Quantile-Based Reliability-Constrained Optimization of a Nonlinear Absorber for Passive Aeroelastic Control under Design Variability

Nassim RAZAALY[a,*], Benjamin CHOUVION[b]

[a]ISAE-ENSMA, Chasseneuil-du-Poitou, FRANCE
[b]CREA, École de l'Air et de l'Espace, F-13660 Salon-de-Provence, FRANCE

Abstract

This paper presents a robust optimization methodology formulated as a quantile minimization problem, that incorporates reliability-based constraints and aleatory variability in the design variables. The proposed approach relies on kriging-based surrogate models constructed in a coupled input space and combined with a sparse adaptive sampling strategy. The optimization method is tested on simple analytical functions to validate and illustrate the methodology. It is then applied to the robust design of a nonlinear absorber, used as a passive means to mitigate flutter instability and reduce the amplitude of limit-cycle oscillations in a two-dimensional airfoil exhibiting geometric nonlinearities. Numerical results demonstrate the effectiveness of the proposed framework while requiring only a limited number of numerical solver evaluations. The proposed methodology is therefore particularly well suited for practical engineering applications, where computational resources are limited and uncertainties must be accounted in the early design process, to develop more robust and reliable systems.

Keywords: Aeroelasticity, Optimization under uncertainties, Quantile, Kriging, Nonlinear tuned vibration absorber, Flutter instability

---

1. Introduction

The interaction between inertial, elastic, and aerodynamic forces that govern the dynamic behavior of aeronautical structures can lead to undesirable phenomena, ranging from performance degradation to the onset of self-sustained oscillations associated with aeroelastic instabilities such as flutter [1]. Aeroelastic flutter control thus remains a critical challenge. Over the years, a wide range of mitigation strategies has been developped, spanning from active, semi-active or purely passive approaches [2–5]. Among the different passive methods, nonlinear absorbers [6–10] appear to be a promising technology, as they are effective over a broader frequency range than a perfectly tuned linear absorber. To combine the benefits of both linear and nonlinear control mechanisms, the concept of the nonlinear tuned vibration absorber (NLTVA) has been introduced [11]. This hybrid approach couples a linear tuned mass damper, responsible for increasing the critical flutter speed with nonlinear elements designed to limit post-critical oscillation amplitudes. Experimental validations have further confirmed the potential of passive flutter control, including studies involving for instance linear absorbers [12] or a nonlinear absorber located within the flap hinge [13]. Optimisation studies of the control system consistently show that the optimal absorber rapidly loses effectiveness when the primary system departs from its nominal design (see for instance [11] or [12]). Consequently, the optimal parameters are strongly dependent on the nominal design, emphasising the practical relevance of Uncertainty Quantification (UQ).

In this context, the study of uncertainties and their effects on the aeroelastic response has gained significant importance in recent years [14]. The characterization of these sources, whether *aleatory* (e.g. manufacturing tolerances, material scatter, gusty winds) or *epistemic* (e.g. modelling assumptions), has been discussed by Pettit [15] and Dai and Yang [16]. Most works have focused on aleatory variability, modelled using random variables. Such uncertainties can be treated via non-probabilistic methods that identify worst-case scenarios within prescribed parameter bounds [16],

---

Email address: nassim.razaaly@ensma.fr

or through probabilistic approaches that propagate known input distributions to obtain response statistics, using techniques such as probabilistic collocation [17], stochastic spectral projection [18], or the Polynomial Chaos Expansion (PCE) method [19]. In contrast, Denimal and Sinou [20] proposed a detailed stability analysis of an airfoil model that accounts for both types of uncertainties through the development of a hybrid surrogate modelling framework combining Kriging and PCE.

Given the critical influence of uncertainties on aeroelastic performance and safety, there is a clear need for design optimization frameworks that explicitly account for sources of uncertainty arising from variability in design parameters, loading conditions, material properties, or modeling assumptions. Within this context, *optimization under uncertainty* seeks design solutions that remain both performant and safe in the presence of variability. Various approaches exist for formulating an optimization problem under uncertainties for a mechanical system. Conventional robust optimization aims at minimizing the expected value of a Quantity of Interest (QoI), but provides limited control over performance variability. To address this issue, several principal strategies have emerged. One can cite for instance the mean-variance methods, such as optimizing a linear combination of the mean and standard deviation, or minimizing both mean and standard deviation in a multi-objective context [21], also known as the Taguchi method. However, within this kind of approaches, the estimation of second-order statistics typically involves an additional computational burden. Another class of methods directly targets tail performance, leading to quantile-based objective functions [22, 23].

Figure 1, extracted from [23], illustrates the different types of optima that can be obtained depending on the chosen formulation. The QoI is represented as a function of the scalar design variable x. The solid curve represents the deterministic QoI in the absence of uncertainty, whereas environmental and operational variability induce a dispersion of the QoI, depicted by the shaded regions. The deterministic optimum x_1^* provides good performance under *nominal conditions* (i.e., in the absence of uncertainty), but results in a highly variable QoI with a poor expected value when uncertainties are present, as indicated by the associated probability density function (PDF) shown in orange, implying severe performance degradation under off-design conditions. The design x_2^*, obtained by optimizing the mean performance, performs well on average while still exhibiting significant variability. By contrast, x_3^* minimizes the standard deviation, resulting in low variability at the expense of degraded performance. Finally, x_4^*, obtained by minimizing the 95% quantile, emphasizes performance over 95% of the realizations while neglecting worst-case scenarios.

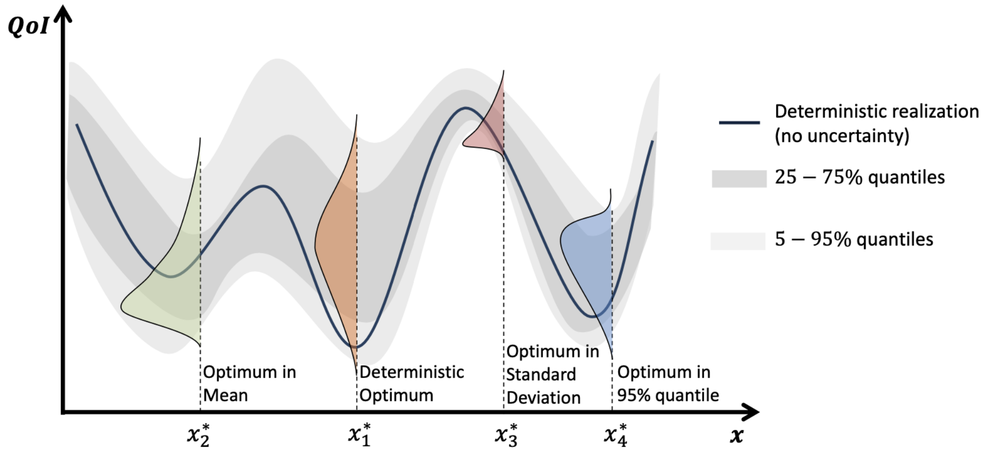

Figure 1: Comparison between deterministic optimization and several robust optimization strategies (from [23]). The colored areas represent the probability density functions associated with the different types of optima.

Other approaches attempt to trade off system cost with reliability, known as reliability-based design optimization (RBDO), for which a thorough review can be found in [24] or [25]. Both robust and reliability-based optimization frameworks often rely on nested, two-level strategies. An outer optimization loop explores the design space, while an inner loop evaluates statistical quantities such as failure probabilities, response moments, or output quantiles through uncertainty propagation or reliability analysis [26]. Among these methods are the Reliability Index Approach, the

Second-Order Reliability Method, and Monte Carlo Simulation [27]. These nested approaches can become computationally expensive, particularly when high-fidelity aeroelastic simulations are required.

To mitigate this computational burden, surrogate modeling techniques have been widely adopted. Surrogate-based strategies provide an efficient alternative by constructing approximations of the system response either in the stochastic space for each design vector [22, 28], or in a global coupled space that includes both design and environmental variables. The latter *coupled-space* approaches enable the reuse of surrogate models across multiple design configurations, further improving computational efficiency, though they are typically limited by the dimensionality of the problem [29].

In this context, the RBDO method proposed by [30] is particularly appealing: it addresses reliability-based constraints using a *coupled-space* surrogate modeling approach and allows the design vector to be affected by variability, albeit in the context of minimizing a *deterministic* objective function. On the other hand, the approach introduced in [29] enables the solution of a related problem in a coupled space with an objective function subject to variability, but assumes an *exact* (i.e., deterministic) design vector. The present work aims to fill this gap by extending the RBDO philosophy of [30] to a fully *robust* formulation.

The objective of this manuscript is therefore to propose an efficient mono-objective robust optimization methodology that simultaneously accounts for reliability-based constraints and design variables affected by *aleatory* variability, with both the objective and constraints formulated in terms of quantiles. Its applicability is then demonstrated on a *realistic* aeroelastic problem: the robust design of a NLTVA to control post-critical oscillation amplitudes and the flutter instability of a two-degrees-of-freedom airfoil model [11] subjected to a quasi-steady aerodynamic flow. The methodology relies on Gaussian Process (GP)–based surrogate modeling [31] of the objective and constraint functions in the coupled space, incorporating both design and external variables from an initial quasi-random design of experiments (DoE). A specific parallel active-learning strategy is proposed to adaptively enrich the metamodels, focusing on promising design vectors and improving the accuracy of quantile estimates.

The remainder of this manuscript is organized as follows. Section 2 presents the formal statement of the problem of interest. Section 3 recalls the theoretical background on surrogate-based optimization and reliability analysis. Section 4 introduces the proposed numerical framework. Section 5 is devoted to numerical applications: the methodology is first validated and illustrated using several analytical test cases, and is then applied to the robust design of an NLTVA used as a passive means to mitigate aeroelastic effects in an airfoil system.

2. Formulation of the robust optimization problem

This section describes the formulation of the optimization under uncertainties problem tackled in this work, and introduces key notations to facilitate reading in subsequent sections. The original contribution of this paper, that is the methodology used to solve this problem, will be detailed in Section 4.

2.1. Problem of interest

We recall below a classical RBDO approach, consisting in minimizing a deterministic cost under probabilistic constraints, which translates into system reliability [30, 32]:

$$
\begin{aligned}
&\min_{\boldsymbol{d}\in\mathcal{D}} \quad f(\boldsymbol{d}), \quad \text{subject to:} \\
&\mathbb{P}\left(g_k(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z}) \leq 0\right) \leq \alpha_{g_k}, \quad k = 1, \dots, p,
\end{aligned}
\tag{1}
$$

where f is the cost function, depending on the d_x-dimensional design vector $\boldsymbol{d} \in \mathcal{D} \subset \mathbb{R}^{\mathrm{d}_x}$ defining some design parameters, and g_k are the constraint functions, also known as limit-state function in the reliability analysis literature, for which a negative value indicates system failure. These depend on random vector variables classified into two groups:

- the *random design variables* $\boldsymbol{X} \sim \mathcal{P}_{\boldsymbol{X}(\cdot;\boldsymbol{d})}$, simply denoted $\mathbf{X}(\mathbf{d})$ in the following, where $\mathcal{P}_{\mathbf{X}(\cdot\,;\,\mathbf{d})}$ is the probability distribution of $\mathbf{X}$ parameterized by the design vector $\mathbf{d}$;

- the d_z-dimensional *environmental* or *external* random variables $\mathbf{Z} \sim \mathcal{P}_{\mathbf{Z}}$, which influence the limit-state without being controlled during the optimization.

Finally, $\mathbb{P}$ refers to the probability operator induced by the variability of the random vector $(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})$. The specificity of RBDO lies in the fact that constraints are expressed in terms of failure probabilities $\mathbb{P}(g_k(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})) \leq 0)$, which must remain below a prescribed threshold $\alpha_{g_k} \in]0, 1[$.

In the following, the probabilistic constraints (Eq. (1)) are equivalently replaced by quantile-based constraints [30], while the deterministic cost function is reformulated using a quantile-based objective [22, 23], resulting in the robust optimization problem below:

$$\min_{\boldsymbol{d} \in \mathcal{D}} \quad Q_{\alpha_f}[f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})], \quad \text{subject to the quantile constraints:} \tag{2}$$
$$Q_{\alpha_{g_k}}[g_k(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})] \leq 0, \quad k = 1, \ldots, p,$$

where $Q_{\alpha_{g_k}}[g_k(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})] = \inf\{q \in \mathbb{R} : \mathbb{P}(g_k(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z}) \leq q) \geq \alpha_{g_k}\}$ denotes the quantile of the constraint function g_k of level α_{g_k}, and $\alpha_f \in]0, 1[$ is the quantile level associated with the cost function f.

This formulation is motivated by two key considerations:

- **Efficiency in surrogate-based optimization.** Quantile estimation is tightly connected to reliability analysis, where adaptive sampling focuses on the classification boundary defined by the limit-state function (see Subsection 3.2). As a result, the surrogate model only needs to be accurate in regions close to the quantile boundary, thereby reducing the overall computational burden. In contrast, formulations involving mean and variance typically require accurate surrogate predictions across the entire domain, which significantly increases the number of required evaluations [33].

- **Engineering interpretability.** Quantile-based constraints and objectives naturally align with practical design specifications. For instance, enforcing that a constraint holds with high probability (e.g. 90%) directly translates into a quantile threshold, which is often how safety margins or performance tolerances are expressed in engineering practice. This interpretability is further illustrated in the next subsection.

2.2. Illustrative example

To conceptualize the above mathematical formalism, consider the case of a simplified aeroelastic system in which a non-linear absorber is coupled to an airfoil. This example serves as a reduced representation of the primary application investigated in this work, which will be analyzed in detail in Section 5.

The motivation is to maximize the flutter speed V_f while ensuring that the amplitude of the limit cycle oscillation (LCO) in plunge, denoted y_{LCO}, remains below a prescribed threshold y^0_{LCO}. The objective and constraint functions are therefore defined as $f = -V_f$ and $g_1 = y_{LCO} - y^0_{LCO}$. For now, V_f and y_{LCO} are assumed to be outputs of a deterministic simulator that takes as inputs two physical parameters: the absorber mass m and the airfoil moment of inertia I, both affected by variability. This optimization is therefore expressed in an uncertain framework, and we adopt the formulation presented in Eq. (2).

The goal is to optimize the absorber design, by finding the best value of m, while explicitly accounting for the variability affecting the airfoil parameter I. In this framework, the optimization is performed with respect to the design variable represented as $\boldsymbol{d} = m$ and defined over the domain $\mathcal{D} = [0, m_{\text{max}}]$. Both sources of uncertainties are modeled as independent Gaussian variables. The design-related uncertainty is described by the random variable $\boldsymbol{X}(\boldsymbol{d}) = \boldsymbol{X}(m) \sim \mathcal{N}(m, \sigma_m^2)$, with mean equal to the design value m, and fixed variance. The *environmental* uncertainty, corresponding to the variability in the airfoil's moment of inertia I, is represented by the random variable $Z \sim \mathcal{N}(\mu_I, \sigma_I^2)$.

To choose the quantile levels, it is assumed that the goal is to maximize a worst-case flutter speed (e.g., its 5% quantile), which corresponds to the 95% quantile of the objective $f = -V_f$. Simultaneously, it is required that the plunge LCO amplitude remains below the prescribed threshold with at least 90% probability. This leads to the following quantile-based robust optimization problem:

$$\min_{\boldsymbol{d} \in \mathcal{D}} \quad Q_{0.95}[-V_f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})], \quad \text{subject to the quantile constraint:} \tag{3}$$
$$Q_{0.9}[y_{LCO}(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z}) - y^0_{LCO}] \leq 0.$$

This example highlights the practical meaning of the proposed quantile-based robust optimization formulation, allowing decision making under uncertainty while focusing the design effort on a specific subset of variables: the absorber mass in this case. To further formalize this framework, the input space and probabilistic modeling assumptions are detailed in the next subsection.

2.3. Input space and probabilistic modeling

Notations.

- The two sources of uncertainty are concatenated into a d_θ-dimensional random vector $(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})$, where $(\boldsymbol{x}, \boldsymbol{z})$ denotes a sample from this distribution. For simplicity, the random vector or its realizations will be referred to as $\theta(\boldsymbol{d})$.

- The objective and constraint functions f and g_k are defined over the input space Ω and are assumed to be the outputs of expensive-to-evaluate numerical solvers, treated as black-box deterministic models.

- An input vector is indicated by $\boldsymbol{\theta} = (\boldsymbol{x}, \boldsymbol{z}) \in \Omega = \Omega_x \times \Omega_z \subset \mathbb{R}^{\mathrm{d}_\theta}$, where $\mathcal{D} \subset \Omega_x \subset \mathbb{R}^{\mathrm{d}_x}$ defines the design space, and $\Omega_z \subset \mathbb{R}^{\mathrm{d}_z}$ denotes the domain of environmental uncertainties. The total input dimension is $\mathrm{d}_\theta = \mathrm{d}_x + \mathrm{d}_z$.

Assumptions.

- The input and design spaces are assumed to have a tensorized structure: $\Omega = \prod_{i=1}^{\mathrm{d}_\theta} [\theta_{min}^{(i)}, \theta_{max}^{(i)}]$ and $\mathcal{D} = \prod_{i=1}^{\mathrm{d}_x} [d_{min}^{(i)}, d_{max}^{(i)}]$.

- All random vectors are assumed to have mutually independent components, modeled using univariate distributions such as uniform, Gaussian, or truncated Gaussian, the latter being used in the numerical examples of this work. For example, $\boldsymbol{X}(\boldsymbol{d})$ can be modeled as $\boldsymbol{X}(\boldsymbol{d}) = (X^{(i)}(d^{(i)}))_{i=1,\dots,\mathrm{d}_x}$, where each component is defined by $X^{(i)}(d^{(i)}) \sim \mathcal{N}(d^{(i)}, \sigma_x^{2(i)})$. In the exemple given in Subsection 2.2, each component of the design-dependent random vector has its mean equal to the corresponding design variable, with a fixed variance $\sigma_x^{2(i)}$.

- To ensure that the generated samples of $\boldsymbol{\theta}(\boldsymbol{d})$ remain within the admissible domain Ω, a *bounding operation* is applied. Specifically, any component of a sample $\boldsymbol{\theta}$ falling outside the prescribed bounds is clamped to the nearest bound, such that $\theta^{(i)} \in [\theta_{\mathrm{min}}^{(i)}, \theta_{\mathrm{max}}^{(i)}], \forall i \in [\![1, \mathrm{d}_\theta]\!]$. In the example of Subsection 2.2, the mass passed to the simulators f and g_1 must remain positive. Considering the extreme case where $\boldsymbol{d} = m = 0$, the design random vector $\boldsymbol{X}(\boldsymbol{d}) \sim \mathcal{N}(0, \sigma_m^2)$ is therefore mathematically equivalent to a positively truncated Gaussian with zero mean and variance σ_m^2, after the described procedure. This guarantees that all input values used in the simulation remain physically meaningful.

These assumptions are adopted in the numerical experiments for clarity, although the proposed framework remains general, assuming a sampler of $\boldsymbol{\theta}(\boldsymbol{d})$ consistent with the physical input space is available.

The present framework is designed for quantile levels $\alpha \in]\beta, 1-\beta[$, where $\beta \in [10^{-3}, 0.5]$. The estimation of *extreme quantiles*, corresponding to $\beta < 10^{-3}$, is not addressed here and is considered beyond the scope of this study.

3. Theoretical background on surrogate-based optimization and reliability analysis

This section provides a concise overview of key concepts in surrogate-assisted optimization and reliability analysis employed throughout this work, while the original methodology is detailed in Section 4. For clarity, a generic scalar-valued model $\mathrm{y} = \mathrm{F}(\mathbf{x})$ is considered, where $\mathbf{x} \in \mathbb{R}^{\mathrm{d}}$ denotes a d-dimensional input vector. This notation is used exclusively in this section.

The theory of Gaussian Process (GP) metamodels, also referred to as Kriging models, is briefly summarized in Subsection 3.1, as the foundation of all surrogate modeling techniques used in this study. In particular, we describe how to construct a GP-based surrogate model $\hat{\mathrm{F}}$ of a computationally expensive function F, from a dataset. Building upon this class of metamodels, Subsection 3.2 reviews the fundamentals of reliability analysis, including failure probability and quantile estimation using GP surrogates. A standard *adaptive sampling strategy* (also referred to as an infill criterion) is introduced, which iteratively refines the surrogate model by selecting new evaluation points that improve accuracy near the quantile boundary. This subsection adopts a probabilistic setting by introducing a random input vector $\mathbf{X}$. Subsection 3.3 concludes this section by outlining the key principles of Bayesian optimization, which relies on a GP surrogate $\hat{\mathrm{F}}$ to minimize the original function F over a predefined design space. The adaptive sampling strategy adopted for mono-objective constrained optimization is detailed, introducing additional notation for the constraint functions G_k.

3.1. Gaussian process metamodeling

Metamodeling (also known as supervised learning) aims to construct a *fast* emulator $\hat{\mathrm{F}}$ of a deterministic function F, based on a Design of Experiments (DoE), here a set of N input–output pairs $\mathcal{D} = \{(\mathbf{x}_i, \mathrm{y}_i)\}_i$, with $\mathbf{x}_i \in \mathbb{R}^{\mathrm{d}}$ and $\mathrm{y}_i = \mathrm{F}(\mathbf{x}_i)$ denoting respectively the input vector and the scalar output. The Kriging approach assumes that F is a realization of an underlying Gaussian Process [34], expressed as:

$$\mathrm{Y}(\mathbf{x}) = \boldsymbol{\beta}^T \mathbf{f}(\mathbf{x}) + \mathrm{Z}(\mathbf{x}), \tag{4}$$

where:

- $\boldsymbol{\beta}^T \mathbf{f}(\mathbf{x})$ is the deterministic part that defines the mean (or trend) of the Kriging. It is modeled as a linear combination of user-defined basis functions $\mathbf{f} = (\mathrm{f}_1, \ldots, \mathrm{f}_q)$, with regression coefficients $\boldsymbol{\beta} \in \mathbb{R}^q$ to be estimated from the data.
- $\mathrm{Z}(\mathbf{x})$ is a zero-mean stationary Gaussian process with a covariance between two points $\mathbf{x}$ and $\mathbf{x}'$ defined by $\mathrm{cov}(\mathrm{Z}(\mathbf{x}), \mathrm{Z}(\mathbf{x}')) = \sigma_{\mathrm{Z}}^2 \mathrm{R}(\mathbf{x}, \mathbf{x}')$, where σ_{Z}^2 denotes the variance of the process, and $\mathrm{R}(\mathbf{x}, \mathbf{x}')$ is the autocorrelation function.

Various approaches exist to define the autocorrelation function R [31, 34]. In this study, the anisotropic power-exponential function is adopted:

$$\mathrm{R}(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^{\mathrm{d}} \exp\left[-\frac{(\mathrm{x}_i - \mathrm{x}_i')^\kappa}{l_i}\right], \tag{5}$$

where $\mathbf{l} = (l_1, \ldots, l_{\mathrm{d}})$ denotes the vector of scaling parameters to be calibrated, and κ the power exponent.

To predict the response of F at a given unobserved location $\mathbf{x}$, the random variable $\mathrm{Y}(\mathbf{x})$ conditioned on the observation vector $\mathbf{y} = (\mathrm{y}_1, \ldots, \mathrm{y}_N)$. This conditional distribution, $\hat{\mathrm{F}}(\mathbf{x}) = \mathrm{Y}(\mathbf{x}) \mid \mathbf{y}$, is known to be Gaussian [34], with its mean and variance given by:

$$\mu_{\hat{\mathrm{F}}(\mathbf{x})} = \boldsymbol{\beta}^T \mathbf{f}(\mathbf{x}) + \mathbf{r}(\mathbf{x})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{F}\boldsymbol{\beta}) \tag{6}$$

$$\sigma_{\hat{\mathrm{F}}(\mathbf{x})}^2 = \sigma_{\mathrm{Z}}^2 \left[1 - \mathbf{r}(\mathbf{x})^T \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}) + \mathbf{u}(\mathbf{x})^T (\mathbf{F}^T \mathbf{R}^{-1} \mathbf{F})^{-1} \mathbf{u}(\mathbf{x})\right] \tag{7}$$

where $\mathbf{u}(\mathbf{x}) = \mathbf{F}^T \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}) - \mathbf{f}(\mathbf{x})$ and $\mathbf{r}_i(\mathbf{x}) = \mathrm{R}(\mathbf{x}, \mathbf{x}_i)$.

The Kriging parameters $\boldsymbol{\beta}, \sigma_{\mathrm{Z}}^2$ are estimated via the generalized least-squares solution of the underlying regression problem:

$$\boldsymbol{\beta} = (\mathbf{F}^T \mathbf{R}^{-1} \mathbf{F})^{-1} \mathbf{F}^T \mathbf{R}^{-1} \mathbf{y} \tag{8}$$

$$\sigma_{\mathrm{Z}}^2 = \frac{1}{N} (\mathbf{y} - \mathbf{F}\boldsymbol{\beta})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{F}\boldsymbol{\beta}), \tag{9}$$

where $\mathbf{R}_{ij} = \mathrm{R}(\mathbf{x}_i, \mathbf{x}_j)$ is the correlation matrix of the experimental design, and $\mathbf{F}_{ij} = \mathbf{f}_j(\mathbf{x}_i)$ is the basis function matrix of size $N \times q$. For conciseness, the dependence on the scaling parameters has been omitted.

The hyperparameters $\mathbf{l}$ are determined via maximum likelihood estimation [34]:

$$\hat{\mathbf{l}} = \arg\min_{\mathbf{l}} \sigma_{\mathrm{Z}}^2 (\det \mathbf{R})^{1/N} \tag{10}$$

To summarize, once the scaling parameters $\mathbf{l}$ are obtained by solving Eq. (10), the Kriging parameters $\boldsymbol{\beta}$ and σ_{Z}^2 are estimated using Eqs. (8)–(9). This enables the definition of the Kriging surrogate model—also referred to as the GP predictor—at a new location $\mathbf{x}$, denoted by:

$$\hat{\mathrm{F}}(\mathbf{x}) \sim \mathcal{N}(\mu_{\hat{\mathrm{F}}(\mathbf{x})}, \sigma_{\hat{\mathrm{F}}(\mathbf{x})}^2). \tag{11}$$

In other words, the surrogate model is inherently stochastic at a given point $\mathbf{x}$, although its *posterior* mean $\mu_{\hat{\mathrm{F}}(\mathbf{x})}$ is directly computable (Eq. (6)), and the associated prediction uncertainty is quantified by the *posterior* variance $\sigma_{\hat{\mathrm{F}}(\mathbf{x})}^2$, as given by Eq. (7), under a Gaussian distribution. More details can be found in [35]. In this work, we use the SMT implementation [36] with a power-exponential kernel and exponent $\kappa = 1.9$, in line with the work of [37].

3.2. *GP-based reliability analysis*

The objective of this subsection is to concisely recall how *reliability analysis* estimations can be performed using a precomputed predictor via Monte Carlo Sampling (MCS), and to present an active learning strategy for proposing additional input vectors to enrich the DoE, as a result enhancing the accuracy of these estimations. This approach follows the framework introduced in [38–40]. It assumes the availability of a GP predictor $\hat{\mathrm{F}}(\mathbf{x}) \sim \mathcal{N}(\mu_{\hat{\mathrm{F}}}(\mathbf{x}), \sigma^2_{\hat{\mathrm{F}}}(\mathbf{x}))$, constructed from a design of experiments (DoE) $\mathcal{D}$, as described in Subsection 3.1.

Given a d-dimensional random vector $\mathbf{X}$, consider the problem of estimating a *moderate* quantile[1] at level $\alpha \in]0,1[$ of $\mathrm{F}(\mathbf{X})$, defined as the real value q such that:

$$\alpha = \mathbb{P}_{\mathbf{X}}(\mathrm{F}(\mathbf{X}) < q), \tag{12}$$

where $\mathbb{P}_{\mathbf{X}}$ denotes the probability measure associated with the random vector $\mathbf{X}$. Evaluation of the quantile $q = Q_\alpha(\mathrm{F}(\mathbf{X}))$ is closely related to estimating the probability $\mathbb{P}_{\mathbf{X}}\,(\mathrm{F}(\mathbf{X}) < q)$, also known as *failure probability estimation*. Both problems fall under the broader category of reliability analysis. To simplify notation, q refers either to the quantile value itself or to the threshold appearing in the probability expression, depending on the context.

MCS Estimators. Both problems can be estimated using the *direct* Monte Carlo method, which relies on a large set of n samples $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ drawn from $\mathbf{X}$, for which the model F is *directly* evaluated. After sorting the outputs in ascending order $\{\mathrm{F}(\mathbf{x}_1), \ldots, \mathrm{F}(\mathbf{x}_n)\}$, such that $\mathrm{F}(\mathbf{x}_1) \leq \cdots \leq \mathrm{F}(\mathbf{x}_n)$, the quantile MCS estimator is obtained via linear interpolation as follows:

$$\hat{q}_{\text{direct}} = \mathrm{F}(\mathbf{x}_{\hat{l}}) + (\alpha(n-1) + 1 - \hat{l})(\mathrm{F}(\mathbf{x}_{\hat{l}+1}) - \mathrm{F}(\mathbf{x}_{\hat{l}})), \tag{13}$$

where $\hat{l} = \lfloor \alpha(n-1) + 1 \rfloor$ and $\lfloor \cdot \rfloor$ denotes the floor function.

The precision of the quantile estimator naturally depends on the Monte Carlo sample size n. A way to assess this accuracy is to use the same evaluations to estimate the associated failure probability $\mathrm{p} = \mathbb{P}_{\mathbf{X}}\,(\mathrm{F}(\mathbf{X}) < q) = \mathbb{E}_{\mathbf{X}}[\mathbb{1}_{\mathrm{F}-q<0}(\mathbf{x})]$, written as an expectation, where $q = \hat{q}_{direct}$ is the MCS quantile estimate. The corresponding MCS estimators are given below:

$$\hat{\mathrm{p}} = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\mathrm{F}-q<0}(\mathbf{x}_i), \tag{14}$$

$$\hat{\sigma}^2_{\hat{\mathrm{p}}} = \frac{\hat{\mathrm{p}}(1-\hat{\mathrm{p}})}{n-1}, \tag{15}$$

where $\hat{\mathrm{p}}$ is the estimator of the failure probability, and $\hat{\sigma}^2_{\hat{\mathrm{p}}}$ is the associated variance estimator. Note that the expression for $\hat{\mathrm{p}}$ reveals that the failure probability problem reduces to *correctly classifying* [38, 40] the sign of the so-called Limit-State Function (LSF) $\mathrm{F} - q$.

These estimators enable the construction of confidence intervals for p, expressed as $[\hat{\mathrm{p}} \pm k_p \hat{\sigma}_{\hat{\mathrm{p}}}]$, where, for example, $k_p = 1.96$ corresponds to a 95% confidence level. In practice, a commonly used normalized indicator of MCS convergence is the ratio of the estimated standard deviation to the estimated failure probability, referred to as the coefficient of variation (CoV):

$$\hat{\delta}_{\hat{\mathrm{p}}} = \frac{\hat{\sigma}_{\hat{\mathrm{p}}}}{\hat{\mathrm{p}}}. \tag{16}$$

This metric is used in the present work to assess the Monte Carlo convergence of the quantile estimator. As a guideline, a CoV of 1% is generally considered very accurate, 5% acceptable, and 0.01% extremely accurate [38, 40, 41].

To reduce the large number of potentially expensive *direct* evaluations of the function F, the latter is replaced by the GP surrogate $\hat{\mathrm{F}}$. The stochastic nature of $\hat{\mathrm{F}}$ allows for the estimation of both the quantile and its associated bounds.

---

[1] It is referred to as *moderate* because the methodology described in this subsection remains practical down to about 10^{-4}–10^{-5} with crude Monte Carlo using 10^{8}–10^{10} realizations. In the proposed robust optimization framework, Monte Carlo evaluations are limited to about 10^{3}–10^{4} samples per function call and reused across thousands of iterations, which restricts the approach to higher quantile levels.

By replacing F with its GP posterior mean $\mu_{\hat{F}}$, and considering the ordered set $\{\mu_{\hat{F}}(\mathbf{x}_1), \dots, \mu_{\hat{F}}(\mathbf{x}_n)\}$, a quantile estimate $\hat{q}$ based on MCS using the GP can be obtained. Lower and upper conservative bounds $\hat{q}^-$ and $\hat{q}^+$ for the quantile can also be computed by substituting F with $\mu_{\hat{F}} - k_q\sigma_{\hat{F}}$ and $\mu_{\hat{F}} + k_q\sigma_{\hat{F}}$, respectively. By analogy with Gaussian confidence intervals, $k_q = 3$ is set throughout this work to obtain conservative bounds, while noting that the quantile estimate itself does not necessarily follow a Gaussian distribution.

To illustrate this, consider the aeroelastic problem presented in Subsection 2.2, and the problem of evaluating the 95th quantile of the opposite of the flutter speed for a given specific value of the design parameter $\boldsymbol{d}^* = m^*$. The corresponding two-dimensional uncertain vector representing the variability of the absorber mass m and the airfoil moment of inertia I would be $\mathbf{X} = (\boldsymbol{X}(m^*), Z)$, with $\mathrm{F} = -V_f$. A GP surrogate $\hat{\mathrm{F}}$ of F is then built from a DoE consisting of (m_i, I_i) pairs and their corresponding flutter speed values (V_{f_i}) obtained by running the simulator model. From a MC set evaluated on the surrogate $\hat{\mathrm{F}}$, the above procedure permits the estimation of the sought quantile and its bounds.

Active Infill Strategy. To improve the accuracy of $\hat{\mathrm{F}}$ in targeting the estimated quantile $\hat{q}$ derived from the GP and MCS, various learning functions have been developed, including the Expected Feasibility Function [42] and the *U-function* [43]. The latter function is defined such that $\Phi(\mathrm{U}(\mathbf{x}))$—where Φ is the cumulative distribution function (CDF) of the standard normal distribution—represents the probability that the sample $\mathbf{x}$ is *correctly classified* according to the sign of the LSF $\mu_{\hat{F}}(\mathbf{x}) - \hat{q}$, and is formulated as:

$$\mathrm{U}(\mathbf{x}) = \left| \frac{\mu_{\hat{F}}(\mathbf{x}) - \hat{q}}{\sigma_{\hat{F}}(\mathbf{x})} \right|. \tag{17}$$

Due to the monotonicity of Φ, the minimal values of U correspond to points with the highest probability of misclassification. Therefore, desirable infill samples for the DoE are obtained by minimizing U, typically over a candidate Monte Carlo population:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathcal{X}} \mathrm{U}(\mathbf{x}),$$

where $\mathbf{x}^*$ identifies the vector most likely to be misclassified by the current surrogate model.

This *acquisition* criterion balances *exploitation*, which refines the surrogate near the estimated quantile, and *exploration*, which favors regions of high uncertainty [42]. This strategy enhances the accuracy of the GP in the vicinity of the LSF, thereby improving the estimation of both the failure probability and the quantile [38, 40].

3.3. Deterministic optimization assisted by Kriging

Independent of the previous Subsection 3.2, this subsection presents the fundamental principles of *Bayesian Optimization* (BO) for solving constrained mono-objective optimization problems with deterministic black-box functions. The general formulation considered is:

$$\min_{\mathbf{x}\in\mathcal{X}\subset\mathbb{R}^d} \mathrm{F}(\mathbf{x}), \quad \text{s.t. } \mathrm{G}_k(\mathbf{x}) < 0, \quad k = 1, \dots, p, \tag{18}$$

where $\mathcal{X}$ denotes the design space, and F and G_k represent the objective and constraint functions, respectively.

BO, also referred to as *Efficient Global Optimization* [37, 44], relies on Kriging surrogates $\hat{\mathrm{F}}$ and $\hat{\mathrm{G}}_k$ of the true functions F and G_k, constructed from an initial DoE. The optimization problem in Eq. (18) is then solved through an iterative process that iteratively enriches the DoE. Based on the current surrogate models $\hat{\mathrm{F}}$ and $\hat{\mathrm{G}}_k$, *acquisition functions* are optimized to identify new candidate design vectors. These candidates represent both potential solutions to Eq. (18) and informative locations where the expensive functions F and G_k should be evaluated to improve the DoE. In practice, this corresponds to solving auxiliary optimization problems that rely only on surrogate model evaluations, whose computational cost is negligible compared with that of the numerical solvers associated with with F and G_k. The remainder of this section introduces these acquisition functions and describes their role in guiding the enrichment of the DoE and improving the accuracy of the surrogate models for solving Eq. (18).

Acquisition functions. While the constrained case will be adressed in the next paragraph, several infill-sampling criteria for *unconstrained* optimization are considered in this framework, including minimizing the surrogate prediction (MSP), minimizing the lower confidence bound (LCB), maximizing the expected improvement (EI), the probability of improvement (PI), the scaled EI (sEI), and WB2. For detailed definitions and formulations, the reader is referred to [37, 45]. Among these, MSP, LCB, and EI are briefly recalled below:

- LCB involves a linear combination of the posterior mean and standard deviation of the GP predictor :

  $$LCB(\mathbf{x}) = \mu_{\hat{\mathrm{F}}(\mathbf{x})} + k_{LCB}\sigma_{\hat{\mathrm{F}}(\mathbf{x})},$$

  where the parameter $k_{LCB} \in [-3, 3]$ controls the exploration–exploitation trade-off. The special case $k_{LCB} = 0$ corresponds to the MSP criterion.

- EI is a *goal*-oriented criterion based on the statistical improvement function $I(\mathbf{x}) = \max(\mathrm{y}_{min} - \hat{\mathrm{F}}(\mathbf{x}), 0)$, where y_{min} typically denotes the best observed value in the current DoE. The expected value of this improvement yields the EI, which gets the following closed-form expression:

  $$EI(\mathbf{x}) = (\mathrm{y}_{min} - \mu_{\hat{\mathrm{F}}(\mathbf{x})})\,\Phi\left(\frac{\mathrm{y}_{min} - \mu_{\hat{\mathrm{F}}(\mathbf{x})}}{\sigma_{\hat{\mathrm{F}}(\mathbf{x})}}\right) + \sigma_{\hat{\mathrm{F}}(\mathbf{x})}\,\phi\left(\frac{\mathrm{y}_{min} - \mu_{\hat{\mathrm{F}}(\mathbf{x})}}{\sigma_{\hat{\mathrm{F}}(\mathbf{x})}}\right)$$

  where $\phi(\cdot)$ is the standard normal probability density function (PDF).

Note that, like EI, the PI, sEI, and WB2 criteria also depend on the value of y_{min}.

The new design point is then obtained by solving:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathcal{X}} \hat{\mathrm{A}}(\mathbf{x}), \tag{19}$$

where $\hat{\mathrm{A}}$ denotes one of the acquisition functions: LCB, −EI, −PI, −sEI, or −WB2. The hat symbol indicates that $\hat{\mathrm{A}}$ is defined based on surrogate models.

This optimization generally requires solving a global optimization problem that may involve multiple local minima. It can be tackled using gradient-free methods such as CMA-ES or Differential Evolution, or through local gradient-based algorithms. In this work, the latter strategy is employed: the SLSQP algorithm is used via the *scipy.optimize.minimize* function [46], with multiple restarts from randomly selected initial designs to mitigate the influence of local minima.

Handling constraints. Several strategies exist for incorporating constraints within Bayesian optimization.

A first common approach employs the mean predictions of the constraint surrogates during acquisition function optimization. This transforms the original unconstrained subproblem (Eq. (19)) into the constrained formulation:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathcal{X}} \hat{\mathrm{A}}(\mathbf{x}), \quad \text{s.t. } \mu_{\hat{\mathrm{G}}_k(\mathbf{x})} < 0, \quad k = 1, \ldots, p. \tag{20}$$

This strategy, however, neglects the predictive variance associated with the constraint models. In the proposed framework, this approach is used in conjunction with the MSP criterion to identify the current best design vector.

An alternative approach accounts for this uncertainty through the concept of the *Probability of Feasibility* (PF), defined as [45]:

$$\widehat{\mathrm{PF}}(\mathbf{x}) = \prod_{k=1}^{p} \mathcal{P}_k(\hat{\mathrm{G}}_k(\mathbf{x}) < 0) = \prod_{k=1}^{p} \Phi\left(-\frac{\mu_{\hat{\mathrm{G}}_k(\mathbf{x})}}{\sigma_{\hat{\mathrm{G}}_k(\mathbf{x})}}\right) \in [0, 1], \tag{21}$$

where $\mathcal{P}_k$ denotes the probability measure associated with the Gaussian posterior of $\hat{\mathrm{G}}_k(\boldsymbol{x})$. This feasibility indicator can then be incorporated into the acquisition process either via a penalization scheme:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathcal{X}} \hat{\mathrm{A}}(\mathbf{x}) + k_{PF}(1 - \widehat{\mathrm{PF}}(\mathbf{x})) \tag{22}$$

where k_{PF} is a large penalization coefficient that discourages constraint violations (i.e., values of $\widehat{\mathrm{PF}}$ far from 1), or through a multiplicative formulation:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathcal{X}} \hat{\mathrm{A}}(\mathbf{x})\,\widehat{\mathrm{PF}}(\mathbf{x}), \tag{23}$$

noting that the multiplicative approach is not compatible with the LCB criterion.

A final option consists in discarding constraints altogether, hence prioritizing the minimization of the objective function.

Random selection. In the present work, the methods described in this subsection are employed to generate a design sample aimed at solving Eq. (18), based on the surrogate models $\hat{\mathrm{F}}$ and $\hat{\mathrm{G}}$ provided as input. To leverage the diversity of the strengths and weaknesses of the formulations developed above and to introduce stochasticity into the overall optimization procedure (fully described in Section 4), a random selection approach is adopted in the formulation of the sub-optimization process, such that the acquisition and constraint-handling methods are randomly chosen from five and four alternatives, respectively.

4. Numerical methodology for quantile-based optimization

This section outlines the proposed methodology for quantile-based optimization as formulated in Eq. (2), based on a *coupled-space* Kriging-assisted approach. It relies on the following steps:

- An initial DoE is generated in the input space (see Subsection 4.1) by evaluating the numerical simulators f and g_k, which is then used to construct the Kriging models $\hat{f}$ and $\hat{g}_k$. These GP predictors, defined over the *full* space, are subsequently used to adaptively propose input vectors to be added to the DoE, with the goal of solving Eq. (2) until a convergence criterion is satisfied.

- Building upon the fundamental concepts presented in Section 3, this coupled-space infill sampling strategy forms the backbone of the proposed algorithm and represents the main novelty of this work. The procedure begins by constructing quantile surrogates from $\hat{f}$ and $\hat{g}_k$ as functions of the design vector $\boldsymbol{d}$ (see Subsection 4.2). The problem is then reformulated as a deterministic optimization task, as given in Eq. (18), and promising design vectors are generated using the BO framework (Subsection 3.3).

- To map these design vectors to K_θ points in the *full* input space, the reliability-based U-function (Eq. (17)) is employed, as described in Subsection 4.3. These vectors are added to the DoE, and a convergence check determines whether the algorithm terminates or restarts from the metamodel reconstruction.

Figure 2 summarizes the algorithmic flowchart, and the following subsections detail the relevant components of the methodology.

4.1. Initial DoE and coupled-space GPs

The methodology proposed begins by generating an initial dataset in the *full coupled-space*. A set of N_0 samples is first drawn in the unit hypercube $[0,1]^{\mathrm{d}_\theta}$ using quasi-random Latin Hypercube Sampling (LHS) with the maximin distance criterion [47], and subsequently mapped to the physical input space $\Omega = \prod_{i=1}^{\mathrm{d}_\theta}[\theta^{(i)}_{\mathrm{min}}, \theta^{(i)}_{\mathrm{max}}]$ via affine scaling. As noted in Subsection 2.3, the tensorized form of Ω is adopted here for simplicity; however, the methodology remains general and can accommodate any quasi-Monte Carlo sampling scheme defined over Ω.

The numerical simulators f and g_k are evaluated at the resulting input vectors, yielding the initial DoE:

$$\mathcal{T} = \left\{\left(\boldsymbol{\theta}_i, f(\boldsymbol{\theta}_i), g_1(\boldsymbol{\theta}_i), \ldots, g_p(\boldsymbol{\theta}_i)\right)\right\}_{i=1}^{N_0} \tag{24}$$

This dataset is then used to construct the Kriging metamodels $\hat{f}$ and $\hat{g}_k$, as described in Subsection 3.1.

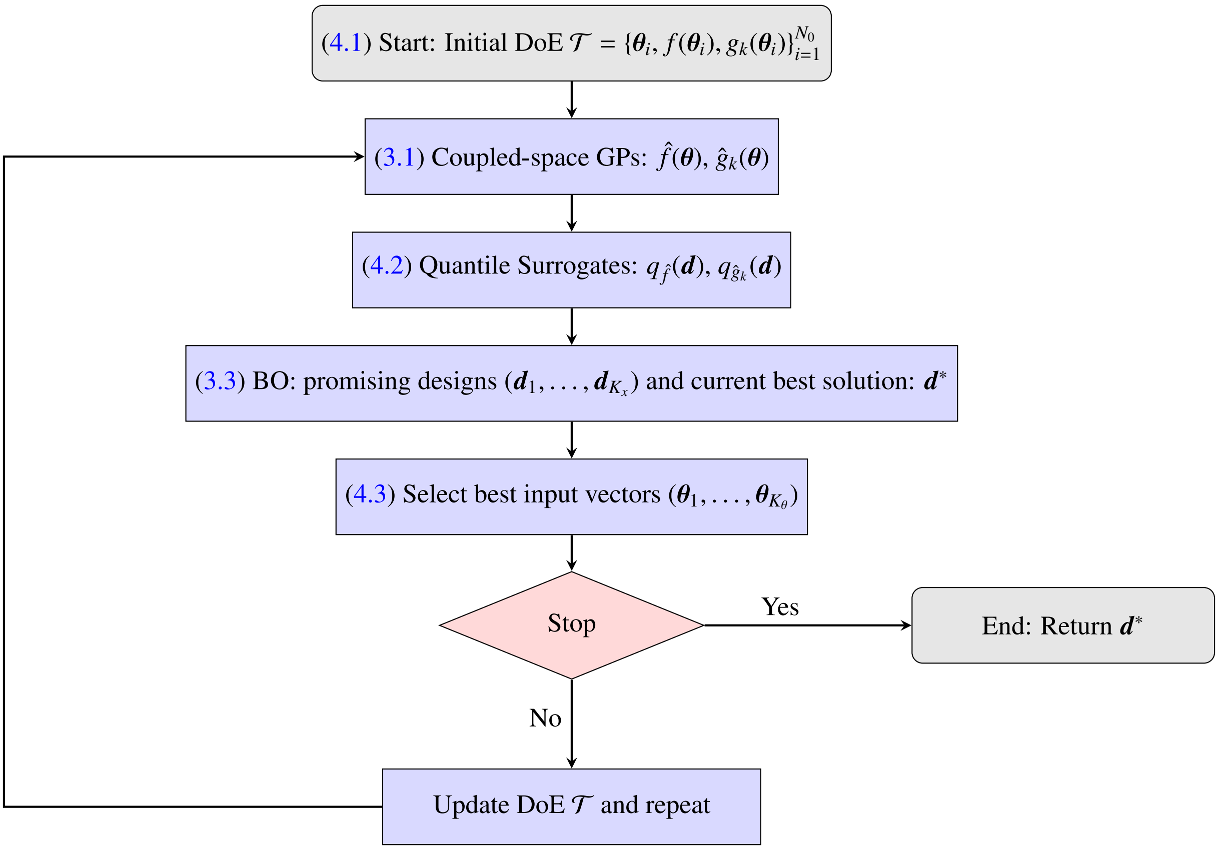

Figure 2: Flowchart of the proposed quantile-based optimization under uncertainty framework. The corresponding subsection for each step is indicated in parentheses.

4.2. *Quantile surrogate modeling and design vector selection*

This subsection describes how the GP predictor $\hat{f}(\boldsymbol{\theta})$, defined over the full input space, is used to estimate the α_f quantile associated with a given design vector $\boldsymbol{d}$, together with lower and upper bounds on this estimate. It concludes by explaining how this procedure is exploited to identify promising design vectors.

For a fixed design vector $\boldsymbol{d}$, a Monte Carlo population $\{\boldsymbol{\theta}_i\}_{i=1}^{n_q}$ is drawn from the conditional distribution $\boldsymbol{\theta}(\boldsymbol{d})$. The reliability analysis procedure presented in Subsection 3.2 is then applied to compute:

- The quantile estimator $\hat{q}_{\hat{f}}(\boldsymbol{d})$,
- Lower and upper bounds $\hat{q}^-_{\hat{f}}(\boldsymbol{d})$ and $\hat{q}^+_{\hat{f}}(\boldsymbol{d})$.

The dependence on α_f is omitted for clarity.

Artificial GP interpretation. For BO-related purposes, an artificial GP representation of the quantile surrogate is adopted, analogous to Eq. (11):

$$q_{\hat{f}}(\boldsymbol{d}) \sim \mathcal{N}\left(\mu_{q_{\hat{f}}(\boldsymbol{d})}, \sigma^2_{q_{\hat{f}}(\boldsymbol{d})}\right), \tag{25}$$

with $\mu_{q_{\hat{f}}(\boldsymbol{d})} := \hat{q}_{\hat{f}}(\boldsymbol{d})$ and $\sigma_{q_{\hat{f}}(\boldsymbol{d})} := \left(\hat{q}^+_{\hat{f}}(\boldsymbol{d}) - \hat{q}^-_{\hat{f}}(\boldsymbol{d})\right) / \sqrt{12}$. This formulation assumes a uniform distribution of the true quantile within its estimated bounds, yielding the stated standard deviation. Although this Gaussian interpretation is heuristic and not theoretically guaranteed, it enables integration with standard BO frameworks. The same procedure applies to each constraint function:

$$q_{\hat{g}_k}(\boldsymbol{d}) \sim \mathcal{N}\left(\mu_{q_{\hat{g}_k}(\boldsymbol{d})}, \sigma^2_{q_{\hat{g}_k}(\boldsymbol{d})}\right). \tag{26}$$

Fixing the seed. Each evaluation of the quantities defined in Eqs. (25) and (26) relies on GP predictions over Monte Carlo samples, which may induce stochasticity in the resulting quantile surrogate *posteriors* $\mu_{q_{\hat{f}}(\boldsymbol{d})}, \sigma_{q_{\hat{f}}(\boldsymbol{d})} : \mathcal{D} \to \mathbb{R}$ and affect subsequent optimization steps. To mitigate this issue, the implementation fixes the random seed when optimization steps are involved, thus transforming the above stochastic functions into deterministic ones. Consequently, for two similar design vectors $\boldsymbol{d} \approx \tilde{\boldsymbol{d}}$, the corresponding Monte Carlo samples $\boldsymbol{\theta}_i \overset{i.i.d.}{\sim} \boldsymbol{\theta}(\boldsymbol{d})$ and $\tilde{\boldsymbol{\theta}}_i \overset{i.i.d.}{\sim} \boldsymbol{\theta}(\tilde{\boldsymbol{d}})$ are expected to be close, yielding a smooth quantile surrogate in both mean and variance.

Promising design vectors. The above protocol allows the original robust optimization problem in Eq. (2) to be reformulated as a deterministic one (Eq. (18)), by adopting $q_{\hat{f}}(\boldsymbol{d})$ and $q_{\hat{g}_k}(\boldsymbol{d})$ as *GP* predictors of $Q_{\alpha_f}[f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})]$ and $Q_{\alpha_{g_k}}[g_k(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})]$, respectively. In line with the BO framework described in Subsection 3.3, this enables the selection of K_x candidate solutions $(\boldsymbol{d}_1, \dots, \boldsymbol{d}_{K_x})$ and the identification of the current optimum $\boldsymbol{d}^*$. A further clarification is needed when considering goal-oriented BO criteria, which require an estimate of the current *best* value in the DoE, denoted $f_{\min}$. This value should correspond to the quantile of the objective function evaluated at the current optimum, which is approximated to lie within the interval $[\hat{q}^- \hat{f}(\boldsymbol{d}^*), \hat{q}^+ \hat{f}(\boldsymbol{d}^*)]$. To account for this, $f_{\min}$ is sampled as K_x evenly spaced values within this interval.

This approach—interpreting quantile estimates as GP predictors and embedding them within a BO-driven design space exploration—represents one of the key contributions of the proposed algorithm.

4.3. Input vector selection

This subsection details the proposed strategy to select K_θ new input vectors to be added to the DoE, based on the K_x design vectors previously identified.

These K_x design vectors $\{\boldsymbol{d}_q\}_{q=1}^{K_x}$ are, by construction, promising candidates from an optimization standpoint, as outlined in the preceding discussion. However, in order to update the coupled-space GP models, full input vectors $\boldsymbol{\theta} = (\boldsymbol{d}, \boldsymbol{z})$ must be evaluated, where $\boldsymbol{z}$ represents the environmental uncertain variables. A naive strategy would consist in simply pairing each design vector $\boldsymbol{d}_q$ with a random realization $\boldsymbol{z}_q$, yielding $\boldsymbol{\theta}_q = (\boldsymbol{d}_q, \boldsymbol{z}_q)$. Through numerical experimentation, however, we observed that this strategy suffers from several drawbacks, such as limited control over the accuracy of quantile estimation and a high likelihood of generating redundant input vectors that are clustered in narrow regions of the input space.

To overcome these limitations, the following method is proposed:

- **Sampling MCS candidate vectors**: For each design vector $\boldsymbol{d} \in \{\boldsymbol{d}_q\}_{q=1}^{K_x}$, a Monte Carlo population of size n_θ is drawn from the conditional distribution $\boldsymbol{\theta}(\boldsymbol{d})$, generating a set of full input vectors $\{\boldsymbol{\theta}_i\}_{i=1}^{n_\theta}$. These vectors inherit the promising nature of their associated design $\boldsymbol{d}$.

- **U-function selection**: Among these n_θ candidates, selection is performed to prioritize those likely to improve the quantile estimates of both the objective and constraint functions—currently approximated by $\mu_{q_{\hat{f}}(\boldsymbol{d})}$ and $\mu_{q_{\hat{g}_k}(\boldsymbol{d})}$, respectively. To this end, the U-function (Eq. (17)) is used as a learning criterion: good candidates are those minimizing the U-function for each objective and constraint, resulting in $p + 1$ different acquisition functions. A Pareto front-based selection is then applied in this multi-objective setting, retaining only the n^q_{pareto} first-rank vectors.

- **Final selection**: Out of the $K_x \sum_{q=1}^{K_x} n^q_{\text{pareto}}$ selected candidates, a K-means clustering algorithm [48] is used to extract the final K_θ input vectors to be added to the DoE. Specifically, the point closest to the centroid of each cluster is selected.

This strategy ensures that selected points originate from promising design vectors, contribute to improving quantile accuracy, and promote diversity in the DoE by avoiding excessive clustering.

Random selection. To promote stochastic exploration—particularly during the early stages when the surrogate models may still be inaccurately fitted—an additional input-space enrichment procedure is introduced. During the first $N_{\text{iteration\_random}}$ iterations, this mechanism is activated with probability $p_{\text{iteration\_random}} \in (0, 1)$. When triggered, the number of design vectors is increased from K_θ to $K_\theta + 1$ for the current iteration. The additional vector is determined by maximizing its minimum Euclidean distance to all existing DoE points, including the K_θ vectors selected through the procedure described above. This maximin-distance criterion enhances the space-filling properties of the design and encourages broader exploration of the input domain during the initial optimization phase.

4.4. Numerical settings

The proposed algorithm involves several numerical parameters used throughout this work, which are summarized in Table 1. These values were selected based on preliminary sensitivity analyses to ensure a satisfactory compromise between computational cost and convergence robustness.

Table 1: Key hyperparameters of the proposed quantile-based optimization algorithm.

Parameter	Description	Typical value(s)
N_0	Initial DoE size	30–100
n_q	MC size for quantile estimation	10^3–10^4
K_x	BO-proposed design vectors per iteration	3–10
K_θ	Full input vectors added per iteration	2–5
n_θ	MC candidates per design vector	100–500
n^q_{pareto}	Pareto points retained per $\boldsymbol{d}_q$	2–5
$N_{\text{iteration\_random}}$	Iterations with random enrichment	5–15
$p_{\text{iteration\_random}}$	Random enrichment probability	0.2–0.4

5. Numerical applications

This section first presents several analytical test cases, studied in [23] and [29], to validate and illustrate the proposed methodology. It then considers the design problem of a NLTVA [11] in order to assess its relevance and performance in a more realistic mechanical engineering application context.

5.1. Analytical functions

5.1.1. Unconstrained quantile minimization

The problem studied in this subsection, extracted from [23], is formulated as a quantile minimization of level $\alpha_f = 80\%$. It can be written as follows:

$$\min_{\boldsymbol{d}\in\mathcal{D}} \quad Q_{80\%}[f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})], \text{ where :} \tag{27}$$

- the cost function f is defined as:

$$f(x, z_1, z_2) = 3 - 4\exp\big(-4(x-4)^2\big) - 5.2\exp\big(-4(x-6)^2\big) + \frac{x-a}{b-a}z_1 + \frac{b-x}{b-a}z_2; \tag{28}$$

- $\boldsymbol{Z} = (Z_1, Z_2)$ is a two-dimensional random vector whose components are independent and follow a uniform and a normal distribution, respectively: $Z_1 \sim \mathcal{U}(0, 10)$, $Z_2 \sim \mathcal{N}(1.01, 0.71)$;
- the one-dimensional random design variable $\boldsymbol{X}(\boldsymbol{d})$ is *deterministic*, so that $f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})$ simply reads $f(\boldsymbol{d}, \boldsymbol{Z})$. Here, for a sake of simplicity and clarity, the optimal design variable will be referred to in terms of x;
- the corresponding design space is $\mathcal{D} = [a, b]$, with $a = 2$ and $b = 8$.

The unconstrained test case is presented in Figure 3a. As this is a stochastic simulator, multiple values of the function f are possible for each x, and the grey points therefore represent a Monte Carlo sample of size 10,000. Since the objective is to minimize the 80% quantile of this function, the corresponding quantile curve is plotted as a black dashed line. This curve was obtained using a Monte Carlo sample of size 10^7 at each point x along the curve, as no analytical expression for the quantile is available. By construction, the 80% quantile of f exhibits two minima: a global minimum at $x^* = 4$ and a local minimum at $x = 6$. These results serve as a reference for assessing the accuracy and efficiency of the proposed method.

When applying the proposed methodology, the evolution of the system at selected iterations of the convergence procedure is presented next, and illustrated in Figures 3b–3e, highlighting the progressive convergence toward the optimum.

1. An initial Design of Experiments (DoE) of size $N_0 = 4$ is generated, constituting the first iteration (see Figure 3b). This allows the construction of a coupled-space metamodel of f as well as a metamodel of the 80% quantile of the objective function (shown as a solid red line), along with a prediction interval corresponding to ±2 *standard deviations*, as detailed in Eq. (25). The proposed infill strategy selects, at each iteration, $K_\theta = 2$ input vectors in the coupled space—here, triplets $\mathbf{t} = (x, z_1, z_2)$—which are then evaluated using the function f to enrich the DoE.
2. After 4 iterations (Figure 3c), the approach appears to have identified the local minimum around $x = 6$: in this region, the uncertainty associated with the metamodel is relatively low compared to the region near the global optimum at $x^* = 4$, where the uncertainty remains very large.
3. After 7 iterations (Figure 3d), the metamodel is significantly more accurate in the region of the global optimum, although some areas of large uncertainty persist, which are unimportant from an optimization perspective.
4. After 10 iterations (Figure 3e), the metamodel has continued to refine itself in the vicinity of the global optimum and more rarely in other regions. Most notably, the lower bound of the confidence interval is globally minimal in the region of the global optimum.

At each iteration, once the quantile metamodels are constructed, it is possible to estimate the best design vector proposed by the methodology, which enables plotting the convergence curve shown in Figure 3f. The global optimum is indicated by a green dotted line. As previously described, the method initially detects the local optimum near $x = 6$ (iterations 3 and 4) before focusing on the global one. The global optimum is very well approximated after only 13 evaluations (5 iterations) using the proposed approach: a substantial improvement compared to the 150 evaluations required in [23], although in the latter the stochastic space is not parametrized.

5.1.2. *Constrained minimization*

The second analytical problem studied is extracted from [29] and considers a constrained stochastic optimization problem with $\mathrm{d}_x = 4$ design variables and $\mathrm{d}_z = 3$ uncertain parameters. It is formulated as:

$$\begin{aligned} &\min_{\boldsymbol{d}\in\mathcal{D}} \quad \mathbb{E}\left[f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})\right], \quad \text{subject to:} \\ &Q_{95\%}[g(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})] \leq 1, \quad \text{where:} \end{aligned} \tag{29}$$

- the objective and constraint functions are defined as

$$f(x_1, x_2, x_3, x_4, z_1, z_2, z_3) = \sum_{i=1}^{3} \Big[(1 - x_i) + 3(x_{i+1} - x_i^2)^2 (1 + 0.5 \arctan(5(z_i - 0.5)))\Big], \tag{30}$$

$$\begin{aligned} g(x_1, x_2, x_3, x_4, z_1, z_2, z_3) = &\left(4 - 2.1x_1^2 + \frac{x_1^4}{3}\right) x_1^2 + x_1 x_2 + 4(x_2^2 - 1)x_2^2 \\ &+ 0.2\Big(\cos(2\pi z_1) - \sin(0.5\pi z_1) - z_1 - [\cos(2\pi \cdot 0.05) - \sin(0.5\pi \cdot 0.05) - 0.05]\Big). \end{aligned} \tag{31}$$

- the uncertain inputs $\boldsymbol{Z} = (Z_1, Z_2, Z_3)$ are modeled as independent and identically distributed uniform distributions: $Z_k \sim \mathcal{U}(0, 1)$, for $k \in [\![1, 3]\!]$,

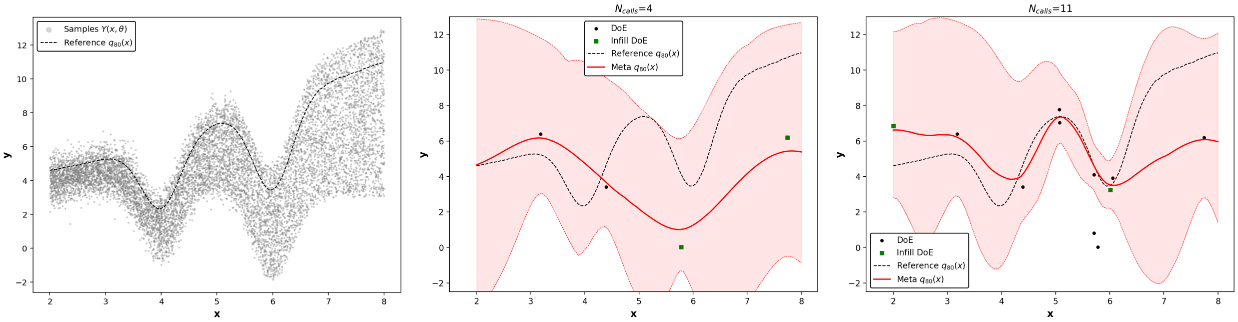

(a) Test case definition

(b) First iteration: Initial DoE (N_0 = 4)

(c) After 4 iterations

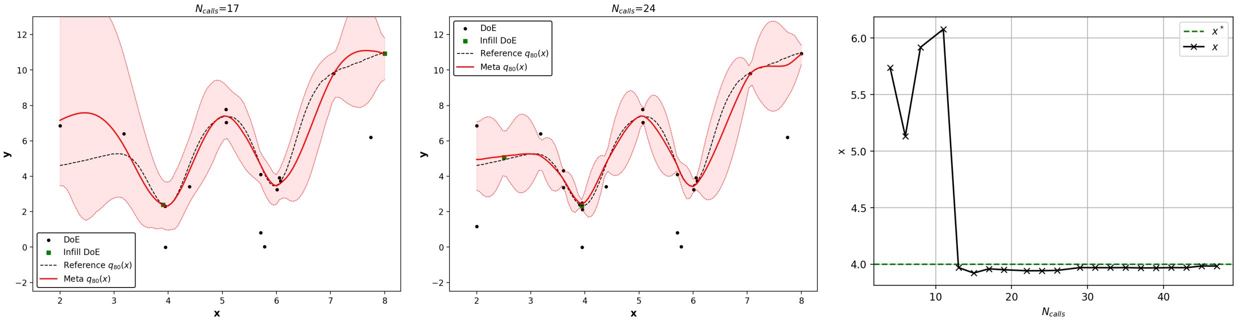

(d) After 7 iterations

(e) After 10 iterations

(f) Convergence history of the best design d_1

Figure 3: Illustration of the unconstrained analytical test case from [23]. The gray points represent the variability of the stochastic model in the design space, while the reference 80% quantile (black dashed line) is obtained by Monte Carlo simulation. An initial metamodel is constructed from a DoE of size $N_0 = 4$, and then updated at each iteration with $K_\theta = 2$ samples $\{x^k, \theta_1^k, \theta_2^k\}_{k=1}^{K_\theta}$ selected adaptively. The 80% quantile metamodel is shown in red, together with its 95% *confidence interval*, and the DoE is indicated by large black dots. The evolution of the *best design* throughout the optimization process is shown in subfigure (f), highlighting its convergence toward the analytical reference.

- the random design mapping $\boldsymbol{X}(\boldsymbol{d})$ is again deterministic, such that the optimization is performed directly over the design vector $\boldsymbol{d} = (x_1, x_2, x_3, x_4)$ within the hyper-rectangular design space $\mathcal{D} = [-0.2,\ 1.2]^4$.

The original problem defined as Eq. (29) in [29] is equivalently rewritten in terms of quantiles as:

$$\begin{aligned} &\min_{\boldsymbol{d}\in\mathcal{D}} \quad Q_{50\%}\left[f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})\right], \quad \text{subject to:} \\ &Q_{95\%}[g(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})] - 1 \leq 0. \end{aligned} \tag{32}$$

This equivalence holds because, for any $Z \sim \mathcal{U}(0, 1)$, the random variable $(1 + 0.5\ \arctan(5(Z - 0.5)))$ is symmetrically distributed around its mean value of 1. Consequently,

$$\mathbb{E}\Big[1 + 0.5\ \arctan(5(Z - 0.5))\Big] = Q_{50\%}\Big[1 + 0.5\ \arctan(5(Z - 0.5))\Big] = 1. \tag{33}$$

It therefore follows, in the present case, that

$$Q_{50\%}[f(\mathbf{X}(\mathbf{d}), \mathbf{Z})] = \mathbb{E}[f(\mathbf{X}(\mathbf{d}), \mathbf{Z})]. \tag{34}$$

The median of the objective function and the 95% quantile of the constraint admit closed-form expressions, that reduce respectively to the classical four-dimensional Rosenbrock and Six-Hump Camel functions:

$$Q_{50\%}\left[f(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})\right] = \sum_{i=1}^{3}\Big[(1 - x_i) + 3(x_{i+1} - x_i^2)^2\Big], \tag{35}$$

$$Q_{95\%}[g(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})] = \left(4 - 2.1x_1^2 + \frac{x_1^4}{3}\right)x_1^2 + x_1 x_2 + 4(x_2^2 - 1)x_2^2. \tag{36}$$

The reference optimal solution reported in [29] is $\boldsymbol{d}^\star = (0.7033,\ 0.7035,\ 0.6212,\ 0.3859)$, with corresponding objective and constraint quantile values $Q_{50\%}^{\mathrm{ref}} = 0.4981$ and $Q_{95\%}^{\mathrm{ref}} = 1$.

For this test case, an initial DoE of size $N_0 = 10$ is used, followed by an enrichment with $K_\theta = 4$ additional design vectors per iteration. The procedure is stopped before reaching 100 function evaluations. The convergence behavior is illustrated in Figure 4(a). It shows the exact 50% quantile value (i.e., the mean here) of the objective function f evaluated at each iteration for the best design vector proposed by the method. The horizontal dashed line indicates the reference optimum value from [29]. Since this is a constrained optimization problem, feasible designs, satisfying $Q_{95\%}[g]) \leq 1$, are marked with green triangles, while infeasible designs are shown as red crosses. The first evaluations yield inconsistent results, with designs exhibiting either high objective function values or constraint violations. From approximately 40 calls to the analytical function, the results appear significantly closer to the reference value and converge toward it. The final designs achieve objective values very close to the reference, with constraint violations that progressively disapear.

To enable direct comparison with results given in [29], the final optimal design vectors, together with the corresponding $Q_{50\%}[f]$ and $Q_{95\%}[g]$ values, are presented in Figure 4(b) and compared against the reference solution. Similarly to Figures 10(a)–(c) of [29], reported after 150, 161, and 200 evaluations respectively, the proposed method visually achieves at least comparable to superior accuracy with a computational budget of approximately one hundred evaluations.

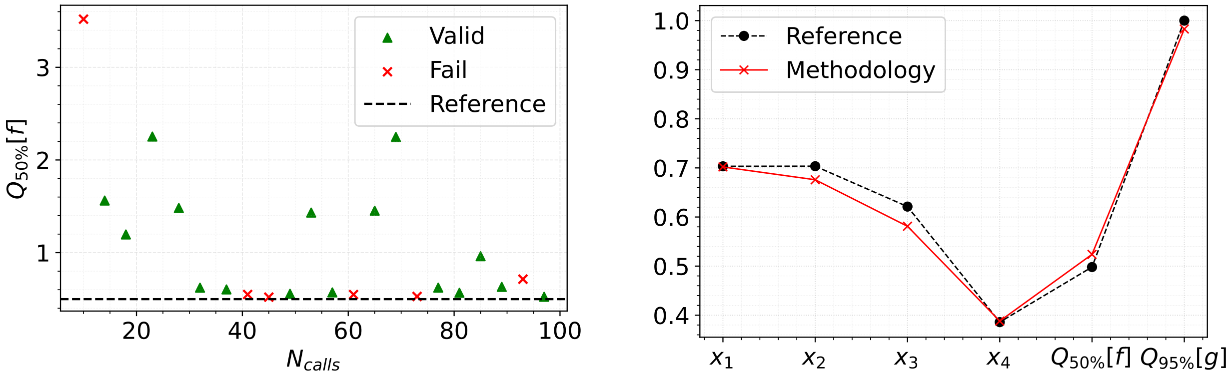

(a)

(b)

Figure 4: Illustration of the constrained analytical test case from [29]. (a): Convergence history showing the evolution of the objective $Q_{50\%}$ evaluated at the best proposed design at each iteration. Feasible and unfeasible designs are marked with green triangles and red crosses respectively. (b): Comparison of the final design and quantiles against the reference optimum [29].

These two analytical test cases therefore demonstrate that the proposed adaptive sampling methodology validates its performance against state-of-the-art approaches for robust optimization assisted by surrogate modelling in coupled spaces, where design vectors are not affected by variability.

5.2. Robust optimization of NLTVA parameters

5.2.1. Case presentation

In this study, a Nonlinear Tuned Vibration Absorber (NLTVA) is implemented to control the flutter instability of a two-degree-of-freedom airfoil model subjected to a quasi-steady aerodynamic flow. Robust optimization of its design is performed following the proposed methodology, while explicitely considering variability in the decision parameters.

The system follows the dimensionless formulation proposed by [11]. It is illustrated in Fig. 5. The airfoil's pitch (rotation about the elastic axis α) and plunge (vertical displacement y) motions are characterized by the natural frequencies ω_α and ω_h, and by the damping ratios ζ_α and ζ_h, respectively. The frequency ratio $\Omega = \omega_h/\omega_\alpha$ between the plunge and pitch motion, and the reduced mass x_α and moment of inertia r_α of the airfoil will be used in the governing equations. Possible geometric nonlinearities are also accounted for by introducing dimensionless cubic stiffness terms

ξ_α and ξ_h, associated with the pitch and plunge motions, respectively. The NLTVA consists of a secondary mass-spring-damper system of displacement $\tilde{x}$, characterized by its natural frequency ω, its damping ratio ζ, and connected to the airfoil at a prescribed dimensionless offset distance λ from its elastic center. The NLTVA-reduced frequency γ defined as $\gamma = \omega^2/\omega_\alpha^2$ will be used. The mass ratio between the mass of the NLTVA and that of the wing is defined as ϵ. The restoring force of the NLTVA also includes a cubic nonlinear stiffness component of dimensionless stiffness ξ. The detailed expression of all these parameters in terms of inertia, stiffness, and damping properties of both the airfoil and the NLTVA are provided in [11]. For brevity, they are not repeated here, but the same notation is deliberately retained.

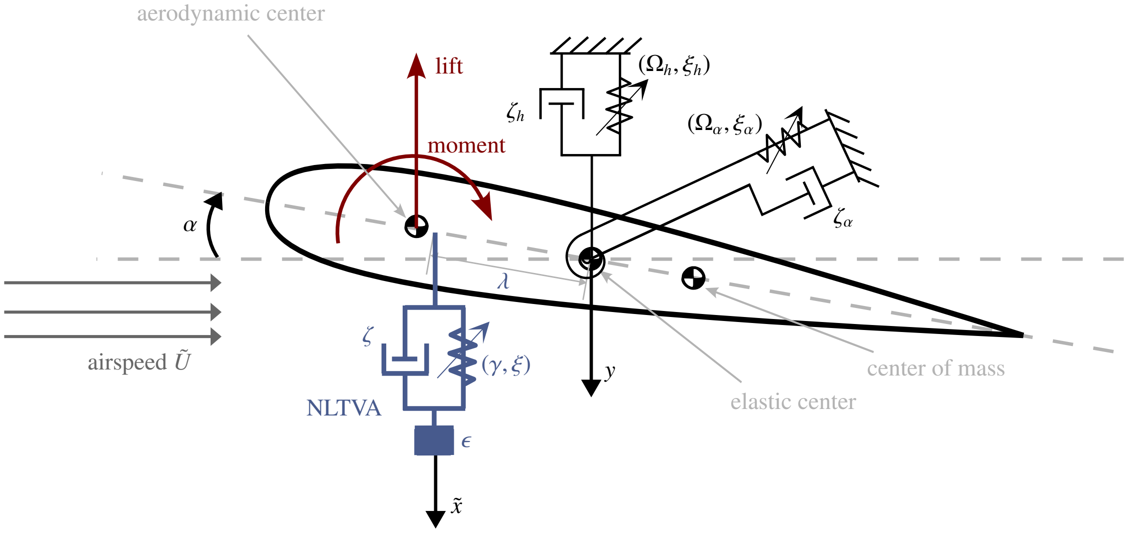

Figure 5: Two-degree-of-freedom airfoil coupled with a nonlinear tuned vibration absorber (NLTVA). All quantities are expressed in dimensionless form; the corresponding physical variables and the relations used for nondimensionalization are provided in [11].

The aerodynamic loads are described using a simple quasi-steady model, assuming a linear dependence of lift and moment on instantaneous displacements and velocities. The dimensionless aerodynamics parameters used are $\tilde{U}$, β and ν. These quantities depend on the airflow velocity, the lift-curve slope coefficient, and the distance between the aerodynamic and elastic centers.

The governing equations of motion of the airfoil–NLTVA system can be written in dimensionless form as [11]:

$$\boldsymbol{M}\ddot{\boldsymbol{q}} + \boldsymbol{C}\dot{\boldsymbol{q}} + \boldsymbol{K}\boldsymbol{q} + \boldsymbol{f}_{\text{nl}}(\boldsymbol{q}) = \boldsymbol{0}. \tag{37}$$

The degrees of freedom are contained in the vector $\boldsymbol{q}$ such that $\boldsymbol{q} = \begin{bmatrix} y & \alpha & \tilde{x} \end{bmatrix}^{\mathrm{T}}$. The linear mass $\boldsymbol{M}$, damping $\boldsymbol{C}$ and stiffness $\boldsymbol{K}$ matrices, including aerodynamics effects are defined with:

$$\boldsymbol{M} = \begin{bmatrix} 1 & x_\alpha & 0 \\ x_\alpha & r_\alpha^2 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{C} = \begin{bmatrix} \zeta_h + \epsilon\zeta + \beta\tilde{U} & -\epsilon\zeta\lambda & -\epsilon\zeta \\ -\nu\tilde{U} - \epsilon\zeta\lambda & \zeta_\alpha + \epsilon\zeta\lambda^2 & \epsilon\zeta\lambda \\ -\zeta & \zeta\lambda & \zeta \end{bmatrix}, \quad \text{and} \quad \boldsymbol{K} = \begin{bmatrix} \Omega^2 + \epsilon\gamma & \beta\tilde{U}^2 - \epsilon\gamma\lambda & -\epsilon\gamma \\ -\epsilon\gamma\lambda & r_\alpha^2 - \nu\tilde{U}^2 + \epsilon\gamma\lambda^2 & \epsilon\gamma\lambda \\ -\gamma & \gamma\lambda & \gamma \end{bmatrix}.$$

The nonlinear force term $\boldsymbol{f}_{\text{nl}}(\boldsymbol{q})$ resulting from both the assumed geometric nonlinearities of the airfoil and the nonlinear restoring force of the NLTVA, is expressed as:

$$\boldsymbol{f}_{\text{nl}}(\boldsymbol{q}) = \begin{bmatrix} \xi_h y^3 + \epsilon\xi(y - \tilde{x} - \lambda\alpha)^3 \\ \xi_\alpha \alpha^3 + \epsilon\lambda\xi(\tilde{x} + \lambda\alpha - y)^3 \\ \xi(\tilde{x} + \lambda\alpha - y)^3 \end{bmatrix}.$$

The original study showed the ability of the NLTVA to delay and suppress flutter onset while effectively limiting post-flutter oscillations. The objective of this section is not to investigate the validity of the modeling assumptions but solely focus on robust optimization following the development made in Section 4. The reference configuration is defined with the parameters given in Table 2, extracted from [11]. Several robust optimization problems are considered and investigated in the subsequent numerical simulations.

Table 2: Dimensionless aeroelastic parameters of the reference configuration [11].

Symbol	**Description**	**Nominal value**
x_α	Reduced static unbalance of the airfoil	0.2
r_α	Reduced radius of gyration of the airfoil	0.5
β	Reduced lift curve slope	0.2
ν	Reduced aerodynamic moment coefficient	0.08
Ω	Frequency ratio ω_h/ω_α	0.5
ζ_α	Pitch damping ratio	0.01
ζ_h	Plunge damping ratio	0.01
ξ_α	Cubic stiffness coefficient in pitch	1.0
ξ_h	Cubic stiffness coefficient in plunge	0.0

5.2.2. Numerical quantities of interest and optimization problems

The primary objective is to maximize the 5% quantile of the *normalized flutter speed* V_f (i.e., neglecting the 5% worst-case scenarios) by tuning a subset of the NLTVA parameters. V_f is defined as the smallest value of the nondimensional airspeed $\tilde{U}$ at which the generalized complex eigenvalue problem

$$(\lambda^2\mathbf{M} + \lambda\mathbf{C}(\tilde{U}) + \mathbf{K}(\tilde{U}))\delta\mathbf{q} = \mathbf{0} \tag{38}$$

admits a purely imaginary conjugate pair of eigenvalues, where $\delta\mathbf{q}$ represents a small perturbation of $\mathbf{q}$ around the $\mathbf{0}$-equilibrium position. This quantity is determined using a combination of grid search and Brent's root-finding method, where the goal is to locate the point at which the real part of an eigenvalue crosses zero.

The optimization is performed under the uncertainty scenario summarized in Table 3, where all NLTVA parameters are subject to inherent uncertainties.The chosen uncertainty scenario has been selected to illustrate the capabilities of the proposed methodology. The mass ratio ϵ and the normalized position λ have their nominal value set at their maximum allowable bound, i.e. their optimal values [11], and therefore act purely as external random inputs.

In some formulations, additional probabilistic constraints are enforced to guarantee the following:

- Supercritical Hopf bifurcation behavior, ensured by requiring the *first Lyapunov coefficient* l_1 to be negative. This coefficient determines the nature of the Hopf bifurcation leading to the onset of flutter. It is computed using the standard center manifold reduction in nonlinear autonomous systems described in [49] (Chapter 3). A negative value of l_1 indicates a supercritical Hopf bifurcation, characterized by a stable limit cycle that emerges smoothly as $\tilde{U}$ crosses its critical value. A positive l_1-value indicates a subcritical bifurcation in which unstable limit cycle oscillations appear before $\tilde{U}$ reaches V_f, often leading to a sudden jump in amplitude potentially dangerous in aeroelastic systems.

- Limitation of the *LCO amplitudes in plunge*. For a given airspeed $\tilde{U}$, the LCO amplitudes in plunge, denoted $y_{\max}(\tilde{U})$, along with those in pitch and for the absorber, are obtained by numerical integration of the governing ordinary differential equations (Eq. (37)), formulated as a first-order ODE, using a fourth-order Runge–Kutta method, and starting from $\boldsymbol{q}(0) = \begin{bmatrix}0.1 & 40\pi/180 & 0.4\end{bmatrix}^{\mathrm{T}}$ and $\dot{\boldsymbol{q}}(0) = \mathbf{0}$. These relatively large initial perturbations are deliberately chosen to increase the likelihood of capturing the worst-case stable solution, e.g., in the presence of bistability. The uniqueness of the stable solution, or the selection of the upper branch when it exists, is verified a posteriori for some critical designs using branch continuation techniques (see more details in Subsection 5.2.4). The plunge amplitude to restrict is evaluated at a selected airflow velocity $\tilde{U}^*$ above the nominal critical flutter speed V_f, and is denoted $y_{\mathrm{LCO}} = y_{\max}(\tilde{U}^*)$.

More formally, the robust above-mentionned optimization problem is formulated as follows:

$$\min_{\boldsymbol{d}\in\mathcal{D}} \quad Q_{0.95}[-V_f(\boldsymbol{X}(\boldsymbol{d}),\boldsymbol{Z})], \qquad \boldsymbol{d} = (\zeta,\gamma,\xi), \qquad \boldsymbol{Z} = (\epsilon,\lambda), \tag{39}$$

Table 3: Stochastic modeling of the random vector $\theta(\boldsymbol{d}) = (\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})$ and associated design space $\mathcal{D}$ for the NLTVA parameters. All uncertain parameters follow a Gaussian distribution truncated at $\pm 3\sigma$ and centered on the nominal (or design) value μ: $\mathcal{U} \sim \mathcal{N}_{\text{trunc}[-3\sigma,+3\sigma]}(\mu, \sigma^2)$.

Symbol	NLTVA design parameter	Nominal value / design range	Standard deviation σ
ϵ	Mass ratio	0.05	5% of $[0.04, 0.05]$
λ	Normalized position	1	5% of $[-1, 1]$
ζ	Damping ratio	$[0, 0.2]$	5% of the range
γ	Frequency ratio	$[0.2, 0.6]$	5% of the range
ξ	Cubic stiffness coefficient	$[0, 0.4]$	5% of the range

optionally subject to the probabilistic constraints

$$\begin{cases} (\mathrm{C}_1) \qquad Q_{0.95}[l_1(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})] + 0.5 < 0, \\ (\mathrm{C}_2) \qquad Q_{0.95}[y_{\mathrm{LCO}}(\boldsymbol{X}(\boldsymbol{d}), \boldsymbol{Z})] - Q_{0.95}\Big[y_{\mathrm{LCO}}(\boldsymbol{X}(\boldsymbol{d}^{DO}), \boldsymbol{Z})\Big] < 0, \end{cases} \tag{40}$$

where constraint (C_1) enforces negativity of the first Lyapunov coefficient, and constraint (C_2) limits LCO amplitudes in plunge evaluated at the nondimensional velocity $\tilde{U}^* = 1.3$, chosen above the flutter speed ($V_f = 1.25$) associated with the deterministic optimum $\boldsymbol{d}^{DO} = (0.110, 0.462, 0)$, and $Q_{0.95}\Big[y_{\mathrm{LCO}}(\boldsymbol{X}(\boldsymbol{d}^{DO}), \boldsymbol{Z})\Big] \approx 0.0685$. As selecting an appropriate constraint value is inherently problem-dependent, we adopt the deterministic optimum reported by Malher et al. [11] as a reference throughout this work. Indeed, in the absence of uncertainty considerations during the design phase, $\boldsymbol{d}^{DO}$ constitutes a meaningful baseline.

Three cases are considered:

- unconstrained robust optimization, denoted RO, consisting of maximizing the 5%-quantile of the flutter speed V_f;
- robust optimization with constraint (C_1), denoted RO-C_1;
- robust optimization with constraints (C_1) and (C_2), denoted RO-C_1-C_2.

For each case, the surrogate-based method starts with an initial DoE of size $N_0 = 10$ and is iteratively enriched with $K_\theta = 4$ additional points per iteration. Convergence is typically achieved after 45 to 80 evaluations of the numerical solver.

A reference solution for the unconstrained case, denoted RO-Ref, is also provided for comparison. It is obtained using the deterministic global optimization algorithm DIRECT (DIviding RECTangles) introduced in [50] and implemented in `scipy.optimize`. Since the flutter speed does not depend on the cubic stiffness parameter ξ, the objective function is evaluated for each design point $\mathbf{d} = (\zeta, \gamma)$ using a Monte Carlo population of size 200. The maximum number of DIRECT iterations is set to 200, corresponding to a total of $200 \times 200 = 40{,}000$ flutter-speed evaluations. This approach is particularly efficient in low-dimensional design spaces, such as the present two-dimensional case. At each iteration, DIRECT partitions the search domain into hyper-rectangles and systematically refines regions likely to contain the global optimum, thereby balancing global exploration and local exploitation. Repeated runs of the algorithm and variations in Monte Carlo sample size did not significantly affect the results, confirming the robustness of this reference solution.

In the following subsections, the results are compared with both this reference solution and the nominal deterministic optimum from [11], hereafter referred to as DO.

5.2.3. Analysis

The variability in performance across the five designs is illustrated by the probability density functions (PDFs) of the quantities of interest V_f, l_1, and y_{LCO} in Fig. 6, while the corresponding design vectors and QoI statistics are reported in Table 4. The distributions and statistics were computed using Monte Carlo sampling: for each design vector listed in the sixth column of Table 4, $n_{\mathrm{MC}} = 1000$ realizations of the NLTVA parameters were drawn and used to evaluate the system QoIs.

In the following, the analysis is restricted to the LCO amplitudes in the plunge degree of freedom for conciseness. This choice is supported by the results of Malher [11], which show that bifurcation diagrams in plunge exhibit significantly richer and more complex behavior than in pitch. It should also be noted that the present model relies on a simplified two-degree-of-freedom system with linear quasi-steady aerodynamic forces, which yields unrealistically large pitch amplitudes (up to 50°). These values fall outside the validity range of the underlying aerodynamic model.

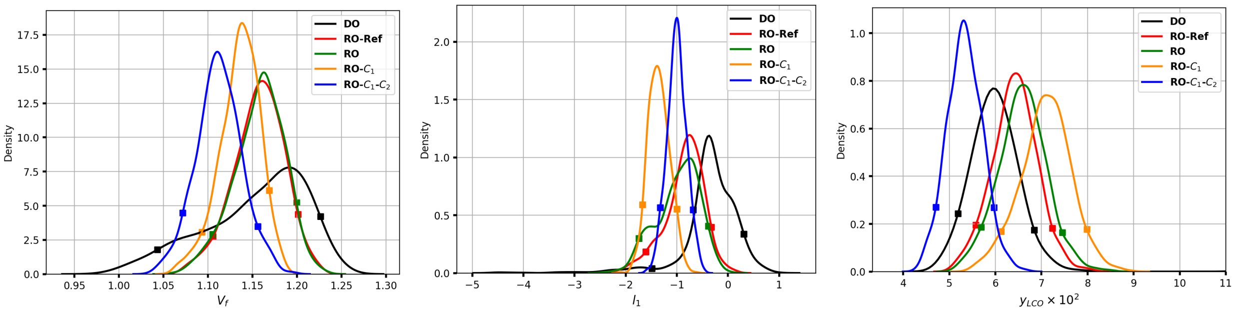

(a) Normalized flutter speed V_f

(b) First Lyapunov coefficient l_1

(c) LCO amplitude in plunge y_{LCO}

Figure 6: PDFs of the three quantities of interest. The curves correspond to the following designs: (—) deterministic optimization (DO) [11], (—) DIRECT reference (RO-Ref), (—) unconstrained robust optimization maximizing the 5%-quantile of V_f (RO), (—) robust optimization with constraint on l_1 (RO-C$_1$), and (—) robust optimization with both constraints on l_1 and y_{LCO} (RO-C$_1$-C$_2$). The 5% and 95% quantiles are indicated by square markers on each curve.

Table 4: Comparison of the designs obtained from Monte Carlo simulations ($n_{\mathrm{MC}} = 1000$). The reported quantities include the 5% and 95% quantiles (denoted respectively q_5 and q_{95}), the nominal values of the QoIs, the optimal design variables (ζ^*, γ^*, ξ^*), and the number of solver evaluations. The definitions of the QoIs and the design labels are given in Fig. 6.

Design	**Quantity**	q_5	q_{95}	Nominal	$(\zeta^*, \gamma^*, \xi^*)$	N_{calls}
DO	V_f	1.043	1.226	1.255	(0.110, 0.462, 0)	-
	l_1	-1.489	0.311	0.982		
	$y_{\mathrm{LCO}} \times 10^2$	5.19	6.84	6.34		
RO-Ref	V_f	1.106	1.201	1.195	(0.134, 0.432, 0)	40 000
	l_1	-1.612	-0.326	-0.423		
	$y_{\mathrm{LCO}} \times 10^2$	5.57	7.24	6.84		
RO	V_f	1.105	1.200	1.201	(0.138, 0.438, 0)	45
	l_1	-1.740	-0.384	-0.667		
	$y_{\mathrm{LCO}} \times 10^2$	5.69	7.45	7.03		
RO-C$_1$	V_f	1.093	1.169	1.166	(0.166, 0.420, 0)	80
	l_1	-1.673	-0.998	-1.178		
	$y_{\mathrm{LCO}} \times 10^2$	6.12	7.98	7.61		
RO-C$_1$-C$_2$	V_f	1.072	1.156	1.134	(0.096, 0.390, 0.180)	83
	l_1	-1.325	-0.684	-0.938		
	$y_{\mathrm{LCO}} \times 10^2$	4.71	5.97	5.67		

The DO case exhibits large variability in flutter speed (Fig. 6a), despite having the highest nominal value of 1.26, but with a 5% quantile as low as 1.04. By first seeking to increase this 5% quantile—effectively neglecting the 5% worst-case scenarios while improving the remaining 95% in terms of flutter speed only—the unconstrained robust optimization (RO) successfully achieves this goal. The reference RO-Ref optimum obtained via DIRECT and the one proposed by the present methodology (RO) are very close in terms of design vector (Table 4), with vectors $(0.134, 0.432, 0)$ versus $(0.138, 0.438, 0)$, corresponding to an L_2 norm error of less than 1.6%. The 5% quantile values are approximately 1.11, with a relative error of less than 1% (1.106 vs. 1.105), representing an improvement of about +6.7% over DO. Visually, the PDFs for RO-Ref and RO show slightly more pronounced variability in the first Lyapunov coefficient l_1, with a minor leftward shift for RO compared to RO-Ref, likely non-significant, as well as a slight increase in LCO plunge amplitudes around 0.2×10^{-2}, with similar PDF shapes. These very minor differences between RO-Ref and RO highlight the accuracy of the proposed methodology, which requires three orders of magnitude fewer objective function evaluations. In the following, only the RO design will be discussed.

To ensure supercriticality for all possible realizations in subsequent designs, RO-C_1 is obtained by imposing that the 95% quantile of the first Lyapunov coefficient must be less than -0.5. The corresponding PDF confirms a supercritical Hopf bifurcation under all uncertain conditions (Fig. 6b), at the expense of a slight reduction in overall flutter speed: the 5% quantile decreases from 1.11 (RO) to 1.09, which remains significantly higher than DO. Visualization of the PDFs for the LCO plunge amplitude evaluated at the airspeed $\tilde{U}^*$ (see Fig. 6c) indicates that, among the considered designs, DO exhibits the lowest amplitudes overall. Its PDF is shifted toward lower values and its 95% quantile, representative of a worst-case scenario, is 6.84×10^{-2}. In comparison, RO reaches 7.45×10^{-2} and RO-C_1 is even higher at 7.98×10^{-2}, corresponding to an increase of approximately 16.7% relative to DO.

To mitigate this degradation, an additional constraint is introduced in the robust formulation. The RO-C_1-C_2 case is formulated similarly to RO-C_1 for ensuring supercriticality overall, with the addition of a constraint enforcing that the 95% quantile of y_{LCO} remains below that of DO. The resulting design satisfies this constraint well, with the 95% quantile of y_{LCO} dropping to 5.97×10^{-2}, well below DO's value, with a PDF that tends to be narrower, indicating greater robustness to variability. The Lyapunov coefficient remains consistently negative, with higher minimum values than in RO or RO-C_1. However, its 5% quantile for flutter speed decreases to 1.07, which is lower than RO and RO-C_1 but still 0.03 higher than DO.

Although DO performs better under nominal conditions, i.e., assuming perfect NLTVA properties, accounting for uncertainties in these parameters, as represented by the scenario in Table 3, reveals that RO-C_1-C_2 is substantially more robust and therefore better suited for industrial applications.

The optimizations performed primarily aim to delay the onset of flutter. Some formulations additionally enforce a supercritical bifurcation in a manner that is robust to uncertainties in the NLTVA parameters. The RO-C_1-C_2 further aims to control LCO amplitudes by enforcing a constraint directly on the LCO plunge amplitude y_{LCO}, as described previously. It should be noted this single operating point served as a proxy metric; the underlying *rationale* is that improvements observed at $\tilde{U}^* = 1.3$—a point located sufficiently beyond the nominal flutter onset for all designs—is expected to yield beneficial effects on LCO amplitudes across a meaningful portion of the post-flutter regime. Consequently, this single-point constraint is anticipated to favorably influence the overall bifurcation behavior of the optimized configurations over the relevant airspeed range. Nevertheless, investigating the full stable branch emerging from the flutter speed remains essential to ensure the absence of higher-amplitude stable responses and to verify that the overall vibratory energy remains within acceptable limits.

5.2.4. Impact on bifurcation diagrams

This subsection presents a detailed analysis of the proposed designs from the perspective of their bifurcation characteristics. For any given design, the incorporation of uncertainties, as quantified in Table 3, generates an ensemble of realizations, each leading to a distinct bifurcation curve.

The PDFs shown in Fig. 6 were obtained from a MC analysis. This implies that, for each design vector, summarized in the sixth column of Table 4, 1000 sets of NLTVA parameters were generated and used to evaluate the system QoIs. Among these realizations, four representative cases are retained for further analysis: the nominal design (i.e., without uncertainties, corresponding to the deterministic design), and three realizations whose flutter speed is closest to the 5%, 50%, and 95% quantiles of the V_f distribution. Additional realizations are selected based on the variability of the LCO plunge amplitude y_{LCO}, obtained at $\tilde{U} = 1.3$, leading to a total of seven realizations per design, denoted by nom, q_5^V, q_{50}^V, q_{95}^V, q_5^y, q_{50}^y and q_{95}^y respectively.

Tables 5 and 6 list the corresponding numerical values of the NLTVA parameters for the selected realizations based on V_f and y_{LCO} variability, respectively.

To provide further UQ insight into the designs through their associated bifurcation diagrams, the following procedure is adopted. For all the realizations discussed above, the bifurcation curves are obtained by computing the LCO amplitudes as a function of the airspeed.

Two numerical approaches are considered:

Table 5: NLTVA parameters ($\zeta, \gamma, \xi, \epsilon, \lambda$) and corresponding performance quantities ($V_f, l_1, y_{\mathrm{LCO}}$) for selected realizations of DO, RO, RO-C_1, and RO-C_1-C_2 designs. The reported cases correspond to realizations whose flutter speed V_f is closest to the 5%, 50%, and 95% quantiles of its distribution.

	DO			RO			RO-C_1			RO-C_1-C_2		
Values	q_5^V	q_{50}^V	q_{95}^V	q_5^V	q_{50}^V	q_{95}^V	q_5^V	q_{50}^V	q_{95}^V	q_5^V	q_{50}^V	q_{95}^V
ζ	0.108	0.102	0.100	0.146	0.134	0.130	0.165	0.151	0.165	0.102	0.0864	0.114
γ	0.491	0.465	0.465	0.408	0.429	0.443	0.401	0.409	0.426	0.413	0.373	0.422
$\xi \times 10^2$	0.527	0.0637	2.43	1.87	3.73	3.15	0.566	0.701	1.86	17.2	17.2	17.4
$\epsilon \times 10^2$	4.99	4.99	4.99	4.93	4.98	5.00	5.04	4.93	4.94	4.95	5.03	5.00
λ	0.987	0.993	0.942	0.837	0.926	0.974	0.818	0.934	0.999	0.717	0.990	0.936
V_f	1.043	1.165	1.227	1.105	1.161	1.200	1.093	1.137	1.169	1.072	1.112	1.156
l_1	-0.326	-0.376	0.375	-1.131	-0.703	-0.559	-1.449	-0.989	-1.284	-1.075	-1.024	-0.843
$y_{\mathrm{LCO}} \times 10^2$	6.57	6.04	5.79	5.94	6.43	6.76	6.17	6.65	7.68	4.69	5.26	6.17

Table 6: NLTVA parameters ($\zeta, \gamma, \xi, \epsilon, \lambda$) and corresponding performance quantities ($V_f, l_1, y_{\mathrm{LCO}}$) for selected realizations of the DO, RO, RO-C_1, and RO-C_1-C_2 designs. The reported cases correspond to realizations whose plunge LCO amplitude y_{LCO} is closest to the 5%, 50%, and 95% quantiles of its distribution.

	DO			RO			RO-C_1			RO-C_1-C_2		
Value	q_5^y	q_{50}^y	q_{95}^y	q_5^y	q_{50}^y	q_{95}^y	q_5^y	q_{50}^y	q_{95}^y	q_5^y	q_{50}^y	q_{95}^y
ζ	0.112	0.0982	0.125	0.133	0.141	0.145	0.158	0.170	0.176	0.0903	0.0835	0.107
γ	0.480	0.478	0.459	0.437	0.441	0.455	0.423	0.434	0.438	0.389	0.404	0.390
$\xi \times 10^2$	0.237	0.417	0.412	0.0143	2.09	0.401	2.48	2.82	0.156	13.1	19.2	18.8
$\epsilon \times 10^2$	4.95	4.87	4.98	5.07	4.96	5.01	4.95	4.92	4.93	5.07	4.99	5.01
λ	0.750	0.995	0.997	0.800	0.910	0.991	0.808	0.894	0.985	0.800	0.988	0.983
V_f	1.151	1.077	1.220	1.124	1.164	1.191	1.106	1.133	1.158	1.076	1.140	1.134
l_1	-3.026	-0.302	-1.128	-0.858	-1.055	-1.204	-1.381	-1.585	-1.450	-1.044	-0.796	-1.036
$y_{\mathrm{LCO}} \times 10^2$	5.19	5.98	6.84	5.69	6.58	7.45	6.12	7.11	7.98	4.71	5.33	5.97

- *Harmonic balance method (HBM) and branch continuation.* The periodic solutions of the governing equation (37) are computed using the HBM. To this end, the ordinary differential equation is first transformed into an algebraic equation through a truncated multi-harmonic periodic approximation combined with *harmonic balance*. In practice, seven harmonics are found to be sufficient to obtain converged results. The resulting algebraic system is then solved while varying $\tilde{U}$ using path-continuation techniques with a pseudo arc-length corrector. The reader is referred to [51] for a general presentation of HBM-continuation and, for example, to [8] for their application to LCO prediction in aeroelastic simulations. One particularity here is that the oscillating frequency of the LCO is not known a priori and must be determined simultaneously with the amplitude. The additional equation is obtained by fixing the phase of the cycle, which is otherwise undertermined in autonomous systems. To initiate the continuation scheme on the periodic branch at the flutter speed—i.e., at a Hopf bifurcation, the branching point is first accurately located by exploiting its singularity property, after which the tangent associated with the unstable eigenvector is used [52, 53].

  This approach allows all periodic solutions and their stability to be tracked efficiently as the airspeed varies. The stability is assessed by monitoring the real part of the characteristics Floquet exponents [52]. It is computationally faster and provides complete branches, including unstable segments. Computation of quasi-periodic solutions created by the NLTVA effects would require accurately locating the Neimark-Sacker bifurcation point responsible for their emergence along the periodc solution branch and employing a two-dimensional Fourier series approximation [54]. These are not considered here.

- *Time integration.* A total of 300 airspeed values, linearly spaced, are considered, with a higher density of points near the flutter speed region to accurately capture the dynamics in this critical zone, while the remaining points extend up to $\tilde{U} = 1.8$. For the highest airspeed, the system is initialized with a small perturbation in plunge, $\boldsymbol{q}(0) = \begin{bmatrix}10^{-4} & 0 & 0\end{bmatrix}^{\mathrm{T}}$ and $\dot{\boldsymbol{q}}(0) = \boldsymbol{0}$, and integrated to a long final time $t = 5000$ to ensure decay of transients and convergence to any stable attractor. The final state of this simulation is then used as the initial condition for the next airspeed value, for which the integration is performed over a shorter time horizon of $t = 500$. This procedure is repeated sequentially for all airspeed values. Although computationally more expensive, this approach provides independent validation of the harmonic-balance results and captures quasi-periodic responses, when they arise. It provides only stable branches, however, and does not trace unstable periodic solutions.

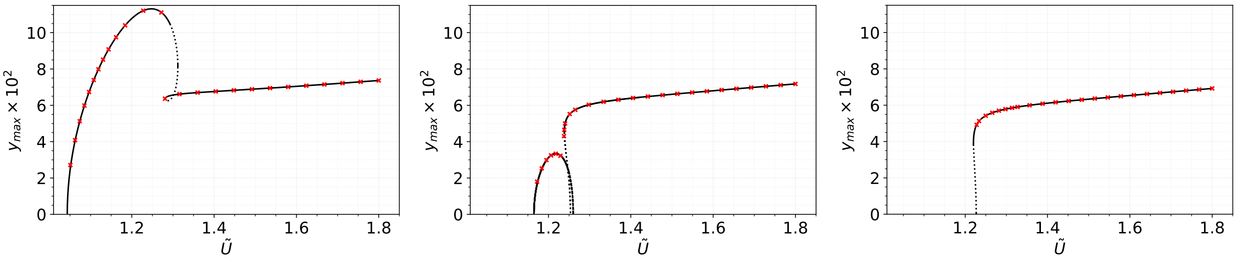

(a) DO-q_5^V (Quantile 5%) (b) DO-q_{50}^V (Quantile 50%) (c) DO-q_{95}^V (Quantile 95%)

Figure 7: Bifurcation diagrams of the plunge LCO amplitude y_{LCO} obtained from DO realizations associated with different quantiles of the V_f distribution. Results obtained using numerical time integration are compared with those from harmonic balance and continuation analysis. (—) Continuation branches, (×) time-integration results (only a subset of computed points is displayed for clarity), (···) unstable branches detected by continuation.

Figure 7 presents the resulting bifurcation curves of the plunge LCO amplitude obtained from numerical time integration and with the HBM and path-continuation for the three realizations associated with the DO design and related to the V_f distribution, namely DO-q_5^V, DO-q_{50}^V and DO-q_{95}^V. The bifurcation curves obtained from both methods are in very close agreement in the stable regions, highlighting only very minor differences. The other cases considered did not show any noticeable discrepancies either. Consequently, in the following, only the results obtained using the HBM are presented.

The DO-q_5^V case exhibits a supercritical bifurcation followed by two main stable branches separated by a complex intermediate structure involving unstable segments. Time-domain integration shows a jump from the lower-amplitude branch to the upper branch, leading to undesirable nonlinear and potentially unpredictable behavior. The medium quantile case, DO-q_{50}^V, displays atypical dynamics. Despite originating from a supercritical Hopf bifurcation, its bifurcation diagram consists of two disconnected branches, each terminating in short unstable segments near $\tilde{U} \approx 1.26$. In contrast, the upper quantile case, DO-q_{95}^V, exhibits a more standard behavior. It is characterized by a subcritical Hopf bifurcation and a single branch of periodic solutions with increasing amplitude.

Figure 8 presents the bifurcation curves associated with the four designs DO, RO, RO-C_1, and RO-C_1-C_2. The analysis is performed by considering the *horizontal variability* that triggers the Hopf instability, i.e., by examining three realizations associated with the distribution of V_f, together with the nominal case. The PDFs of V_f correspond to those shown in Fig. 6a, using the same color scheme. The DO design (Fig. 8a) exhibits widely spread behavior, with some realizations leading to rapid amplitude growth and undesirable nonlinear responses (see the 5% quantile). In contrast, the RO (Fig. 8b) design yields a more compact distribution, with the 5% quantile shifted toward higher values, as sought by the optimization problem. RO-q_{95}^V and RO-nom designs are almost indistinguishable up to slightly above $\tilde{U} = 1.2$. The RO-C$_1$ design (Fig. 8c) displays a behavior similar to RO but with a smoother onset of instability. This is consistent with their significantly smaller first Lyapunov coefficients, as reported in Table 5. This behavior is directly linked to constraint (C$_1$), which enforces the 5% quantile of l_1 to remain below -0.5, a condition initially introduced to guarantee supercritical behavior across all realizations of the associated design. However, this comes at the cost of generally higher LCO amplitudes. The RO-C$_1$-C$_2$ design (Fig. 8d) shows a similar qualitative behavior but significantly improves performance, yielding consistently lower LCO amplitudes across all realizations beyond the onset of instability. Notably, the ordering in flutter speed is generally inversely correlated to the amplitude ordering at higher airflow speeds, with the 5% quantile yielding the smallest LCO amplitude and the 95% quantile the largest.

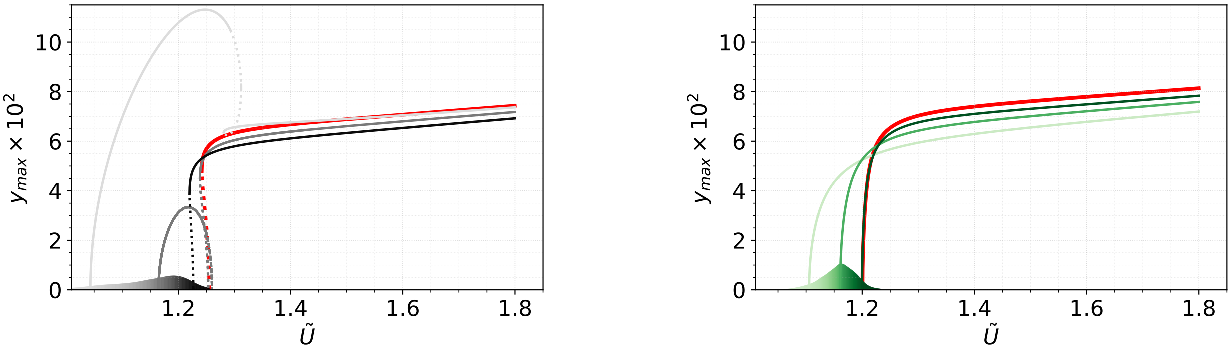

(a) DO

(b) RO

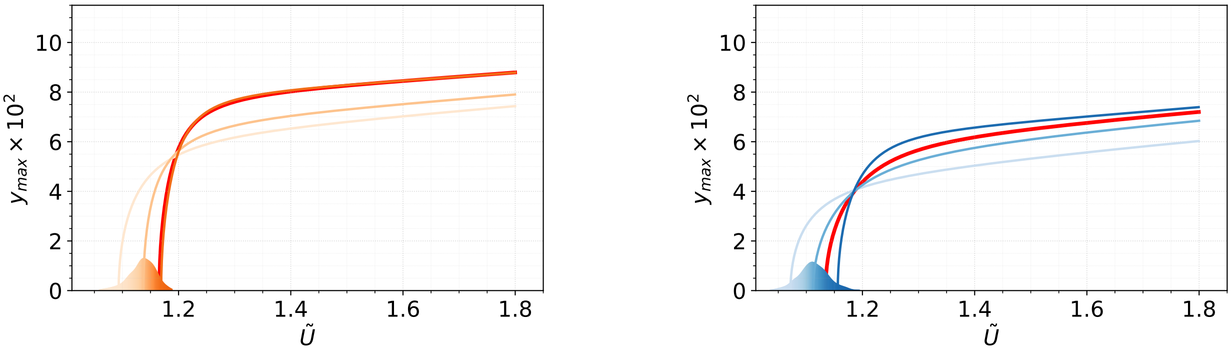

(c) RO-C_1

(d) RO-C_1-C_2

Figure 8: Bifurcation curves of the plunge LCO amplitude as a function of airspeed, obtained via harmonic balance and numerical path-continuation for four designs: (—, —, —) DO, (—, —, —) RO, (—, —, —) RO–C_1 and (—, —, —) RO–C_1–C_2. The different colors correspond to realizations whose flutter speed V_f is closest to the 5%, 50%, and 95% quantiles of the V_f distribution (from light to dark shading). The PDFs of V_f are also shown. (—) Nominal realization for each design.

Figure 9 enables a complementary analysis, this time focusing on the variability of $y_{\mathrm{LCO}} = y_{\max}(\tilde{U}^*)$. This is illustrated by the vertical line at $\tilde{U}^* = 1.3$ and the associated vertical PDFs of y_{LCO}. Such a representation provides a convenient way to assess variability *vertically*, in direct relation to the (C_2) constraint. The PDFs correspond to the curves shown in Fig. 6c, following the same color code. Similarly to Fig. 8, this analysis provides a global view of the entire bifurcation curve. This highlights that the PDF associated with RO-C_1-C_2 is shifted downward in Fig. 9[2] compared to RO and RO-C_1. The previously made observations are largely confirmed: the gap between the curves corresponding to different realizations, remains approximately constant from $\tilde{U}^* = 1.3$ up to $\tilde{U} = 1.8$. A notable distinction, however, is that the nominal curve lies close to the 75% quantile of y_{LCO} for each design. Although the ordering between realizations is less clear at lower airspeeds, a consistent hierarchy is recovered for $\tilde{U} > \tilde{U}^*$, with DO-q^y_{95} above DO-*nom*, followed by DO-q^y_{50} and DO-q^y_5. Overall, the RO-C_1-C_2 configuration yields the lowest $y_{\max}$ values.

The previous figures enabled the characterization of variability, for a given design, along the associated plunge bifurcation diagram. To further deepen the analysis and facilitate comparison between designs in terms of bifurcation behavior, Fig. 10a presents, for each design, the realization closest to the 5% quantile of the flutter speed V_f, along with the associated PDFs. These curves correspond to the lightest-colored curves shown in Fig. 8. The PDFs displayed here are simplified and reduced versions of those shown in Fig. 6a, which represents the distributions of V_f for all designs. Focusing on realizations closest to the 5% quantile reflects the objective of increasing this lower bound of the flutter speed. This trend is clearly observed: the unconstrained RO-q^V_5 (in green) achieves the highest values, followed closely by RO-C_1-q^V_5, which includes a constraint on l_1, then RO-C_1-C_2-q^V_5, and finally DO-q^V_5, the baseline design,

---

[2] equivalently, to the left in Fig. 6c

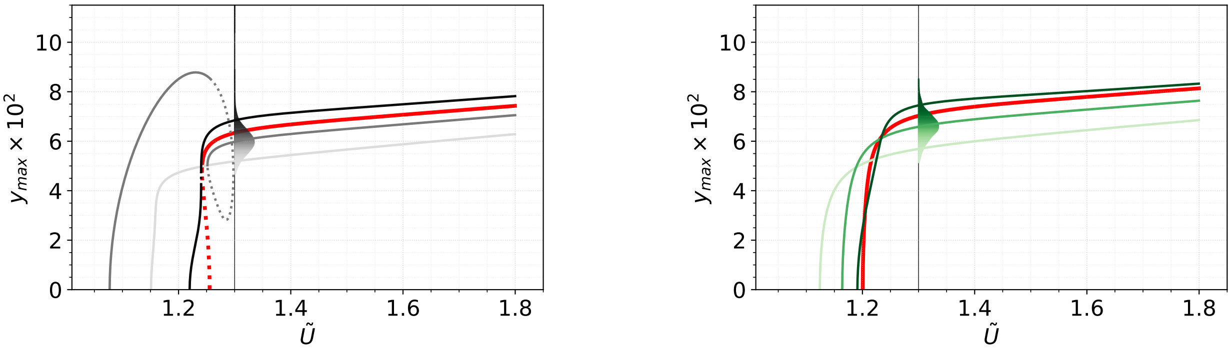

(a) DO

(b) RO

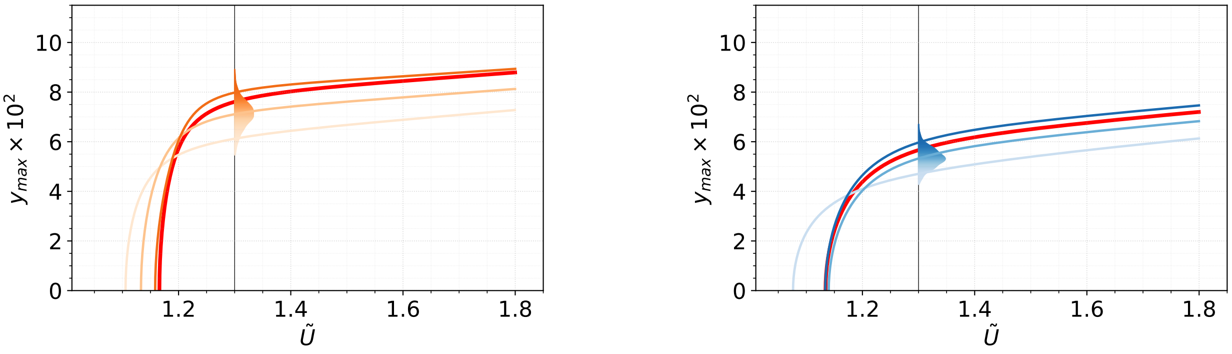

(c) RO-C_1

(d) RO-C_1-C_2

Figure 9: Bifurcation curves of the plunge LCO amplitude as a function of airspeed, obtained via harmonic balance and numerical path-continuation for four designs: (—, —, —) DO, (—, —, —) RO, (—, —, —) RO–C_1 and (—, —, —) RO–C_1–C_2. The different colors correspond to realizations whose post-flutter plunge LCO amplitude evaluated at $\tilde{U}^* = 1.3$ is closest to the 5%, 50%, and 95% quantiles of the y_{LCO} distribution (from light to dark shading). The PDFs of y_{LCO} are also shown. (—) Nominal realization for each design.

whose corresponding realization exhibits the lowest performance. Although RO-q_5^V yields lower $y_{\max}$ values at low airspeeds, it is overtaken by RO-C$_1$-C$_2$-q_5^V beyond $\tilde{U} \approx 1.15$, with a large and nearly constant gap. The RO-C$_1$-q_5^V curve remains slightly above RO-q_5^V, while DO-q_5^V generally exhibits the highest amplitude. Although no definitive conclusions can be drawn from these observations, they provide partial insight into the impact of the optimization formulation on the resulting bifurcation behavior.

On the other hand, Fig. 10b focuses on the value of the plunge LCO amplitude at $\tilde{U}^*$. For each of the four designs, the realization closest to the 95% quantile of y_{LCO} is shown, along with its PDF. These are simplified, vertically oriented versions of the distributions in Fig. 6c, following the same color code. It combines the darkest-colored lines shown in Fig. 9. As observed in Fig. 6c and in Table 4, the 95% quantiles of y_{LCO} for the four designs range from lowest to highest as RO-C$_1$-C$_2$-q_{95}^y, DO-q_{95}^y, RO-q_{95}^y, and RO-C$_1$-q_{95}^y. This ordering is clearly visible in Fig. 10b by examining the intersections of the bifurcation curves with the vertical line $\tilde{U} = \tilde{U}^*$. As expected, RO-C$_1$-C$_2$, the only design explicitly constructed to satisfy constraint (C$_2$), which enforces a 95% quantile of y_{LCO} lower than that of the baseline DO design, fulfills this requirement. In contrast, the RO-q_{95}^y and RO-C$_1$-q_{95}^y realizations yield higher $y_{\max}$ values over most of the airspeed range, except before they intersect the RO-C$_1$-C$_2$-q_{95}^y curve, whose instability occurs at a lower airspeed. This representation provides understanding into how the (C$_2$) constraint applied at a single airspeed $\tilde{U}^*$ and targeting a specific realization—here set at the 95% quantile but potentially set to other values, affects the entire bifurcation curve for that realization over the full airspeed range. More broadly, it facilitates interpretation of the global impact of the constraint among the ensemble of realizations, as previously illustrated in Fig. 9.

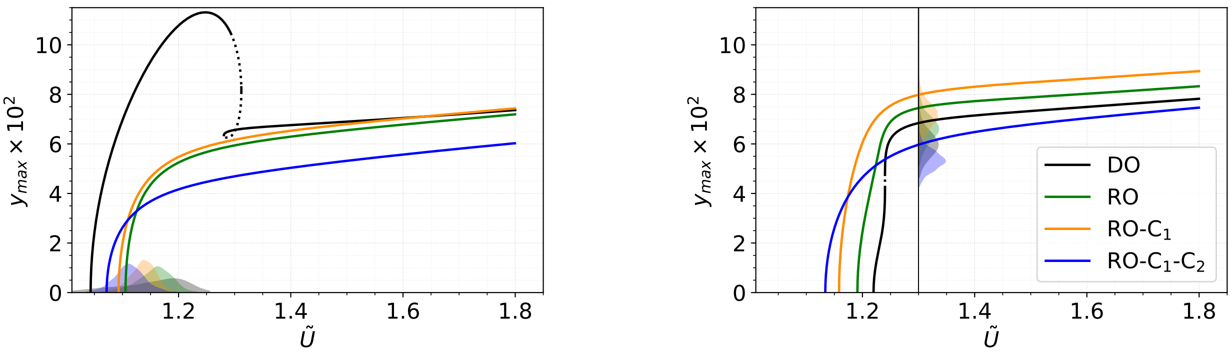

(a) Realizations closest to the 5% quantile of V_f

(b) Realizations closest to the 95% quantile of y_{LCO}

Figure 10: Bifurcation curves of the plunge LCO as a function of airspeed, obtained via harmonic balance and numerical path-continuation for the four designs (—) DO, (—) RO, (—) RO-C_1, and (—) RO-C_1-C_2. (a): curves corresponding to the realization whose flutter speed V_f is closest to the 5% quantile of the V_f distribution. (b): curves corresponding to the realization whose plunge LCO amplitude evaluated at $\tilde{U}^* = 1.3$ (vertical black line) is closest to the 95% quantile of the y_{LCO} distribution. The PDFs are indicated accordingly.

6. Conclusion

This work presented a surrogate-assisted methodology for robust optimization under uncertainty, with a focus on problems where design variables are subject to inherent variability.

The methodology was first validated on two analytical test cases without variability in the design variables. The results confirmed the ability of the approach to accurately identify global optima in both unconstrained and constrained stochastic settings, with comparable or improved accuracy and a similar or reduced number of evaluations compared to recent state-of-the-art methods.

The main application of this study was the robust design of a nonlinear tuned vibration absorber (NLTVA) for passive flutter control under aleatory variability. Five random inputs were considered to represent aleatory variability in the absorber parameters, among which three were also treated as design variables and thus considered uncertain. Several optimization problems were formulated, all maximizing a given quantile of the flutter speed, with different configurations of quantile-based constraints on the first Lyapunov coefficient—to ensure supercritical behavior, and LCO amplitudes. The resulting designs were obtained with very few model evaluations and compared to a deterministic optimal baseline. While the deterministic design performs well under nominal conditions, it exhibits large variability in flutter speed and may lead to subcritical Hopf instabilities. The unconstrained robust optimization results in a design that significantly enhances worst-case performance and shows excellent agreement with a reference solution obtained at a higher computational cost. Enforcing supercritical Hopf bifurcation behavior yields a well-behaved distribution of flutter speeds without jump phenomena, at the cost of increased plunge LCO amplitudes. This limitation was mitigated by introducing an additional constraint on the plunge amplitude.

Further analysis of bifurcation curves for selected NLTVA parameter realizations confirmed the effectiveness of the LCO constraint, demonstrating improved control of amplitudes across both airspeed and variability, and illustrated how a constraint imposed at a single airspeed influences the entire bifurcation curve. This realistic engineering application highlights the importance of accounting for uncertainties in the parameters defining the vibration control system and of using appropriate constraints. The focus on plunge amplitude limitation was motivated by its richer bifurcation behaviour compared to pitch. While this complexity is effectively controlled by the present methodology, other choices could have been considered, such as targeting the pitch response, both degrees of freedom simultaneously, or alternative quantities such as the absorber position. Future work could also extend this approach to enforce the constraint explicitly over multiple airspeeds or the full range, for example using a multipoint or area-based formulation.

The proposed methodology relies on a coupled design-stochastic space metamodel based on kriging, combined with an original parallel adaptive infill strategy. This enables efficient exploration of the joint space and the optimization of quantile-based objectives and constraints with a limited number of expensive function evaluations. The framework could be extended to any surrogate model providing mean and variance predictions. As a result, it inherits typical limitations of kriging, in particular reduced efficiency in high-dimensional spaces.

Quantiles were estimated through Monte Carlo sampling, and the use of AK-MCS-type infill strategies [33, 38] restricts the range of accessible quantile levels to moderate values: estimating extreme quantiles would require prohibitively large Monte Carlo sample sizes, leading to a excessively high computational cost for surrogate evaluations and candidate selection. Nevertheless, this work opens perspectives for the development of reliability-based robust optimization methods capable of handling very small quantile levels and failure probabilities (e.g., down to 10^{-9}), for instance by using advanced extensions of AK-MCS such as those proposed in [39, 40].

References

[1] D. H. Hodges, G. A. Pierce, Introduction to Structural Dynamics and Aeroelasticity, 2nd Edition, Cambridge Aerospace Series, Cambridge University Press, 2011. doi:10.1017/CBO9780511997112.

[2] K. V. Singh, Active aeroelastic control with time delay for targeted flutter modes, Aerospace science and technology 43 (2015) 281–288. doi:10.1016/j.ast.2015.03.009.

[3] E. Livne, Aircraft active flutter suppression: State of the art and technology maturation needs, Journal of Aircraft 55 (1) (2018) 410–452. doi:10.2514/1.c034442.

[4] V. Mukhopadhyay, Historical perspective on analysis and control of aeroelastic responses, Journal of Guidance, Control, and Dynamics 26 (5) (2003) 673–684. doi:10.2514/2.5108.

[5] N. Jalili, A comparative study and analysis of semi-active vibration-control systems, Journal of Vibration and Acoustics 124 (4) (2002) 593–605. doi:10.1115/1.1500336.

[6] J. G. Pérez, A. Ghadami, L. Sanches, G. Michon, B. I. Epureanu, Data-driven optimization for flutter suppression by using an aeroelastic nonlinear energy sink, Journal of Fluids and Structures 114 (2022) 103715. doi:10.1016/j.jfluidstructs.2022.103715.

[7] K. Dhital, B. Chouvion, Passive aeroelastic control of a near-ground airfoil with a nonlinear vibration absorber, Aerospace 11 (2024) 1043. doi:10.3390/aerospace11121043.

[8] R. Alcorta, B. Chouvion, G. Michon, O. Montagnier, On the use of frictional dampers for flutter mitigation of a highly flexible wing, International Journal of Non-Linear Mechanics 156 (2023) 104515. doi:10.1016/j.ijnonlinmec.2023.104515.

[9] Y. S. Lee, G. Kerschen, D. M. McFarland, W. J. Hill, C. Nichkawde, T. W. Strganac, L. A. Bergman, A. F. Vakakis, Suppressing aeroelastic instability using broadband passive targeted energy transfers, part 2: experiments, AIAA journal 45 (10) (2007) 2391–2400. doi:10.2514/1.28300.

[10] W. Tian, Y. Li, P. Li, Z. Yang, T. Zhao, Passive control of nonlinear aeroelasticity in hypersonic 3-d wing with a nonlinear energy sink, Journal of Sound and Vibration 462 (2019) 114942. doi:10.1016/j.jsv.2019.114942.

[11] A. Malher, C. Touzé, O. Doaré, G. Habib, G. Kerschen, Flutter control of a two-degrees-of-freedom airfoil using a nonlinear tuned vibration absorber, Journal of Computational and Nonlinear Dynamics 12 (5) (2017) 051016. doi:10.1115/1.4036420.

[12] E. Verstraelen, G. Habib, G. Kerschen, G. Dimitriadis, Experimental passive flutter suppression using a linear tuned vibration absorber, AIAA Journal 55 (5) (2017) 1707–1722. doi:10.2514/1.J055397.

[13] C. Fernandez-Escudero, S. Prothin, E. Laurendeau, A. Ross, G. Michon, Nonlinear flap for passive flutter control of bidimensional wing, Journal of Vibration and Control 0 (0) (2023). doi:10.1177/10775463231223778.

[14] J.-C. Chassaing, C. Nitschke, A. Vincenti, P. Cinnella, D. Lucor, Advances in parametric and model-form uncertainty quantification in canonical aeroelastic systems, Aerospace Lab 14 (2018) 1–19. doi:10.12762/2018.AL14-07.

[15] C. L. Pettit, Uncertainty quantification in aeroelasticity: Recent results and research challenges, Journal of Aircraft 41 (5) (2004) 1217–1229. doi:10.2514/1.3961.

[16] Y. Dai, C. Yang, Methods and advances in the study of aeroelasticity with uncertainties, Chinese Journal of Aeronautics 27 (3) (2014) 461–474. doi:10.1016/j.cja.2014.04.003.

[17] H. Tran, C. G. Webster, G. Zhang, A sparse grid method for bayesian uncertainty quantification with application to large eddy simulation turbulence models, in: J. Garcke, D. Pflüger (Eds.), Sparse Grids and Applications - Stuttgart 2014, Springer International Publishing, Cham, 2016, pp. 291–313. doi:10.1007/978-3-319-28262-6_12.

[18] C. Wu, H. Zhang, T. Fang, Flutter analysis of an airfoil with bounded random parameters in incompressible flow via gegenbauer polynomial approximation, Aerospace Science and Technology 11 (7-8) (2007) 518–526. doi:10.1016/j.ast.2007.01.004.

[19] S. Hosder, R. W. Walters, M. Balch, Point-collocation nonintrusive polynomial chaos method for stochastic computational fluid dynamics, AIAA Journal 48 (12) (2010) 2721–2730. doi:10.2514/1.39389.

[20] E. Denimal, J.-J. Sinou, Efficient parametric study of a stochastic airfoil system based on hybrid surrogate modeling with advanced automatic kriging construction, European Journal of Mechanics - A/Solids 99 (2023) 104926. doi:10.1016/j.euromechsol.2023.104926.

[21] M. Rivier, N. Razaaly, P. M. Congedo, Non-parametric measure approximations for constrained multi-objective optimisation under uncertainty, International Journal for Numerical Methods in Engineering 125 (7) (2024) e7403. doi:10.1002/nme.7403.

[22] N. Razaaly, G. Persico, G. Gori, P. M. Congedo, Quantile-based robust optimization of a supersonic nozzle for organic rankine cycle turbines, Applied Mathematical Modelling 82 (2020) 802–824. doi:10.1016/j.apm.2020.01.048.

[23] C. Sabater, O. L. Maitre, P.-M. Congedo, S. Görtz, A bayesian approach for quantile optimization problems with high-dimensional uncertainty sources, Computer Methods in Applied Mechanics and Engineering 358 (2020) 112623. doi:10.1016/j.cma.2019.112623.

[24] Y. Aoues, A. Chateauneuf, Benchmark study of numerical methods for reliability-based design optimization, Structural and multidisciplinary optimization 41 (2) (2010) 277–294. doi:10.1007/s00158-009-0412-2.

[25] M. A. Valdebenito, G. I. Schuëller, A survey on approaches for reliability-based optimization, Structural and Multidisciplinary Optimization 42 (2010) 645–663. doi:10.1007/s00158-010-0518-6.

[26] S. Missoum, C. Dribusch, P. Beran, Reliability-based design optimization of nonlinear aeroelasticity problems, AIAA Journal 47 (3) (2009) 741–751. doi:10.2514/1.46665.

[27] R. E. Melchers, A. T. Beck, Structural reliability analysis and prediction, John wiley & sons, 2018.

[28] C. Sabater, P. Bekemeyer, S. Görtz, Robust design of transonic natural laminar flow wings under environmental and operational uncertainties, AIAA Journal 60 (2) (2022) 767–782. doi:10.2514/1.J060676.

[29] M. Rivier, P. M. Congedo, Surrogate-assisted bounding-box approach applied to constrained multi-objective optimisation under uncertainty, Reliability Engineering & System Safety 217 (2022) 108039. doi:10.1016/j.ress.2021.108039.

[30] M. Moustapha, B. Sudret, Surrogate-assisted reliability-based design optimization: a survey and a unified modular framework, Structural and Multidisciplinary Optimization 60 (5) (2019) 2157–2176. doi:10.1007/s00158-019-02290-y.

[31] C. K. Williams, C. E. Rasmussen, Gaussian Processes for Machine Learning, the MIT Press, 2006.

[32] V. Dubourg, B. Sudret, J.-M. Bourinet, Reliability-based design optimization using kriging surrogates and subset simulation, Structural and Multidisciplinary Optimization 44 (2011) 673–690. doi:10.1007/s00158-011-0653-8.

[33] M. Moustapha, B. Sudret, J.-M. Bourinet, B. Guillaume, Quantile-based optimization under uncertainties using adaptive kriging surrogate models, Structural and Multidisciplinary Optimization 54 (6) (2016) 1403–1421. doi:10.1007/s00158-016-1504-4.

[34] T. J. Santner, B. J. Williams, W. I. Notz, The Design and Analysis of Computer Experiments, Springer Series in Statistics, Springer, New York, 2003. doi:10.1007/978-1-4757-3799-8.

[35] O. Roustant, D. Ginsbourger, Y. Deville, DiceKriging, DiceOptim: Two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization, Journal of statistical software 51 (2012) 1–55. doi:10.18637/jss.v051.i01.

[36] M. A. Bouhlel, N. Bartoli, A. Otsmane, J. Morlier, An improved approach for estimating the hyperparameters of the kriging model for high-dimensional problems through the partial least squares method, Mathematical Problems in Engineering 2016 (1) (2016) 6723410. doi:10.1155/2016/6723410.

[37] N. Bartoli, T. Lefebvre, S. Dubreuil, R. Olivanti, R. Priem, N. Bons, J. R. Martins, J. Morlier, Adaptive modeling strategy for constrained global optimization with application to aerodynamic wing design, Aerospace Science and technology 90 (2019) 85–102. doi:10.1016/j.ast.2019.03.041.

[38] R. Schöbi, B. Sudret, S. Marelli, Rare event estimation using polynomial-chaos kriging, ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering 3 (2) (2017) D4016002. doi:10.1061/AJRUA6.0000870.

[39] N. Razaaly, P. M. Congedo, Extension of AK-MCS for the efficient computation of very small failure probabilities, Reliability Engineering and System Safety 202 (2020) 107057. doi:10.1016/j.ress.2020.107057.

[40] N. Razaaly, P. M. Congedo, Efficient estimation of extreme quantiles using adaptive kriging and importance sampling, Numerical Methods for Engineering 121 (1) (2020) 1–22. doi:10.1002/nme.6300.

[41] J. Bect, D. Ginsbourger, L. Li, V. Picheny, E. Vazquez, Sequential design of computer experiments for the estimation of a probability of failure, Statistics and Computing 22 (3) (2012) 773–793. doi:10.1007/s11222-011-9241-4.

[42] B. J. Bichon, M. S. Eldred, L. P. Swiler, S. Mahadevan, J. M. McFarland, Efficient global reliability analysis for nonlinear implicit performance functions, AIAA journal 46 (10) (2008) 2459–2468. doi:10.2514/1.34321.

[43] B. Echard, N. Gayton, M. Lemaire, AK-MCS: An active learning reliability method combining kriging and monte carlo simulation, Structural Safety 33 (2) (2011) 145–154. doi:10.1016/j.strusafe.2011.01.002.

[44] D. R. Jones, M. Schonlau, W. J. Welch, Efficient global optimization of expensive black-box functions, Journal of Global Optimization 13 (4) (1998) 455–492. doi:10.1023/A:1008306431147.

[45] Z.-H. Han, SurroOpt: a generic surrogate-based optimization code for aerodynamic and multidisciplinary design, in: Proceedings of ICAS, Vol. 2016, 2016, pp. 2016–0281.

[46] P. Virtanen, et al., SciPy 1.0: Fundamental algorithms for scientific computing in python, Nature Methods 17 (2020) 261–272. doi:10.1038/s41592-019-0686-2.

[47] M. D. McKay, R. J. Beckman, W. J. Conover, A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, Technometrics 21 (2) (1979) 239–245. doi:10.2307/1268522.

[48] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in python, Journal of Machine Learning Research 12 (2011) 2825–2830. doi:10.5555/1953048.2078195.

[49] Y. A. Kuznetsov, Elements of applied bifurcation theory, Springer, 1998.

[50] D. R. Jones, C. D. Perttunen, B. E. Stuckman, Lipschitzian optimization without the lipschitz constant, Journal of Optimization Theory and Applications 79 (1) (1993) 157–181. doi:10.1007/BF00941892.

[51] M. Krack, J. Gross, Harmonic balance for nonlinear vibration problems, Springer, 2019. doi:10.1007/978-3-030-14023-6.

[52] R. Seydel, Practical Bifurcation and Stability Analysis, Vol. 5 of Interdisciplinary Applied Mathematics, Springer-Verlag New York, 2010. doi:10.1007/978-1-4419-1740-9.

[53] N. Di Palma, B. Chouvion, F. Thouverez, Parametric study on internal resonances for a simplified nonlinear blade model, International Journal of Non-Linear Mechanics 141 (2022) 103941. doi:10.1016/j.ijnonlinmec.2022.103941.

[54] R. Alcorta, B. Chouvion, O. Montagnier, Dynamics of a non-linear jeffcott rotor in supercritical regime, International Journal of Non-Linear Mechanics 148 (2023) 104272. doi:10.1016/j.ijnonlinmec.2022.104272.