

Horsetail Matching for Optimization Under Probabilistic, Interval and Mixed Uncertainties

Laurence W. Cook* and Jerome P. Jarrett†

University of Cambridge, Cambridge, England, CB21PZ.

Karen E. Willcox‡

Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, 02139

The importance of including uncertainties in the design process of aerospace systems is becoming increasingly recognized, leading to the recent development of many techniques for optimization under uncertainty. Most existing methods represent uncertainties in the problem probabilistically; however, in many real life design applications it is often difficult to assign probability distributions to uncertainties without making strong assumptions. Existing approaches for optimization under different types of uncertainty mostly rely on treating combinations of statistical moments as separate objectives, but this can give rise to stochastically dominated designs. Horsetail matching is a flexible approach to optimization under any mix of probabilistic and interval uncertainties that overcomes some of the limitations of existing approaches. The formulation delivers a single, differentiable metric as the objective function for optimization. It is demonstrated on algebraic test problems and the design of a flying wing using a coupled aero-structural analysis code.

I. Introduction

Optimization techniques are becoming increasingly integrated within the design process of complex aerospace systems when a computational simulation of the system is available. Traditionally, an optimization considers a quantity of interest of a system q (e.g., efficiency, cost, weight) as a function of controllable design variables \mathbf{x} , and the objective is to find a design \mathbf{x}^* such that

$$\begin{aligned} \forall \mathbf{x} \in \mathcal{X}, q(\mathbf{x}^*) &\leq q(\mathbf{x}) \\ \text{s.t. } x_k^L &< x_k < x_k^U \quad k = 1, \dots, n_x \\ g_j(\mathbf{x}) &\leq 0 \quad j = 1, \dots, n_g \end{aligned} \tag{1}$$

where x_k^L and x_k^U are respectively the upper and lower bounds on the k^{th} design variable that make up the design space \mathcal{X} , n_x is the number of design variables, $g_j(\mathbf{x})$ is the j^{th} inequality constraint, and n_g is the number of constraints. However, in practical problems, estimating the objective q will be subject to many uncontrollable uncertainties from a variety of sources.^{1,2} We denote the uncertainties by $\mathbf{u} \in \mathcal{U}$, where \mathcal{U} is the uncertainty space. Uncertainty quantification (UQ) has emerged as an important field for computational analysis, where effective and efficient methods for quantifying the influence of the uncertainties on the quantity of interest are being developed. Furthermore, the importance of including uncertainties in the design process is becoming increasingly recognized, since deterministically optimized designs often lie in extreme regions of design space and see severely degraded performance when realized.^{3,4} Therefore effective computational methods of handling uncertainties within the optimization process are needed, this is addressed by the field of optimization under uncertainty (OUU).

Handling uncertainties in optimization is not a challenge unique to engineering design, and has been considered extensively in other contexts. The field of operations research has developed mathematically

*PhD Student, Department of Engineering. AIAA Student Member.

†Lecturer, Department of Engineering. AIAA Senior Member.

‡Professor, Department of Aeronautics and Astronautics. AIAA Associate Fellow.

rigorous robust optimization techniques for linear problems.^{5,6} Similarly the field of stochastic programming has developed methods for probabilistic optimization under uncertainty that preserve mathematical structure of the underlying problem such as linearity and convexity.⁷ Engineering design differs from these fields mainly because the design problems are often non-linear, and so we cannot exploit any underlying structure of the problem when formulating the OUU problem. Nonetheless, there are important concepts from these fields for formulating the OUU problem effectively that should not be overlooked in engineering design.

The first step in any quantification of uncertainty, whether for analysis or design, is representing the uncertainties in \mathbf{u} mathematically. A typical approach is to assign a probability distribution to each entry in the vector \mathbf{u} (and so \mathbf{u} becomes a multidimensional random variable with independent components). However, uncertainties arise from a variety of sources and so this is not always the most appropriate choice. Indeed, various representations of uncertainty are available, giving the designer a richer set of options for characterizing the differing types of uncertainty that may be present in a given problem.⁸ A common alternative to representing uncertainties probabilistically is to use intervals, where bounds on the values that the uncertainty can take are all that is known. For these uncertainties assigning a probability distribution assumes structure that is unknown in reality;^{9,10} an interval representation avoids imposing this extra structure.

In this paper the uncertainties in \mathbf{u} are split into uncertainties that are treated probabilistically, θ , and uncertainties that are treated with interval analysis, ϕ ; the quantity of interest q is now a function of the design variables as well as these uncertainties: $q = q(\mathbf{x}, \theta, \phi)$. In general problems, both of these types of uncertainties can exist and so UQ for mixed probabilistic and interval uncertainties can be used. This mixed UQ analysis gives rise to a cumulative distribution function (CDF) for every possible value of the interval uncertainties, and then the envelope gives the interval at every probability level and the upper and lower bounds on the true CDF.^{10,11} This uncertainty information is illustrated in Figure 1 (where a small subset of infinite possible CDFs are plotted), and is referred to as a “horsetail plot” (also known as a p-box), giving rise to the name of our proposed approach.

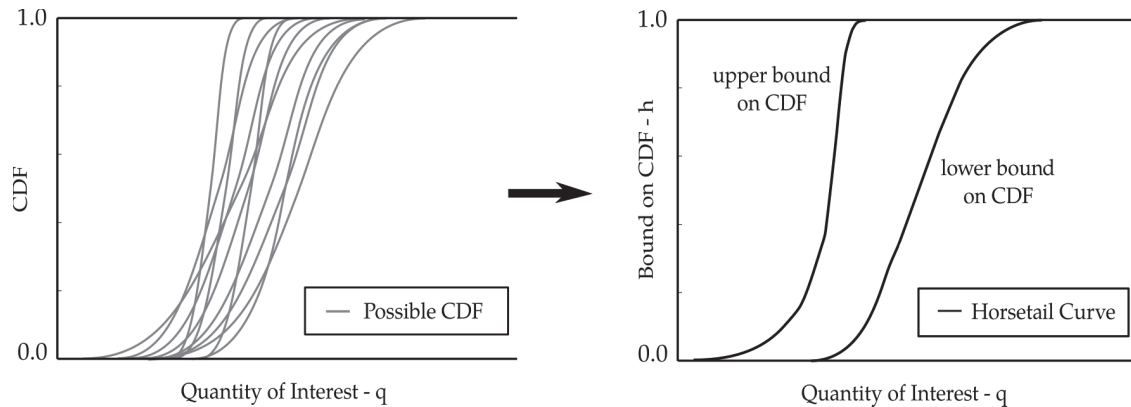


Figure 1: Example of a horsetail plot - the envelope of possible CDFs giving the two horsetail curves that represent the upper and lower bounds on the true CDF.

Other methods of propagating uncertainties beyond just probability and intervals are available, such as the Dempster-Shafer theory of evidence,¹² and possibility theory.¹³ These methods also result in a horsetail plot, in that they give upper and lower bounds on the true CDF of the quantity of interest. A horsetail plot is thus a rather general form of the output of an uncertainty analysis. A probabilistic analysis yields a special case of a horsetail plot where the upper and lower bounds on the CDF are coincidental. A purely interval based analysis yields another special case where the bounds are step functions.

The following definition of a horsetail plot is used in this paper (where the functions F_u and F_l are referred to as the horsetail curves):

Definition I.1. A horsetail plot consists of two non-decreasing functions $F_u: \mathcal{Q} \rightarrow [0, 1]$ and $F_l: \mathcal{Q} \rightarrow [0, 1]$ that represent the upper and lower bounds on the CDF respectively, where \mathcal{Q} is the set of all possible values of the quantity of interest q .

Much of the previous work into OUU in aerospace applications has considered exclusively probabilistic uncertainties where most commonly the first two statistical moments of q (namely mean and variance) are

considered in an optimization. For robust optimization these can be optimized separately in a multi-objective formulation,^{14–16} or can be combined into a single objective using a weighted sum approach.^{17–20} Various other methods for robust optimization have also been developed.¹ A notable alternative to statistical moment based methods is the “robust regularization” approach,^{1,21} which is a minimax strategy optimizing the worst case performance over the uncertainty space (it is referred to as the “robust counterpart” in the context of convex programming²²), so it treats the uncertainties as intervals.

Recently, matching approaches have been developed in order to try and overcome some of the disadvantages (discussed below) of statistical moment based methods. For example, a CDF based matching approach minimizes the area between a design’s cumulative density function and a step function at the deterministic optimum.²³ A PDF matching approach minimizes the distance between a design’s probability density function and a target, measured by a differentiable squared L_2 -norm metric.²⁴

Examples of robust optimization under mixed uncertainties to date almost exclusively use a weighted sum combination of average statistical moments and intervals of statistical moments.^{25–28} For example,

$$f_{obj} = w_1\bar{\mu} + w_2\bar{\sigma} + w_3\delta\mu \quad (2)$$

where f_{obj} is the objective function to be minimized, $\bar{\mu}$ is the average mean of the CDFs in the horsetail plot, $\bar{\sigma}$ is the average standard deviation, $\delta\mu$ is the interval of possible μ values, and w_1, w_2, w_3 are weights specified by the designer. Another suggested approach for robust design under mixed uncertainties is to minimize the interval of the 50% quantile.¹¹

Under probabilistic uncertainties, methods that consider statistical moments as separate objectives assume it is always worth trading off mean performance for reduced variance (increased robustness), but this is not the case in many problems if it gives rise to a design that is stochastically dominated by another design.²⁹ Stochastic dominance indicates that for any given value of the quantity of interest q , a dominated design is less likely to obtain this value or better than another design; this is a concept the authors argue is often overlooked in the engineering OOU literature. Stochastic dominance under probabilistic uncertainties is illustrated in Figure 2, and is defined by Definition I.2; it essentially means that the CDFs for the designs do not cross at any point.




Figure 2: Stochastic dominance under probabilistic uncertainties (left) and mixed uncertainties (right). In both cases design x_A is stochastically dominated by design x_B , according to definitions I.2 and I.3.

Definition I.2. A design x_A *stochastically dominates* a design x_B (or design x_B is *stochastically dominated by* design x_A) under probabilistic uncertainties if

$$\forall q \in \mathcal{Q}, F(q)_{x_A} > F(q)_{x_B}$$

or equivalently

$$\forall h \in [0, 1], F^{-1}(h)_{x_A} < F^{-1}(h)_{x_B}$$

where $F(q)_x: \mathcal{Q} \rightarrow [0, 1]$ and $F^{-1}(h)_x: [0, 1] \rightarrow \mathcal{Q}$ are respectively the CDF and inverse CDF for design x .

The concept of stochastic dominance extends straightforwardly to mixed uncertainties, by defining it as when each horsetail curve for a design stochastically dominate the corresponding curves for another design according to Definition I.2. This is illustrated in Figure 2 and defined in Definition I.3:

Definition I.3. A design x_A *stochastically dominates* a design x_B (or design x_B is *stochastically dominated* by design x_A) under mixed uncertainties if

$$\forall q \in \mathcal{Q}, F_u(q)_{x_A} > F_u(q)_{x_B} \text{ AND } F_l(q)_{x_A} > F_l(q)_{x_B}$$

or equivalently:

$$\forall h \in [0, 1], F_u^{-1}(h)_{x_A} < F_u^{-1}(h)_{x_B} \text{ AND } F_l^{-1}(h)_{x_A} < F_l^{-1}(h)_{x_B}$$

where F_u and $F_l: \mathcal{Q} \rightarrow [0, 1]$ are respectively the upper and lower bound of the CDF, and F_u^{-1} and $F_l^{-1}: [0, 1] \rightarrow \mathcal{Q}$ are respectively the corresponding inverses.

Stochastic dominance under mixed uncertainties means that for any value of q , the bounds on the probability that a dominating design will achieve this performance or better are strictly higher than the dominated design.

The basic robust optimization philosophy aims to maximize the likelihood of achieving as good performance as possible. A stochastically dominated design is less likely to achieve a given value of q or better than another design, and thus the dominated design is objectively worse than the dominating design. Treating combinations of statistical moments such as $\bar{\mu}$, $\bar{\sigma}$, and $\delta\mu$ as separate objectives in an optimization does not consider the possibility of obtaining stochastically dominated designs. The implication of this is that, for some design problems, many of the designs on a Pareto front trading off these objectives are stochastically dominated by other designs and computational effort has been wasted obtaining them. Similarly if a weighted sum of these objectives is being optimized, some combinations of weightings will give rise to dominated designs, and since it is difficult to know *a priori* where on the Pareto front a given set of weightings will end up, this is a limitation of optimizing under mixed uncertainties using combinations of statistical moments.

Additionally, alternative approaches for optimization under mixed uncertainties are few; it was noted in Ref. 1 that optimization under epistemic uncertainty has received relatively little attention in the literature (epistemic uncertainties are most commonly modeled with intervals). Therefore there is room for significant improvement in techniques for optimization under different types of uncertainty.

These arguments motivate the development of horsetail matching under mixed uncertainties presented in this paper. In Section II the general concept is presented, the advantages of the approach compared to existing methods are discussed, and the numerical implementation is outlined. In Section III, the effectiveness of the implementation is demonstrated, and in Section IV horsetail matching is applied to a physical design problem: the low-fidelity aero-structural optimization of a flying wing, where it is compared to alternative optimization approaches. Finally Section V concludes the paper and proposes future work.

II. Horsetail Matching

This section defines the difference metric underlying our horsetail matching approach, discusses its flexibility and describes its implementation.

II.A. The Difference Metric

The concept of horsetail matching is to minimize the difference between the horsetail plot of the current design and a target, as illustrated in Figure 3.

The measure of this difference is given by the following horsetail matching metric:

$$d_{hm}(\mathbf{x}) = \left(\int_0^1 (F_u^{-1}(h) - t_u(h))^2 dh + \int_0^1 (F_l^{-1}(h) - t_l(h))^2 dh \right)^{1/2} \quad (3)$$

where $F_u^{-1}(h)$ is the inverse of the upper bound of the CDF, $F_l^{-1}(h)$ the inverse of the lower bound (which both exist because by definition the bounds are non-decreasing), and $t_u(h)$ and $t_l(h)$ are respectively the targets for the upper and lower bound horsetail curves. Note that this target does not necessarily have to consist of the inverse of valid CDFs; it can be any pair of functions of h . The metric d_{hm} is the L_2 norm of the difference between the horsetail curves and their targets integrated over h . In our horsetail matching formulation, the overall OUU problem becomes finding \mathbf{x}^* , such that:

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmin}} d_{hm}(\mathbf{x}; t_u, t_l) \quad (4)$$

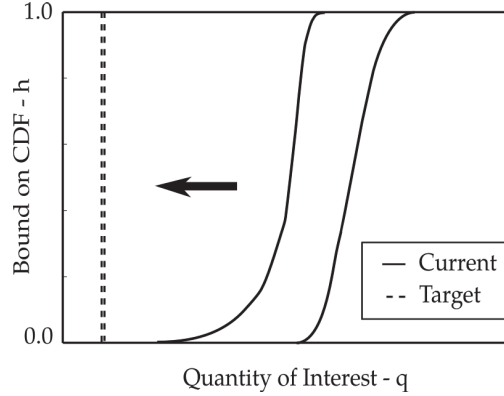


Figure 3: The horsetail matching concept: minimizing the difference between the horsetail plot of the quantity of interest q for the current design and a target.

where the design given by \mathbf{x}^* corresponds to the optimal design under uncertainty - its behavior under uncertainty is as close as possible to that specified in the target. Although this formulation might appear to be restricting, we show in the next section that it is fact flexible enough to recover more traditional robust optimization formulations.

Special Cases

Under exclusively probabilistic uncertainties, the metric reduces to:

$$d_{hm}(\mathbf{x}) = \left(2 \int_0^1 (F_{cdf}^{-1}(h) - t_{cdf}(h))^2 dh \right)^{1/2} \quad (5)$$

where $F_{cdf}^{-1}(h)$ is the inverse of the CDF and $t_{cdf}(h)$ is its target. Horsetail matching in the probabilistic case has been explored in detail by the first two authors.³⁰ Compared to density matching,²⁴ it has been shown to be superior both in terms of computational efficiency and producing good designs from a robust design point of view. It has also been shown to give designs that are not stochastically dominated (see Definition I.2) but which are guaranteed to lie on the mean/variance Pareto front, at a similar computational cost to the weighted sum of mean and variance approach.³⁰

Under exclusively interval uncertainties, the metric reduces to:

$$d_{hm}(\mathbf{x}) = \left((q_{min} - t_u)^2 + (q_{max} - t_l)^2 \right)^{1/2} \quad (6)$$

where q_{min} and q_{max} are the minimum and maximum values of q over all possible values of the uncertainties, and t_u and t_l are their respective target values.

II.B. Flexibility of the Metric

It might seem that requiring a target places a lot of responsibility on the designer, which was an issue discussed for the density matching approach.²⁴ However, using horsetail matching, it is straightforward to provide a target that captures the robust design philosophy of maximizing the likelihood of achieving as good performance as possible. This is achieved by setting both $t_u(h)$ and $t_l(h)$ to be the same constant value of q beyond what is feasible (e.g., zero weight, 100% efficiency). This is the target illustrated in Figure 3, and is referred to as the “standard target” for horsetail matching. Under this target, the minimizer of d_{hm} from Eq. 3 will not be stochastically dominated by any other feasible design according to definition I.3, a property the authors argue is an improvement upon relying on combinations of statistics such as in Eq. 2.

Furthermore, being able to provide a target in the formulation in Eq. 4 adds flexibility to the approach since we are able to specify more than just the basic robust optimization philosophy. This flexibility is a result of the metric in Eq. 3 integrating the L_2 norm over h , since it allows arbitrary functions of h to be used as targets. Other possible metrics that integrate the difference between the CDFs as functions of q (e.g., the Cramer-Von Mises test³¹) would restrict the targets to being inverses of valid CDFs. Additionally,

integrating the L_2 norm over h as opposed to q penalizes sections of a CDF further from the target more than it rewards sections closer to the target, giving the metric an intrinsic preference for robust designs and allowing the target to influence the optimization even when it gives non-feasible values of q ; this would not be the case if the metric was integrated over q . A few uses of the target for different design scenarios are highlighted here:

- *Standard target.* To follow the basic robust optimization philosophy of maximizing the likelihood of achieving as good performance as possible, the targets for both curves should be set to a single value: $t_u(h) = t_l(h) = q_{ideal}$, where q_{ideal} is a value of the quantity of interest beyond what is achievable (e.g., zero weight, 100% efficiency).
- *Risk-averse.* By modifying the shape of the standard target a designer can specify a preference for risk-averse designs so that robustness is preferred over possible high performance: skewing the shape of $t_l(h)$ and $t_u(h)$ for h close to 1 to lower values of q or shifting the whole of t_l to lower values of q emphasizes minimizing the worst cases of q over the horsetail plot. Figure 4 illustrates examples of risk averse targets created like this under probabilistic, interval and mixed uncertainties.

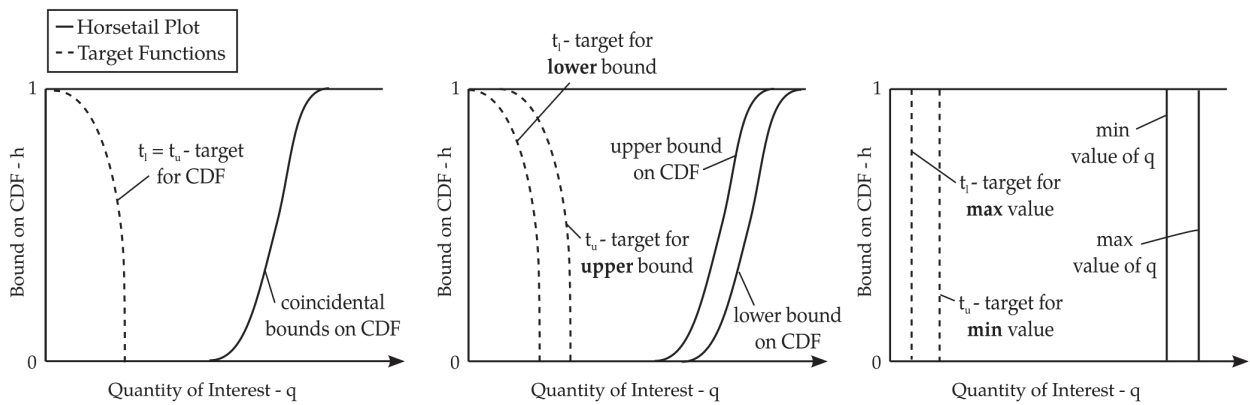


Figure 4: Risk-averse targets (dotted) along with typical horsetail plots (solid) under probabilistic (left), mixed uncertainties (center), and interval uncertainties (right).

- *Risk-seeking.* A designer can alternatively emphasize possible performance over robustness by modifying the standard target in the opposite sense to the risk averse targets as illustrated in Figure 5, instructing the optimization to favor more ambitious designs.




Figure 5: Risk-seeking targets (dotted) along with typical horsetail plots (solid) under probabilistic (left), mixed uncertainties (center), and interval uncertainties (right).

- *Feasible distribution.* In some applications, a designer might care more about higher moments of individual CDFs such as skewness than mean and variance. In this case target distributions over feasible ranges of q with desirable higher-order moment properties can be provided (this was the main advantage of the density matching approach²⁴).

- *Specific value.* In other applications, for example where a component of a larger system is being designed, pure minimization may not be what is required and instead a target value of q_{target} is desired: this is implemented in the horsetail matching formulation by setting $t_u(h) = t_l(h) = q_{\text{target}}$.
- *Worst case optimization.* If the risk averse target is taken to the extreme, such that $t_l(1) \rightarrow \infty$, then only the worst case value is optimized and the horsetail matching formulation reduces back to the robust regularization approach.^{1,21}

These possible uses of the target demonstrate the flexibility offered by horsetail matching in that by varying the target, the optimization under uncertainty problem can be formulated for a variety of scenarios under any mix of probabilistic and interval uncertainties. This is a powerful feature of the approach, and it also allows the formulation to recover existing methods of robust optimization by choosing appropriate targets.

II.C. Implementation

In order to evaluate d_{hm} in Eq. 3, integrals of the form:

$$D = \int_0^1 (F^{-1}(h) - t(h))^2 dh \quad (7)$$

need to be evaluated, where $F^{-1}(h)$ and $t(h)$ are respectively the inverse for either the upper or lower bound on the CDF (one of the two horsetail curves) and its target. The method used in this paper to numerically approximate this integral builds upon density matching²⁴ and previous work into horsetail matching under probabilistic uncertainties,³⁰ where kernels are used to find a differentiable approximation to the CDF. Here it is extended to mixed uncertainties by making use of differentiable approximations to the minimum and maximum functions.

The method of obtaining an approximation, \hat{D} , to integrals in the form of Eq. 7 at a given design is outlined in the following. Recall that we are interested in a quantity of interest, q , that is a function of probabilistic uncertainties, θ , and interval uncertainties, ϕ . Note that prior to a horsetail matching optimization, N fixed points q_i are selected for use in the numerical integration, M_θ samples θ_j are drawn from the distribution of θ , and M_ϕ samples ϕ_l are drawn from the hyper-rectangle defined by the intervals of ϕ . These sets of values $\{q_i: i = 1, \dots, N\}$, $\{\theta_j: j = 1, \dots, M_\theta\}$, and $\{\phi_l: l = 1, \dots, M_\phi\}$ are used in the following algorithms for every design point \mathbf{x} throughout the optimization.

First we obtain an expression for the CDF of q over the probabilistic uncertainties and its gradient with respect to the design variables at each value of ϕ_l using Algorithm 1, where K and Φ are kernel functions.

Algorithm 1 Evaluating a single CDF, f , and its gradient at q with ϕ fixed

- 1: **for** $j = 1, \dots, M_\theta$ **do**
 - 2: $q_j \leftarrow$ quantity of interest $q(\mathbf{x}, \theta_j, \phi)$ from simulation
 - 3: **for** $k = 1, \dots, n_x$ **do**
 - 4: $\frac{\partial q_j}{\partial x_k} \leftarrow$ gradient of quantity of interest $q(\mathbf{x}, \theta_j, \phi)$ from simulation.
 - 5: $f \leftarrow \frac{1}{M_\theta} \sum_{j=1}^{M_\theta} \Phi(q - q_j)$
 - 6: **for** $k = 1, \dots, n_x$ **do**
 - 7: $\frac{\partial f}{\partial x_k} \leftarrow \frac{1}{M_\theta} \sum_{j=1}^{M_\theta} K(q - q_j) (-1) \frac{\partial q_j}{\partial x_k}$
-

We use a Gaussian kernel such that Φ is the error function for a Gaussian and K is its derivative (which would be used as the kernel function itself in a Kernel density estimation of the PDF³²):

$$K(r) = \frac{d\Phi(r)}{dr} = \frac{1}{\sqrt{2\pi b^2}} \exp\left(-\frac{r^2}{2b^2}\right) \quad \Phi(r) = \frac{1}{2} \left(1 + \operatorname{erf}\left(\frac{r}{\sqrt{2b^2}}\right)\right) \quad (8)$$

where $r \in \mathcal{Q}$ is just the argument of the kernel functions, and b is the bandwidth parameter of the kernel functions, which is fixed throughout the optimization.

Next an expression for the bounds on these CDFs (the horsetail curves) over all samples ϕ_l and the gradients of these bounds are found using Algorithm 2.

Algorithm 2 Evaluating a horsetail curve, h , and its gradient at q

- 1: **for** $l = 1, \dots, M_\phi$ **do**
 - 2: $f_l \leftarrow$ value of CDF at q with $\phi = \phi_l$ from Algorithm 1.
 - 3: **for** $k = 1, \dots, n_x$ **do**
 - 4: $\frac{\partial f_l}{\partial x_k} \leftarrow$ gradient of CDF at q with $\phi = \phi_l$ from Algorithm 1.
 - 5: $h \leftarrow S_\alpha = \left(\sum_{l=1}^{M_\phi} f_l e^{\alpha f_l} \right) / \left(\sum_{l=1}^{M_\phi} e^{\alpha f_l} \right)$
 - 6: **for** $k = 1, \dots, n_x$ **do**
 - 7: $\frac{\partial h}{\partial x_k} \leftarrow \sum_{l=1}^{M_\phi} \frac{\partial S_\alpha}{\partial f_l} \frac{\partial f_l}{\partial x_k} = \sum_{l=1}^{M_\phi} \left(\frac{e^{\alpha f_l}}{\sum_{m=1}^{M_\phi} e^{\alpha f_m}} \left[1 + \alpha (f_l - S_\alpha) \right] \frac{\partial f_l}{\partial x_k} \right)$
-

In Algorithm 2, S_α is a differentiable approximation to the maximum or minimum of the M_ϕ values f_l , where α is a fixed parameter whose sign controls which extreme it approximates and whose magnitude controls the how good the approximation is. Large positive values of α make S_α an approximation to the maximum function, and large negative values make S_α an approximation to the minimum function: as $\alpha \rightarrow \infty$, $S_\alpha \rightarrow \max(\{f_l : l = 1, \dots, M_\phi\})$ and as $\alpha \rightarrow -\infty$, $S_\alpha \rightarrow \min(\{f_l : l = 1, \dots, M_\phi\})$.

The function S_α uses exponentials to heavily weight the most extreme values of f_l in a normalized sum of all the values; the approximation becomes better as higher values of α are used since the difference between the most extreme and next most extreme values in the sum become more exaggerated. In general a value of α should be chosen to give a good approximation, but if it is too large then the $f e^{\alpha f}$ terms can exceed the largest floating point numbers able to be stored in memory so some caution should be exercised. In this case since f can only ever take values in the interval $[0, 1]$, fairly large values of α can be used, and this is explored further in Section III.

Next, quadrature is employed to find an approximation, \hat{D} , of the integral D , as a weighted sum of $(F^{-1}(h) - t(h))^2$ at N values of h :

$$D \simeq \hat{D} = \sum_{i=1}^N (F^{-1}(h_i) - t(h_i))^2 w_i \quad (9)$$

where h_i are the quadrature points and w_i are the quadrature weights. Algorithm 3 outlines the method used in horsetail matching to obtain \hat{D} and its gradient; this is a trapezium rule integration where the quadrature points h_i depend on the function $F(q)$. Note that `zeros(N)` initializes an array of size N with zeros, and `zeros(N, N)` an $N \times N$ matrix.

The overall method of numerically evaluating \hat{D} is illustrated in Figure 6, where the figures from left to right outline Algorithms 1 to 3. Algorithm 3 is used to evaluate D_u and D_l , the approximation to the integral in Eq. 7 for each horsetail curves. Then finally the metric and its gradient are evaluated by $d_{hm} \simeq (\hat{D}_u + \hat{D}_l)^{1/2}$ and $\frac{\partial d_{hm}}{\partial x_k} \simeq 0.5 \left(\frac{\partial \hat{D}_u}{\partial x_k} + \frac{\partial \hat{D}_l}{\partial x_k} \right) (\hat{D}_u + \hat{D}_l)^{-1/2}$.




Figure 6: Evaluation of a single CDF at a fixed ϕ_l by summing kernel functions centered at M_θ samples q_j (left); differentiable approximation of the maximum of the CDFs h from M_ϕ values of f (center); trapezium rule integration of \hat{D} using N values of q_i , h_i and t_i (right).

Algorithm 3 Evaluating \hat{D} and its gradient

```
1:  $\mathbf{q} \leftarrow \text{zeros}(N)$ 
2:  $\mathbf{t} \leftarrow \text{zeros}(N)$ 
3:  $\mathbf{W} \leftarrow \text{zeros}(N, N)$ 
4: for  $i = 1, \dots, N$  do
5:    $\mathbf{q}_i \leftarrow \text{fixed point } q_i$ 
6:    $h_i \leftarrow \text{horsetail curve at } q_i \text{ from Algorithm 2}$ 
7:    $\mathbf{t}_i \leftarrow t(h_i)$  value of  $t$  at  $h_i$  from the specified target
8:    $i_1 \leftarrow \min(i + 1, N)$ 
9:    $i_2 \leftarrow \max(i - 1, 1)$ 
10:   $\mathbf{W}_{i,i} \leftarrow 0.5( h_{i_1} - h_{i_2} )$ 
11:  $\hat{D} \leftarrow (\mathbf{q} - \mathbf{t})^T \mathbf{W} (\mathbf{q} - \mathbf{t})$ 
12: for  $k = 1, \dots, n_x$  do
13:   $\frac{\partial \mathbf{t}}{\partial x_k} \leftarrow \text{zeros}(N)$ 
14:   $\frac{\partial \mathbf{W}}{\partial x_k} \leftarrow \text{zeros}(N, N)$ 
15:  for  $i = 1, \dots, N$  do
16:     $\frac{\partial h_i}{\partial x_k} \leftarrow \text{gradient of the horsetail curve at } q_i \text{ from Algorithm 2}$ 
17:     $t' \leftarrow \text{derivative of } t \text{ with respect to the argument, } \frac{\partial t(h)}{\partial h}, \text{ at } h_i \text{ from the specified target}$ 
18:     $\left( \frac{\partial \mathbf{t}}{\partial x_k} \right)_i \leftarrow t' \frac{\partial h_i}{\partial x_k}$ 
19:     $i_1 \leftarrow \min(i + 1, N)$ 
20:     $i_2 \leftarrow \max(i - 1, 1)$ 
21:     $\left( \frac{\partial \mathbf{W}}{\partial x_k} \right)_{i,i} \leftarrow 0.5 \left( \frac{\partial h_{i_1}}{\partial x_k} - \frac{\partial h_{i_2}}{\partial x_k} \right)$ 
22:   $\frac{\partial \hat{D}}{\partial x_k} \leftarrow 2(\mathbf{q} - \mathbf{t})^T \mathbf{W} \frac{\partial \mathbf{t}}{\partial x_k} + (\mathbf{q} - \mathbf{t})^T \frac{\partial \mathbf{W}}{\partial x_k} (\mathbf{q} - \mathbf{t})$ 
```

Under purely probabilistic uncertainties, since $h(q)$ is the CDF of $q(\mathbf{x}_d, \boldsymbol{\theta})$ and so $\hat{D}_1 = \hat{D}_2$, Algorithm 2 is skipped and the numerical integration in Algorithm 3 is applied directly to the CDF found using Algorithm 1; this is the method used for probabilistic horsetail matching.³⁰

Under purely interval uncertainties, the horsetail curves become step functions and so kernels are not used to propagate the CDFs and a numerical integration is not performed; S_α in Algorithm 2 is used directly to obtain q_{max} (using a large positive value of α) and q_{min} (using a large negative value of α) and their gradients from M_ϕ sampled values of $q_j = q(\mathbf{x}_d, \phi_j)$.

Being able to obtain the metric in a differentiable form means that solving a horsetail matching optimization problem using a gradient-based algorithm only requires the ability to obtain values of q and $\frac{\partial q}{\partial x_k}$ at given values of \mathbf{x} , $\boldsymbol{\theta}$ and ϕ . If the values of q and $\frac{\partial q}{\partial x_k}$ are readily available, horsetail matching can be implemented as a wrapper treating the available code as a black box. Additionally, in many design problems, $\frac{\partial q}{\partial x_k}$ can be obtained efficiently (e.g., via adjoints in CFD codes³³), and so being able to propagate this information through to the gradient of d_{hm} is important for keeping the computational cost of the optimization under uncertainty problem low.

It is worth noting that a high number of samples (both of ϕ and $\boldsymbol{\theta}$) are required to accurately propagate this metric and the gradient within an optimization, which may be infeasible to do via direct sampling if an expensive simulation is used to find q and $\frac{\partial q}{\partial x_k}$. Therefore surrogate models can first be fitted to q as well as to $\frac{\partial q}{\partial x_k}, k = 1, \dots, n_x$ as a function of $\boldsymbol{\theta}$ and ϕ for each design \mathbf{x}_d . These surrogate models can then be sampled as many times as required at negligible computational cost (assuming the computational expense of the simulation vastly outweighs evaluating the surrogate model); propagating mixed uncertainties by sampling surrogate models has been demonstrated to be effective.^{11,26} This also means that it is no more expensive to propagate d_{hm} under mixed uncertainties than probabilistic or interval uncertainties since fitting a surrogate model for q as a function of the uncertainties is the same cost regardless of how it is sampled afterwards.

The use of Kernels requires the selection of a fixed bandwidth, b , prior to a horsetail matching optimiza-


tion. This can be done, for example, by using Scott’s rule³² at the initial design, but it is worth noting that a poor choice of bandwidth can lead to a highly non-smooth gradient (if b is too small) or can give smooth but erroneous values of d_{hm} and its gradient (if b is too large).

III. Experiments on an Algebraic Test Problem

In order to assess the effectiveness of the numerical implementation of the approach, here horsetail matching under mixed uncertainties is tested on an algebraic test problem. The test problem in Eq. 10 is used (an algebraic problem allows many optimizations to be run at low computational cost), where q is a function of two design variables x_1 and x_2 that are each bounded by $[-5, 5]$, one probabilistic uncertainty, θ , which is uniformly distributed over the range $[-1, 1]$, and one interval uncertainty, ϕ , which is contained in the interval $[-1, 1]$.

$$\begin{aligned} y &= x_1/2 \\ z &= x_2/2 + 12 \\ q &= 0.25((y^2 + z^2)/40 + 5y\theta\phi - z\phi^2) + 0.2(z\phi^3) + 10 \end{aligned} \tag{10}$$

Figure 7 shows horsetail plots for this test problem at the design point $\mathbf{x} = (4, 2)$. The method in Section II.C is used to propagate the CDF from $M_\theta = 100$ samples of θ for each value of ϕ (plotted in grey) and then find the envelope of the $M_\phi = 30$ CDFs at $N = 1000$ integration points q_i evenly spaced between $q = -5$ and $q = 30$. To determine an appropriate value of α to use in S_α , Algorithm 2 is performed using different values of α .



(a) Mixed, $\alpha = 100$

(b) Mixed, $\alpha = 10$

Figure 7: Propagated horsetail plots of the test problem using $\alpha = 100$ and $\alpha = 10$.

It can be seen that the method for $\alpha = 100$ obtains the envelope of the CDFs to visible accuracy, whereas $\alpha = 10$ gives a poor approximation to the envelope. Choosing a higher value $\alpha = 1000$ resulted in an overflow error, so $\alpha = 100$ is used for the horsetail matching optimizations in this section and in Section IV. Furthermore, for this test problem, the mixed case for $\alpha = 0$ is equivalent to the purely probabilistic case where ϕ is uniformly distributed since the samples of ϕ are drawn uniformly from the interval $[-1, 1]$, and these two cases are compared in Figure 8.

Optimization Tests

A gradient-based optimizer (SLSQP, implemented using the NLopt toolbox^a) is run on the test problem from 50 random starting points in design space, where the standard target at $q = 0$ is used. The propagated gradient is found from the sensitivity of q to design variables, $\frac{\partial q}{\partial x_k}$, which is obtained from the analytic

^a<http://ab-initio.mit.edu/wiki/index.php/NLopt>




Figure 8: Propagated horsetail plot of the test problem using $\alpha = 0$ and propagated probabilistic CDF

derivative of Eq. 10 at the sample points. Convergence histories are plotted in Figure 9. Note that the difference between the HM metric and the best value obtained from all 50 optimizations is plotted, where a difference of 10^{-3} is considered converged.

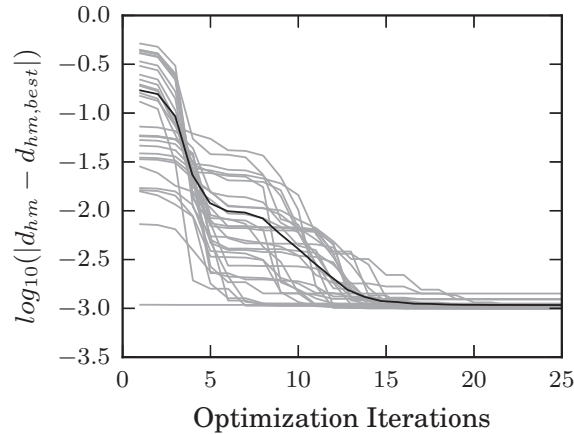


Figure 9: Horsetail matching optimization convergence histories under mixed uncertainties using the propagated gradient on the algebraic test problem.

It can be observed that the optimizer solves the horsetail matching problem from all 50 starting points, indicating the effectiveness of the numerical implementation outlined in section II.C.

IV. Application to Physical Aero-Structural Design Problem

In this section, horsetail matching is applied to a physical aero-structural design test problem. The design problem involves determining the geometry of a flying wing in order to minimize the fuel burn for a given mission subject to lift and failure constraints. It uses a low fidelity coupled aero-structural analysis code^a that utilizes a vortex lattice method aerodynamics model expanded from a modern lifting line theory, and a linear 6-DOF-per-element spatial beam structural model. The code is built on the OpenMDAO architecture³⁴ to efficiently analyze the coupled simulation and its gradient. An illustration of a flying wing analyzed using

^a<https://github.com/hwangjt/OpenAeroStruct>

this code is given in Figure 10.




Figure 10: Illustration of the OpenAeroStruct flying wing model^a.

For the design problem considered here, the wing mesh is built from a total of 11 evenly spaced spanwise points and 4 chordwise points. The vehicle design is parametrized by controlling the span, taper ratio, sweep, and angle attack of the wing, along with the thickness and twist at three spanwise control points for half of the span (the wing is symmetric), giving a total of $n_x = 10$ design variables. The overall design space, \mathcal{X} , for this problem is given in Table 1.

Notation	Design Variable	Lower Bound (x_k^L)	Upper Bound (x_k^U)	Units
x_1	Span	40	70	m
x_2	Taper Ratio	0	0.9	-
x_3	Sweep	10	40	deg
x_4	Nominal Angle of Attack	0	20	deg
x_5	Twist at Loc. 1	-5	20	deg
x_6	Twist at Loc. 2	-5	20	deg
x_7	Twist at Loc. 3	-5	20	deg
x_8	Thickness at Loc. 1	0.1	0.3	m
x_9	Thickness at Loc. 2	0.1	0.3	m
x_{10}	Thickness at Loc. 3	0.1	0.3	m

Table 1: Design space for the flying wing design problem.

We consider two probabilistic uncertainties, θ_1 and θ_2 , and one interval uncertainty, ϕ . The probabilistic uncertainties represent uncertainty in the operating conditions of the flying wing: the Mach number is uniformly distributed from 0.65 to 0.75, and the actual angle of attack is distributed uniformly over a 4 degree range centered on the nominal angle of attack which is denoted by γ_{nom} and given by design variable x_4 . The interval uncertainty represents uncertainty in the material properties: the failure stress is in the interval [15, 25] MPa. The uncertainties are detailed in Table 2.

Notation	Uncertain Parameter	Type	Range	Units
θ_1	Mach Number	Uniform Distribution	[0.65, 0.75]	m
θ_2	Actual Angle of Attack	Uniform Distribution	$[\gamma_{nom} - 2, \gamma_{nom} + 2]$	deg
ϕ	Failure Stress	Interval	[15, 25]	MPa

Table 2: Uncertainties for the flying wing design problem.

The objective is to minimize the fuel burn for a given mission, subject to the following two constraints: the lift should be greater than or equal to the total weight, and the maximum Von-mises stress in the wing should not be greater than the failure stress of the material. These constraints are taken into account in the optimizations using squared penalty functions so that, denoting the fuel burn as m_{fuel} (in units of $kg \times 10^6$),

the lift constraint (normalized by the initial weight) as g_1 , and the failure constraint (normalized by the failure stress) as g_2 , the total quantity of interest for this design problem becomes:

$$q(\mathbf{x}, \theta_1, \theta_2, \phi) = m_{fuel} + 3(\max(g_1, 0))^2 + 3(\max(g_2, 0))^2 \quad (11)$$

The remaining parameters used to set up the OpenAeroStruct analysis are given in Table 3.

OpenAeroStruct Notation	Parameter	Value	Units
Alt	Altitude	30000	feet
E	Elastic Modulus	200×10^9	Pa
G	Shear Modulus	30×10^9	Pa
mrho	Material Density	3×10^3	kg/m ³
W0	Base Weight	1.00×10^5	kg
CL0	Base Lift Coefficient	0.25	-
CD0	Base Drag Coefficient	0.015	-
SFC	Specific Fuel Consumption	17×10^6	1/s
chord	Chord at Wing Root	12	m
R	Range	10×10^6	m

Table 3: Parameters used in the OpenAeroStruct analysis of the flying wing.

Optimization Results

We compare designs resulting from different types of optimizations on the design problem to compare horsetail matching to alternative approaches. For all these optimizations, the gradient-based SLSQP algorithm is used, implemented with NLOpt^a, and is run for 20 iterations (representing a fixed computational budget) after which the best design is taken as the optimum. Figures 11 - 15 give a plan view of the mesh and the horsetail plot of the quantity of interest q from Equation 11 for the optimum design resulting from each optimization case. The mesh vertices are colored to show the twist at each spanwise location: from red corresponding to -5° to blue corresponding to $+15^\circ$. The values of the design variables for the optimum design in each case are given in Table 4.

First we run a deterministic optimization from an initial design given by $s_k = 0.5, k = 1, 2, \dots, 10$, where the nominal value of the quantity of interest (with $\theta_1 = \theta_2 = \phi = 0$) is minimized:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && q(\mathbf{x}, \theta_1 = 0, \theta_2 = 0, \phi = 0) \\ & \text{s.t.} && x_k^L < x_k < x_k^U \quad k = 1, \dots, 10 \end{aligned} \quad (12)$$

The resulting optimum design along with the horsetail plot of q for this design is given in Figure 11.

Now we run horsetail matching optimizations to minimize the horsetail matching metric d_{hm} from Equation 3, where the value and the gradient are evaluated using the approach outline in Section II.C, given the specified target functions $t_u(h)$ and $t_l(h)$:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && d_{hm}(q(\mathbf{x}, \theta_1, \theta_2, \phi); t_u, t_l) \\ & \text{s.t.} && x_k^L < x_k < x_k^U \quad k = 1, \dots, 10. \end{aligned} \quad (13)$$

For this design problem, a 3 dimensional, 3rd order polynomial surrogate is fitted to the quantity of interest, q , and each component of its gradient, $\frac{\partial q}{\partial x_k}$, as a function of the uncertainties, θ_1 , θ_2 and ϕ , at each design point. This corresponds to $4^3 = 64$ calls of the OpenAeroStruct analysis code at each iteration, giving a total computational budget of $64 \times 20 = 1280$ calls for each optimization.

First we run a horsetail matching optimization using a standard target by setting $t_l(h) = t_u(h) = 0$; the resulting optimum design and horsetail plot of q for this design are given in Figure 12.

^a<http://ab-initio.mit.edu/wiki/index.php/NLOpt>




Figure 11: Wing mesh and horsetail plot for the optimum design of a deterministic optimization




Figure 12: Wing mesh and horsetail plot for the optimum design of a horsetail matching optimization using the standard target.

Considering the horsetail plot in Figures 11 we can see that since the deterministic optimization did not take into account the uncertainties, the bounds on the CDF both have tails at large values of the quantity of interest (which represents fuel burn plus penalties for constraint violation), indicating a high likelihood that this design will give degraded performance in reality. In contrast, the horsetail plot for the horsetail matching optimum design (under the standard target) in Figure 12 gives bounds on the CDFs whose tails do not reach as large values of q , indicating improved robustness under mixed uncertainties, whilst only taking a small penalty to the best possible value.

Comparing the meshes for these two designs in Figures 11 and 12 and the design variables themselves in Table 1, we can see that, when compared to the deterministic optimum, the horsetail matching optimum has a smaller span, but a larger taper ratio and nominal angle of attack, along with a less aggressive variation in twist along the wing. These design changes drive the horsetail matching optimum further away from the constraint boundaries, giving a slightly worse nominal performance but reducing the penalty to q due to violation of the constraints under variation due to the uncertainties.

We also run horsetail matching optimizations under different targets. Firstly a risk averse target is used, obtained by skewing the top of the target for the lower bound of the CDF (the right-most horsetail curve) to a value of -10 from the base standard target at $q = 0$; the optimum design and its corresponding horsetail plot are given in Figure 13. Secondly a target at a specific value of $q = 1.3$ is given to minimize variation around this value; the optimum design and its corresponding horsetail plot are given in Figure 14.




Figure 13: Wing mesh and horsetail plot for the final design of a horsetail matching optimization using a risk averse target.

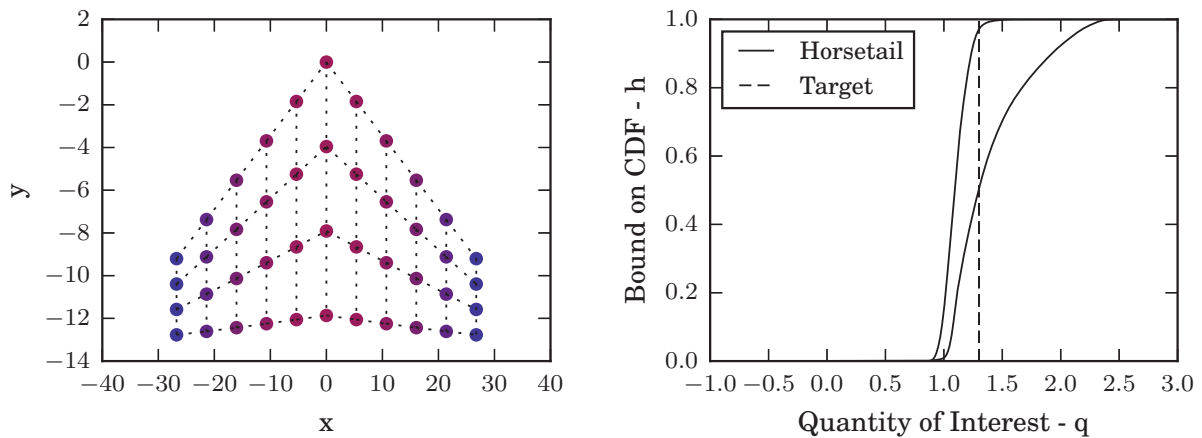


Figure 14: Wing mesh and horsetail plot for the final design of a horsetail matching optimization using a target at a specific value of $q = 350$.

From Figure 13 and Table 1, we can see that the risk averse design is similar to the standard horsetail matching optimum design, but it has a slightly better worse-case performance (the highest value of q reached by the lower bound on the CDF). From Figure 14 and Table 1, we can see that the single value target design is significantly different to the other optimum designs, since the optimizer is instead looking to minimize the variation of both bounds around the target value of $q = 1.3$.

These two cases clearly demonstrate the flexibility of the horsetail matching formulation, as designs with horsetail plots with different properties, representing optimization under different design scenarios, were obtained simply by changing the target used in the optimization.

Comparison to the Weighted Sum Approach

Finally we compare horsetail matching to the method commonly suggested in the literature for optimization under mixed uncertainties: using a weighted sum of averages and intervals of statistical moments of the CDFs that make up the horsetail plot.²⁵⁻²⁸ The following weighted sum of three objectives is optimized:

$$f_{obj} = w_1\bar{\mu} + w_2\bar{\sigma} + w_3\delta\mu \quad (14)$$

where $\bar{\mu}$ is the average of μ_{min} & μ_{max} and $\delta\mu$ is the difference between μ_{min} & μ_{max} , referring to the minimum or maximum over all possible values of the interval uncertainties ϕ (with the same meaning for the

standard deviation σ). To propagate these three objectives, the method outlined in Ref. 26 is used, where polynomial chaos expansions of q as a function of the uncertainties are used to find expressions for μ and σ as functions of the interval uncertainties ϕ , which are then optimized over ϕ to obtain μ_{min} , μ_{max} , σ_{min} , and σ_{max} for use in Eq. 14. Since this method uses a surrogate model at each design point, each optimization iteration has the same computational cost as the horsetail matching implementation; 3rd order polynomials in each dimension are used in both cases.

The weights are selected using the method suggested in Refs. 28 and 26: by looking at the relative magnitudes of the three objectives at the deterministic optimum design and selecting weights so that they all have the same contribution to f_{obj} for this design. For this design problem, at the deterministic optimum the three objectives are $\bar{\mu} = 1.00$, $\bar{\sigma} = 0.392$, and $\delta\mu = 0.0745$, so the weights are selected as $w_1 = 1.00$, $w_2 = 2.55$, and $w_3 = 13.43$. Thus the following problem is optimized:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \bar{\mu} + 2.56 \bar{\sigma} + 13.43 \delta\mu && (15) \\ & \text{s.t.} && x_k^L < x_k < x_k^U && k = 1, \dots, 10 \end{aligned}$$

The results of this weighted sum optimization are given in Figure 15, where the horsetail plot (Weighted Sum) is compared to the plot for the horsetail matching optimum using the standard target (labeled Standard HM).

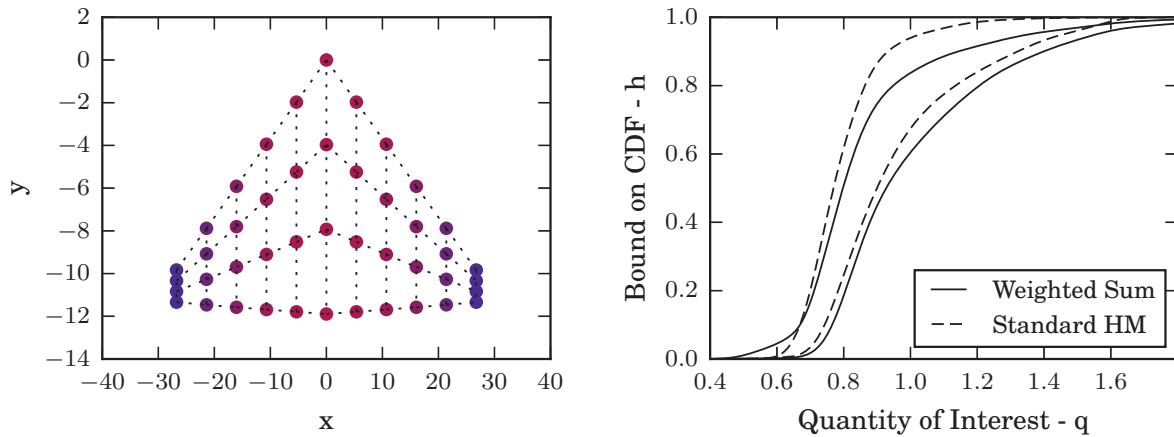


Figure 15: Wing mesh and horsetail plot for the final design of a weighted sum optimization, along with the horsetail plot from the standard horsetail matching optimum design.

It can be observed from Figure 15 that the optimum design using these weights is close to being stochastically dominated by the standard horsetail matching optimum design: the lower bound on the CDF lies entirely to the right of the lower bound for the standard HM optimum and the upper bound is only better for $h < 0.1$. This highlights the limitations that were discussed in Section I of using a weighted sum of combinations of statistical moments as the objective to be minimized in an optimization as in Equation 15.

Table 4 gives the design variables for the optimum designs from the deterministic optimization, the horsetail matching optimizations under the standard target ($t_l(h) = t_u(h) = 0$), the risk-averse target, and the target at a specific feasible value ($t_l(h) = t_u(h) = 1.3$) as well as the weighted sum optimization.

	Deterministic	Standard	Risk Averse	Specific Value	Weighted Sum
Span	55	53.63	54.14	53.5	52.6
Taper Ratio	0.1	0.118	0.119	0.305	0.118
Sweep	10	10	10	10.9	10.0
Nominal Angle of Attack	9.5	10.2	10.19	10.4	10.2
Twist at Loc. 1	9.20	9.14	8.84	10.0	9.13
Twist at Loc. 2	2.0	2.47	2.72	3.56	2.47
Twist at Loc. 3	-1.0	0.14	-0.21	0.76	0.15
Thickness at Loc. 1	0.1	0.1	0.1	0.13	0.1
Thickness at Loc. 2	0.1	0.1	0.1	0.13	0.1
Thickness at Loc. 3	0.1	0.1	0.1	0.12	0.1

Table 4: Design parameters for the optimum designs from the different optimization cases

V. Conclusion

Horsetail matching is a flexible approach to optimization under any mix of probabilistic and interval uncertainties, where the difference between a design’s horsetail plot and a target is minimized. The difference is measured by the total L_2 norm integrated over h over both horsetail curves, which represent the upper and lower bound on the CDF. As well as giving the option to optimize towards a feasible target, since the target is not required to consist of the inverses of valid CDFs, a designer can also use the target to specify their preferences for how risk averse/seeking the optimization should be. Further, since the metric intrinsically favors robust solutions a standard target at a value of q beyond feasible is proposed, for which the optimal solution will not be *stochastically dominated by* other possible designs. This is not a feature of the most commonly proposed approach for robust optimization under mixed uncertainties of minimizing a weighted sum of combinations of statistical moments.

This paper outlines a numerical implementation that delivers a single objective, differentiable metric that is shown to accurately capture the envelope of the possible CDFs (the horsetail plot) and propagate both the metric and its gradient under mixed uncertainties on an algebraic test problem. The approach is applied to the design of a flying wing analyzed using a low fidelity, coupled aero-structural analysis code; on this physical problem the flexibility of horsetail matching is demonstrated by using different targets to optimize for different design scenarios at low computational cost. It was also compared to the weighted sum of combinations of statistical moments approach and shown to give preferable designs at similar computational cost. Although in this paper the approach has been demonstrated on a single problem, since the only requirement for using the formulation is the ability to obtain a quantity of interest and its gradient it can be applied to any such design problem.

Planned future work into horsetail matching primarily involves investigating and improving the formulation when a high number of uncertainties are involved in the problem, since the effectiveness of using surrogate models to propagate the uncertainty deteriorates rapidly as the dimensionality of the uncertainty space increases; this is considered a limitation of the current approach.

Acknowledgments

This work is funded by the Engineering and Physical Sciences Research Council (EPSRC) UK. The third author acknowledges support of the Air Force Office of Scientific Research (AFOSR) MURI on managing multiple information sources of multi-physics systems, Program Manager Jean-Luc Cambier, Award Number FA9550-15-1-0038.

References

- ¹Beyer, H. G. and Sendhoff, B., "Robust Optimization - A Comprehensive Survey," *Computer Methods in Applied Mechanics and Engineering*, Vol. 196, No. 33-34, 2007, pp. 3190–3218.
- ²Kennedy, M. C. and O'Hagan, A., "Bayesian Calibration of Computer Models," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, Vol. 63, No. 3, 2001, pp. 425–464.
- ³Du, X. and Chen, W., "Methodology for Managing the Effect of Uncertainty in Simulation-Based Design," *AIAA Journal*, Vol. 38, No. 8, 2000.
- ⁴Keane, A. J. and Nair, P. B., *Computational Approaches for Aerospace Design: The Pursuit of Excellence*, Wiley, New York, 2005.
- ⁵Ben-Tal, A. and Nemirovski, A., "Robust Convex Optimization," *Mathematics of Operations Research*, Vol. 23, No. 4, 1998, pp. 769–805.
- ⁶Bertsimas, D. and Sim, M., "The Price of Robustness," *Operations Research*, Vol. 52, No. 1, 2004, pp. 35–53.
- ⁷Shapiro, A., Dentcheva, D., and Ruszczyński, A., *Lectures on Stochastic Programming: Modeling and Theory*, SIAM, 2009.
- ⁸Helton, J. C., Johnson, J. D., and Oberkampf, W. L., "An Exploration of Alternative Approaches to the Representation of Uncertainty in Model Predictions," *Reliability Engineering and System Safety*, Vol. 85, No. 1-3, 2004, pp. 39–71.
- ⁹Rao, S. S. and Cao, L., "Optimum Design of Mechanical Systems Involving Interval Parameters," *Journal of Mechanical Design*, Vol. 124, No. 3, 2002.
- ¹⁰Eldred, M. S., Swiler, L. P., and Tang, G., "Mixed Aleatory-Epistemic Uncertainty Quantification With Stochastic Expansions and Optimization-Based Interval Estimation," *Reliability Engineering and System Safety*, Vol. 96, No. 9, 2011, pp. 1092–1113.
- ¹¹Hosder, S. and Bettis, B., "Uncertainty and Sensitivity Analysis for Reentry Flows with Inherent and Model-Form Uncertainties," *Journal of Spacecraft and Rockets*, Vol. 49, No. 2, 2012, pp. 193–206.
- ¹²Shafer, G., *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- ¹³Zadeh, L., Jagoita, C., and Zimmermann, H., "Fuzzy Sets as a Basis for a Theory of Possibility," *Fuzzy Sets and Systems*, Vol. 1.3, No. 28, 1978.
- ¹⁴Dodson, M. and Parks, G. T., "Robust Aerodynamic Design Optimization Using Polynomial Chaos," *Journal of Aircraft*, Vol. 46, No. 2, 2009, pp. 635–646.
- ¹⁵Keane, A. J., "Comparison of Several Optimization Strategies for Robust Turbine Blade Design," *Journal of Propulsion and Power*, Vol. 25, No. 5, 2009, pp. 1092–1099.
- ¹⁶Ghisu, T., Jarrett, J. P., and Parks, G. T., "Robust Design Optimization of Airfoils with Respect to Ice Accretion," *Journal of Aircraft*, Vol. 48, No. 1, 2011, pp. 287–304.
- ¹⁷Park, G.-J., Lee, T.-H., Lee, K. H., and Hwang, K.-H., "Robust Design: An Overview," *AIAA Journal*, Vol. 44, No. 1, 2006, pp. 181–191.
- ¹⁸Padulo, M., Campobasso, M. S., and Guenov, M. D., "Novel Uncertainty Propagation Method for Robust Aerodynamic Design," *AIAA Journal*, Vol. 49, No. 3, 2011, pp. 530–543.
- ¹⁹Zaman, K., McDonald, M., Mahadevan, S., and Green, L., "Robustness-Based Design Optimization under Data Uncertainty," *Structural and Multidisciplinary Optimization*, Vol. 44, No. 2, 2011, pp. 183–197.
- ²⁰Lee, S. W. and Kwon, O. J., "Robust Airfoil Shape Optimization Using Design for Six Sigma," *Journal of Aircraft*, Vol. 43, No. 3, 2006, pp. 4–7.
- ²¹Ryan, K. M., Lewis, M. J., and Yu, K. H., "Comparison of Robust Optimization Methods Applied to Hypersonic Vehicle Design," *Journal of Aircraft*, Vol. 52, No. 5, 2015, pp. 1510–1523.
- ²²Ben-Tal, A., Ahaoui, L. E., and Nemirovski, A., *Robust Optimization*, Princeton University Press, 2009.
- ²³Petrone, G., Iaccarino, G., and Quagliarella, D., "Robustness Criteria In Optimization Under Uncertainty," *Evolutionary and Deterministic Methods for Design, Optimization and Control*, 2011.
- ²⁴Seshadri, P., Constantine, P., Iaccarino, G., and Parks, G., "A Density-Matching Approach for Optimization Under Uncertainty," *Computer Methods in Applied Mechanics and Engineering*, 2016.
- ²⁵Rumpfkeil, M. P., "Robust Design Under Mixed Aleatory / Epistemic Uncertainties Using Gradients and Surrogates," *Journal of Uncertainty Analysis and Applications*, 2013.
- ²⁶Zhang, Y. and Hosder, S., "Robust Design Optimization Under Mixed Uncertainties With Stochastic Expansions," *Journal of Mechanical Design*, Vol. 135, 2013.
- ²⁷Shah, H., Hosder, S., Koziel, S., Tesfahunegn, Y., and Leifsson, L., "Multi-fidelity Robust Aerodynamic Design Optimization Under Mixed Uncertainty," *Aerospace Science and Technology*, Vol. 45, 2015, pp. 17–29.
- ²⁸Du, X., Venigella, P. K., and Liu, D., "Robust Mechanism Synthesis with Random and Interval Variables," *Mechanism and Machine Theory*, Vol. 44, No. 7, 2009, pp. 1321–1337.
- ²⁹Levy, H., *Stochastic Dominance, Investment Decision Making Under Uncertainty*, Springer, 2015.
- ³⁰Cook, L. W. and Jarrett, J. P., "Horsetail Matching: A Flexible Approach to Optimization Under Uncertainty," (*Submitted for Publication*) *Engineering Optimization*.
- ³¹Anderson, T. W., "On the Distribution of the Two-Sample Cramér-von Mises Criterion," *The Annals of Mathematical Statistics*, Vol. 33, No. 3, 1962, pp. 1148–1159.
- ³²Scott, D. W., *Multivariate Density Estimation: Theory, Practice, and Visualization*, John Wiley & Sons, Inc., 1992.
- ³³Jameson, A., Martinelli, L., and Pierce, N. A., "Optimum Aerodynamic Design Using the Navier-Stokes Equations," *Theoretical and Computational Fluid Dynamics*, Vol. 10, No. 1-4, 1998, pp. 213–237.
- ³⁴Gray, J., Moore, K., and Naylor, B., "OpenMDAO: An Open Source Framework for Multidisciplinary Analysis and Optimization," *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, 2010.