspond to some physical experiment, resulting in a noise perturbed observation \mathbf{y}_i . A GP surrogate model ξ of f is a tool to make inference about the value of $f(\mathbf{x}^*)$ for any new input $\mathbf{x}^* \in \mathbb{X}$, conditioned on the set of observations $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_k, \mathbf{y}_k)$.

A Gaussian process ξ indexed by some set \mathbb{X} is defined by the property that for any finite subset $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ of \mathbb{X} , $(\xi(\mathbf{x}_1), \dots, \xi(\mathbf{x}_N))$ is an N -dimensional Gaussian random variable. We will view ξ as a Gaussian distribution over real-valued functions defined on \mathbb{X} (such as $f(\mathbf{x})$). Here \mathbb{X} can be arbitrary but typically \mathbb{X} is a subset of \mathbb{R}^n . The GP ξ is uniquely defined by its mean function $\mu(\mathbf{x}) = E[\xi(\mathbf{x})]$ and covariance function $c(\mathbf{x}, \mathbf{x}') = E[(\xi(\mathbf{x}) - \mu(\mathbf{x}))(\xi(\mathbf{x}') - \mu(\mathbf{x}'))]$. Hence, any function $\mu : \mathbb{X} \rightarrow \mathbb{R}$ paired with a positive semidefinite function $c : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ defines a GP, which we will denote $\xi \sim \mathcal{GP}(\mu, c)$.

Let $X = (\mathbf{x}_1, \dots, \mathbf{x}_k), Y = (\mathbf{y}_1, \dots, \mathbf{y}_k)$ denote the observations and assume that \mathbf{y}_i comes with additive Gaussian noise, $\mathbf{y}_i = f(\mathbf{x}_i) + \epsilon_i$ where ϵ_i are i.i.d. zero-mean Gaussian with common variance σ^2 . In this scenario, the conditional process $\xi|X, Y$ is still a Gaussian process. In particular, if $X^* = (\mathbf{x}_1^*, \dots, \mathbf{x}_m^*)$ contains m new input locations in \mathbb{X} , then the distribution of $\xi^* = \xi(X^*) = (\xi(\mathbf{x}_1^*), \dots, \xi(\mathbf{x}_m^*))$ given the observations X, Y is Gaussian with the following mean

$$E[\xi^*|X, Y] = \mu(X^*) + c(X^*, X)[c(X, X) + \sigma^2 I_m]^{-1}(Y - \mu(X)), \tag{50}$$

and covariance

$$\text{Cov}(\xi^*|X, Y) = c(X^*, X^*) - c(X^*, X)[c(X, X) + \sigma^2 I_m]^{-1}c(X, X^*)^T. \tag{51}$$

Here $\mu(X^*)$ and $\mu(X)$ are vectors with elements $\mu(\mathbf{x}_i^*)$ and $\mu(\mathbf{x}_i)$ respectively, I_m is the $m \times m$ identity matrix, and $c(X^*, X^*), c(X^*, X)$ and $c(X, X)$ have elements $c(X^*, X^*)_{i,j} = c(\mathbf{x}_i^*, \mathbf{x}_j^*), c(X^*, X)_{i,j} = c(\mathbf{x}_i^*, \mathbf{x}_j)$ and $c(X, X)_{i,j} = c(\mathbf{x}_i, \mathbf{x}_j)$.

For the scenario where observations are noiseless, $\mathbf{y}_i = f(\mathbf{x}_i)$, the distribution of $\xi^*|X, Y$ is obtained with $\sigma = 0$ in (50)–(51).

To define a GP prior $\xi \sim \mathcal{GP}(\mu, c)$ over functions $f : \mathbb{X} \rightarrow \mathbb{R}$, we need to specify the mean and covariance function. These are generally given as $\mu(\mathbf{x}|\theta)$ and $c(\mathbf{x}, \mathbf{x}'|\theta)$, conditioned on some parameter θ . An appropriate value for θ is usually found through maximum likelihood estimation or cross validation using the set of observations X, Y . A fully Bayesian approach could also be pursued, where the posterior

calculations typically involve Markov chain Monte Carlo as the formulation in (50)–(51) is not sufficient. In the numerical experiments presented in this paper, we have made use of a constant mean function and a Matérn 5/2 covariance function using plug-in hyperparameters $\theta = (\sigma_c, l_1, \dots, l_n)$ determined from maximum likelihood estimation. The Matérn 5/2 covariance function for $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$ is defined as

$$c(\mathbf{x}, \mathbf{x}') = \sigma_c^2 \left(1 + \sqrt{5}r + \frac{5}{3}r^2 \right) e^{-\sqrt{5}r},$$

$$r = \sqrt{\sum_{i=1}^n \left(\frac{x_i - x'_i}{l_i} \right)^2}. \tag{52}$$

B The sampling distribution $q_{\mathbf{x}}$

Here we present some further details on how the set of samples $\{\mathbf{x}_i, w_i\}$ in Sect. 4.3 can be generated. We start by reviewing some classical techniques from structural reliability analysis that are based on finding ‘important’ regions in \mathbb{X} . The sampling distribution $q_{\mathbf{x}}$ used in this paper is then defined in Section B.2. It is based on the assumption that \mathbf{X} can be transformed to a standard multivariate Gaussian variable \mathbf{U} , and that $q_{\mathbf{U}}$ can be constructed by solving a set of constrained optimization problems in \mathbf{U} -space. For the scenario where these assumptions do not hold, we present an alternative approach in Section B.3, which is based on a naive exploration of the \mathbb{X} -space. Although this will require evaluation of a larger set of samples of \mathbf{X} , no optimization is required and numerical implementation is straightforward.

B.1 Local approximations in SRA

In Sect. 4 we briefly discussed the challenges with estimation of the failure probability $\bar{\alpha}(g)$ in (1). A different alternative often used in structural reliability analysis, is to approximate the performance function $g(\mathbf{x})$ with a function \hat{g} where $\bar{\alpha}(\hat{g})$ can be computed analytically. In this scenario, it is convenient to transform \mathbf{X} to a standard normal variable \mathbf{U} . We will let

$$\mathbf{X} \xrightarrow{\mathcal{T}} \mathbf{U} \sim N(\mathbf{0}, \mathbf{I}) \tag{53}$$

denote an isoprobabilistic transformation, where $\mathbf{U} = \mathcal{T}(\mathbf{X})$ is multivariate standard Gaussian with $\dim(\mathbf{U}) = \dim(\mathbf{X})$. Note that for any univariate random variable X with CDF $F(X)$, a transformation of this type available as $\mathcal{T}(X) = \Phi^{-1}(F(X))$. The generalization to multivariate \mathbf{X} is the Rosenblatt transformation, where $U_i = \Phi^{-1}(F_i(\mathbf{X}_i | \mathbf{X}_1, \dots, \mathbf{X}_{i-1}))$. In structural reliability problems, it is often natural to define \mathbf{X} in terms of the marginal distributions and a copula, in which case the isoprobabilistic transformation (53) can be simplified. A common alternative

is to use a Gaussian copula, where (53) can be obtained using the Nataf transformation (Lebrun and Dutfoy 2009).

In the following we let $g(\mathbf{u})$ denote the function $g(\cdot)$ applied to $\mathbf{x} = \mathcal{T}^{-1}(\mathbf{u})$. Methods such as FORM (First Order Reliability Method) and SORM (Second Order Reliability Method) make use of local approximations in the form of a linear or quadratic surface fitted to $g(\mathbf{u}^*)$ at a certain point $\mathbf{u}^* \in \mathbb{R}^n$. This point \mathbf{u}^* is often called the *design point* or *most probable point* (MPP), and it is defined as

$$\mathbf{u}^* = \arg \min_{\mathbf{u} \in \mathbb{R}^n} \{ \|\mathbf{u}\| \mid g(\mathbf{u}) \leq 0 \}. \tag{54}$$

Observe that if $\hat{g}(\mathbf{u})$ is the first-order Taylor approximation of $g(\mathbf{u})$ at \mathbf{u}^* , i.e. $\hat{g}(\mathbf{u}) = g(\mathbf{u}^*) + \nabla_{\mathbf{u}} g(\mathbf{u}^*)(\mathbf{u} - \mathbf{u}^*)$, then $\bar{\alpha}(\hat{g}) = \Phi(-\|\mathbf{u}^*\|)$, and this is an upper bound on the failure probability if the failure set is convex in \mathbf{U} -space.

In Sect. 4.3 we discussed the importance sampling estimate of the failure probability given some proposal distribution q . A natural candidate is to let q be a distribution centered around the design point, \mathbf{u}^* in \mathbf{U} -space or $\mathbf{x}^* = \mathcal{T}(\mathbf{u}^*)$ in \mathbf{X} -space. The alternative where the estimation is performed in \mathbf{U} -space with $q_{\mathbf{U}}(\mathbf{u}) = \phi(\mathbf{u} + \mathbf{u}^*)$ is often used in practice. For a more detailed discussion around this kind of sampling, the local approximations and structural reliability analysis in more general, see for instance (Madsen et al. 2006) or (Huang et al. 2017).

The constrained optimization problem (54) plays an important role in structural reliability analysis. Although any general-purpose algorithm can be used, customized algorithms that take advantage of the special form of the objective function are recommended. Various alternatives have been developed for this purpose, see for instance (Gong and Yi 2011) and the references therein. For the applications in this paper we have made use of the iHL-RF method from (Zhang and Der Kiureghian 1995).

B.2 The design point mixture

We observe first that a solution to (54) is not necessarily unique, and also that multiple *local* minima may exist when the performance function is nonlinear. Most algorithms designed to solve (54) numerically start with some initial guess \mathbf{u}_0 , and take iterative steps until a minimum is obtained. To reduce the risk of overestimating $\|\mathbf{u}^*\|$, multiple restarts with different (possibly randomized) initial guesses \mathbf{u}_0 is often applied.

Given a finite-dimensional approximation of a performance function $\hat{\xi}(\mathbf{x}, \mathbf{E})$, we want to find a proposal distribution q that is appropriate for a range of different realizations \mathbf{e} of \mathbf{E} . In particular, if $\{(v_j, \mathbf{e}_j) \mid j = 1, \dots, M\}$ is the set of sigma-points for \mathbf{E} as introduced in Sect. 4.2, we want a set of samples from q to be applicable for estimation of $\alpha(\hat{\xi}(\mathbf{x}, \mathbf{e}_j))$ for any $1 \leq j \leq M$.

For any \mathbf{e}_j , we will let $\mathbf{u}_{1,j}^*, \dots, \mathbf{u}_{N,j}^*$ denote N design points in U -space corresponding to $\hat{\xi}(\mathbf{x}, \mathbf{e}_j)$, obtained using randomized initialization. (Note that for methods such as iHL-RF, it is also reasonable to use $\mathbf{u}_{i,j}^*$ as an initial guess in the search for $\mathbf{u}_{i,j+1}^*$). We then define \mathcal{Q} as the equal-weighted Gaussian mixture of the NM random variables $\mathcal{Q}_{i,j} = \mathbf{U}_{i,j} + \mathbf{u}_{i,j}^*$, where $\mathbf{U}_{i,j}$ are i.i.d. standard multivariate Gaussian. Sampling from \mathcal{Q} is then straightforward, and importance sampling estimates can be obtained in the U -space using $p_U(\mathbf{u}) = \phi(\mathbf{u})$ and $q_U(\mathbf{u}) = \frac{1}{NM} \sum_{i,j} \phi(\mathbf{u} - \mathbf{u}_{i,j}^*)$, where ϕ is the multivariate standard normal density.

B.3 A simple alternative

The sampling strategy presented in Sect. 4.3 is based on (1) generating a set of samples that should "cover relevant locations" in the input space \mathbb{X} , and (2) *prune* the set of samples using a threshold on the measure of insignificance (21).

The "relevant locations" in the first step is typically somewhere in the "tail" of the distribution of \mathbf{X} , where also the (uncertain) performance function $\hat{\xi}_k(\mathbf{x})$ may be close to zero. In Section B.2 we made use of importance sampling around design points, which is a common technique in structural reliability analysis. As a simple alternative, we can let q be any distribution from which it is easy to generate samples covering the effective support of $p_{\mathbf{X}}$ (i.e. a bounded domain where \mathbf{X} lies with probability ≈ 1). For instance, assuming U is n -dimensional standard normal (e.g. $U = \mathcal{T}(\mathbf{X})$ if the isoprobabilistic transformation is still applicable), we could let q be a uniform density on the hypercube $[-b, b]$ where $b = \Phi^{-1}(1 - p_{min})$ for some absolute lower bound on the failure probability p_{min} .

Because the initial set of N samples from q will be reduced to a fixed number of n samples after the pruning step, this is a viable alternative. However, in order to obtain similar importance sampling variances [see (29)] as with the method in Section B.2, the initial number of samples N (and hence the number of evaluations of the pruning criterion $\eta(\mathbf{x})$) will have to be larger.

C Selecting sigma-points for the unscented transform

Here we briefly review the method for sigma-point selection by Merwe (2004) and present the sigma-points used for the numerical experiments in Sect. 6.

According to Labbe (2014), research and industry have mostly settled on the version published in (Merwe 2004). Here, the sigma-points are given as a function of the mean and covariance matrix of the input variable, together with three real-valued parameters α , β and κ . In the case where U is a

standardized n -dimensional random variable with $E[U] = \mathbf{0}$ and $E[U^2] = I$, we obtain $2n + 1$ points \mathbf{u}_i are as follows

$$\begin{aligned} \mathbf{u}_0 &= \mathbf{0}, \\ \mathbf{u}_i &= \alpha\sqrt{n + \kappa} \mathbf{v}_i, \\ \mathbf{u}_{i+n} &= -\mathbf{u}_i, \end{aligned}$$

for $i = 1, \dots, n$ where $\mathbf{v}_i = (0, \dots, 1, \dots, 0)$ is the standard unit vector in \mathbb{R}^n . Two different sets of weights are used with this procedure, one for the mean and one for the covariance in (17). We denote these v_i^m and v_i^c respectively, and they are given as

$$\begin{aligned} v_0^m &= 1 - \frac{n}{\alpha^2(n + \kappa)}, \quad v_0^c = v_0^m + 1 - \alpha^2 + \beta, \\ v_i^m &= v_i^c = \frac{1}{2\alpha^2(n + \kappa)} \text{ for } i = 1, \dots, 2n. \end{aligned}$$

For Gaussian distributions, it is often recommended to set $\beta = 2$, $\kappa = 3 - n$ and let $\alpha \in (0, 1]$. In the numerical examples presented in this paper we have used this set of parameters with $\alpha = 0.9$.

References

- Bect, J., Ginsbourger, D., Li, L., Picheny, V., Vazquez, E.: Sequential design of computer experiments for the estimation of a probability of failure. *Stat. Comput.* **22**(3), 773–793 (2012)
- Bect, J., Bachoc, F., Ginsbourger, D.: A supermartingale approach to Gaussian process based sequential design of experiments. *Bernoulli* (2019)
- Bichon, B., Eldred, M., Swiler, L., Mahadevan, S., McFarland, J.: Efficient global reliability analysis for nonlinear implicit performance functions. *AIAA J.* **46**, 2459–2468 (2008)
- Dahl, K., Mohammed, S.E., Øksendal, B., Røse, E.E.: Optimal control of systems with noisy memory and BSDEs with Malliavin derivatives. *J. Funct. Anal.* **271**(2), 289–329 (2016)
- Dahl, K.R., Huseby, A.B.: Buffered environmental contours. *Safety and Reliability - Safe Societies in a Changing World Proceedings of ESREL 2018* (2019)
- Defourny, B., Ernst, D., Wehenkel, L.: Multistage stochastic programming: A scenario tree based approach to planning under uncertainty. *Decis. Theory Model. Appl. Artif. Intell.: Concepts Solut.* (2011)
- Der Kiureghian, A., Ditlevsen, O.: Aleatory or epistemic? does it matter? *Struct. Saf.* **31**(2), 105–112 (2009)
- DNV, G.L.: Recommended Practice: Corroded Pipelines DNVGL-RP-F101. DNV GL, Høvik, Norway (2017)
- Echard, B., Gayton, N., Lemaire, M.: AK-MCS: An active learning reliability method combining Kriging and Monte Carlo simulation. *Struct. Saf.* **33**(2), 145–154 (2011)
- Fernandez, G., Park, C., Kim, N., Haftka, R.: Review of multi-fidelity models. [arXiv:1609.07196v3](https://arxiv.org/abs/1609.07196v3) (2017)
- Goldstein, M., Wooff, D.: *Bayes Linear Statistics: Theory and Methods*. John Wiley & Sons, New York (2007)
- Gong, J.X., Yi, P.: A robust iterative algorithm for structural reliability analysis. *Struct. Multidiscip. Optim.* **43**, 519–527 (2011)

- Huan, X., Marzouk, Y.: Sequential bayesian optimal experimental design via approximate dynamic programming. [arXiv:1604.08320v1](https://arxiv.org/abs/1604.08320v1) (2016)
- Huang, C., ELHami, A., Radi, B.: Overview of Structural Reliability Analysis Methods-Part I, II, p. 17. III, Incertitudes et fiabilité des systèmes multiphysiques (2017)
- Jian, W., Zhili, S., Qiang, Y., Rui, L.: Two accuracy measures of the Kriging model for structural reliability analysis. *Reliab. Eng. Syst. Safety* **167**, 494–505 (2017)
- Jones, M., Goldstein, M., Jonathan, P., Randell, D.: Bayes linear analysis of risks in sequential optimal design problems. *Electron. J. Stat.* **12**, 4002–4031 (2018)
- Julier, S., Uhlmann, J.: Unscented filtering and nonlinear estimation. *Proc. IEEE* **92**, 401–422 (2004)
- Kennedy, M.C., O'Hagan, A.: Bayesian calibration of computer models. *J. Royal Stat. Soc.: Ser. B (Stat. Methodol.)* **63**(3), 425–464 (2001)
- Kyzyurova, K.N., Berger, J.O., Wolpert, R.L.: Coupling computer models through linking their statistical emulators. *SIAM/ASA J Uncertain Quantif.* **6**(3), 1151–1171 (2018)
- Labbe, R.: Kalman and Bayesian Filters in Python. In: <https://github.com/rlabbe/Kalman-and-Bayesian-Filters-in-Python>, GitHub eBook (2014)
- Lebrun, R., Dutfoy, A.: A generalization of the Nataf transformation to distributions with elliptical copula. *Probab. Eng. Mech.* **24**(2), 172–178 (2009)
- Loeppky, J., Sacks, J., Welch, W.: Choosing the sample size of a computer experiment: A practical guide. *Technometrics* **51**, 366–376 (2009)
- Madsen, H., Krenk, S., Lind, N.: *Methods of Structural Safety*. Dover Civil and Mechanical Engineering Series. Dover Publications, New York (2006)
- Menegaz, H.M.T., Ishihara, J.Y., Borges, G.A., Vargas, A.N.: A systematization of the unscented kalman filter theory. *IEEE Trans. Autom. Control* **60**(10), 2583–2598 (2015)
- Merwe, R.: *Sigma-Point Kalman Filters for Probabilistic Inference in Dynamic State-Space Models*. PhD thesis, OGI School of Science and Engineering (2004)
- Merwe, R., Wan, E.: Sigma-point kalman filters for probabilistic inference in dynamic state-space models. *Proceedings of the Workshop on Advances in Machine Learning* (2003)
- Perrin, G.: Active learning surrogate models for the conception of systems with multiple failure modes. *Reliab. Eng. Syst. Safety* **149**, 130–136 (2016)
- Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge (2006)
- Rockafellar, R.T., Royset, J.O.: On buffered failure probability in design and optimization of structures. *Reliab. Eng. Syst. Safety* **95**(5), 499–510 (2010)
- Rozsas, A., Slobbe, A.: Repository and Black-box Reliability Challenge 2019. <https://gitlab.com/rozsasarpi/rprepo/> (2019)
- Schueremans, L., Gemert, D.V.: Benefit of splines and neural networks in simulation based structural reliability analysis. *Struct. Saf.* **27**(3), 246–261 (2005)
- Sun, Z., Wang, J., Li, R., Tong, C.: LIF: A new Kriging based learning function and its application to structural reliability analysis. *Reliab. Eng. Syst. Safety* **157**, 152–165 (2017)
- Turner, R., Rasmussen, CE.: Model based learning of sigma points in unscented kalman filtering. In: *2010 IEEE International Workshop on Machine Learning for Signal Processing*, pp 178–183 (2010)
- Uhlmann, J.: *Dynamic map building and localization : New theoretical foundations*. PhD thesis, University of Oxford (1995)
- Wang, L.: *Karhunen-Loève expansions and their applications*. PhD thesis, London School of Economics and Political Science (2008)
- Zhang, Y., Der Kiureghian, A.: *Two Improved Algorithms for Reliability Analysis*, pp. 297–304. Springer, US, Boston, MA (1995)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.