Using Stochastic Dominance in Multi-Objective Optimizers for Aerospace Design Under Uncertainty

Laurence W. Cook* and Jerome P. Jarrett†

University of Cambridge, Cambridge, England, CB21PZ.

In optimization under uncertainty for aerospace design, statistical moments of the quantity of interest are often treated as separate objectives and are traded off in a multi-objective optimization formulation. However, in many design problems the trade-off between statistical moments can be large and the Pareto front representing this trade-off can include designs with undesirable behavior, such as being robust but being guaranteed to give a worse performance than another design. When a simulation of a system is computationally expensive, obtaining the full Pareto front is unfeasible and so spending optimization time obtaining such undesirable designs wastes time that could be spent obtaining more desirable alternatives. As a remedy, we propose an optimization formulation that can use multiple dominance criteria to avoid generating potentially inferior designs. We consider various orders of stochastic dominance as criteria to use alongside statistical moment based Pareto dominance, and illustrate how this gives rise to improved designs using a limited computational budget in an acoustic horn design problem and a transonic airfoil design problem.

I. Introduction

Optimization is now widely used as a design tool for engineering systems.¹ Its use is often not intended as a means of finding a truly “optimal design” in any sense, rather it is used as a guided method of searching design space in order to improve an initial design(s) given a computational budget. Since simulations of aerospace systems are often time-consuming, this computational budget often becomes the termination criterion for design optimizations instead of a traditional measure of convergence.

As a consequence of the uptake of optimization methods, the importance of considering the influence of uncertainties in the computational design of aerospace systems is becoming increasingly recognized,^{1,2} and optimization under uncertainty (OUU) has emerged as the framework through which this is achieved.³ If performance of the system is described by a quantity of interest, q , such as efficiency, cost, weight, drag, etc., then in OUU metrics that describe the behaviour of q under uncertainty are defined and optimized. Often, statistical moments are used as these metrics: the mean and variance (or standard deviation) of q are treated as separate objectives to be simultaneously minimized. This multi-objective framework seeks a Pareto front (a set of designs for which an objective cannot be improved without worsening another) trading off these objectives, from which the designer selects their preferred design(s) *a posteriori*. This approach has been used extensively in the aerospace literature, for example in the design of airfoils under icing conditions,⁴ windmill airfoil design,⁵ turbine blade design,⁶ subsonic airfoil design,⁷ transonic airfoil design,^{8,9} and the design of trailing edge flaps.¹⁰

However, when the quantity of interest is the output of a non-linear simulation of a physical system, its distribution is rarely Gaussian and so cannot be fully described by two statistical moments. Therefore treating the first two statistical moments of q as separate objectives can give rise to designs with potentially undesirable behavior. Consider, for example, the hypothetical designs described by their statistical moments and cumulative distribution functions (CDF) of the quantity of interest q given in Figure 1, where q is to be minimized. Whilst designs A, B and D all lie on the Pareto front of moments, design D is guaranteed to give

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

†Lecturer, Department of Engineering. AIAA Senior Member.

a worse q than either A or B, and so in many design scenarios will be inferior. Additionally, the CDF of A is not crossed by any of the other CDFs, indicating that for any given value of q , design A is more likely to achieve a performance at least as good as this value; this is known as (first order) stochastic dominance. An optimizer obtaining the Pareto front of moments will consider design D to be equally as optimal as design A and so will spend time seeking designs similar to design D, which under a limited computational budget may be better spent seeking out designs similar to design A.

Figure 1: Hypothetical designs. Designs A, B, and D lie on the Pareto front of moments. Designs A and B Pareto dominate design C. Design D is guaranteed to give a worse value of q than designs A and B. Design A first order stochastically dominates the rest.

The limitations of relying on statistical moments has been recognized in fields such as operations research and finance,¹¹ and the field of “stochastic programming”¹² has given rise to formulations for optimization under uncertainty using coherent risk measures¹³ as alternatives to statistical moments. Similarly, methods for optimization under a constraint that the solution stochastically dominates a reference solution have been developed,^{14,15} for example optimizing the expected profit of a portfolio subject to it stochastically dominating an existing investment.¹⁶ With these formulations it is possible to exploit underlying structure of the quantity of interest such as convexity and maintain mathematically favorable optimization problems. Additionally, alternative approaches to OUU for engineering design based on optimizing the distribution directly have been recently proposed.^{17–21}

Engineering design considers a broad range of problems that often lack the mathematical structure exploited in operations research (e.g., engineering design problems are typically non-convex, and in many cases the underlying analysis codes may be black-box). Therefore the optimizers of choice for the multi-objective framework for OUU in engineering design have typically been heuristics such as genetic algorithms and Tabu search. Such heuristics do not assume any structure of the underlying problem and do not necessarily guarantee any measure of convergence, but they can make significant design improvements using the given computational budget.^{4–8,10}

This paper considers an approach that seeks to obtain a superior set of designs compared to the traditional Pareto front of moments given a fixed computational budget. This is achieved by enforcing alternative (or additional) non-dominance criteria to filter out undesirable designs from an optimization. The primary non-dominance criterion investigated is stochastic dominance, but we proceed in a general sense so that any combination of general non-dominance criteria can be used subject to certain restrictions.

Multi-objective optimization heuristics are well-suited to exploring design space in order to obtain a set of non-dominated designs, since they already compare designs via a non-dominance criteria (defined by only Pareto dominance). In addition, such heuristics can obtain a set of designs (as opposed to just a single optimal design) in a single run. Therefore in this work two common heuristics are modified to obtain non-dominated sets based on alternative criteria: a multi-objective genetic algorithm (NSGA-II)²² and multi-objective Tabu search (MOTS).²³

Section II formally defines alternative non-dominance criteria such as stochastic dominance, discusses their relevance to optimization under uncertainty in aerospace design, and describes how multi-objective algorithms can be modified to obtain a general non-dominated set of designs. Section III performs detailed investigations of these modified optimizers on algebraic test problems, then Section IV and V applies these optimizers to practical design problems. Finally Section VI concludes the paper.

II. Optimizing for Alternative Dominance Criteria

A. Problem Description

In this work, we consider an engineering system whose performance is described by a quantity of interest that is to be minimized. This quantity of interest, q , is a function of controllable design variables $\mathbf{x}$ and uncertain inputs $\mathbf{u}$: $q = q(\mathbf{x}, \mathbf{u}) \in \mathbb{Q}$, where $\mathbb{Q}$ is the set of feasible values of q . The vector $\mathbf{u}$ is a realization of a random vector $\mathbf{U}(\omega)$, where ω is the underlying random event with sample space Ω such that $\omega \in \Omega$. Thus q for a given $\mathbf{x}$ is itself a random variable, denoted $q_x(\omega) := q(\mathbf{x}, \mathbf{U}(\omega))$.

The mean and variance of q_x are defined in the usual sense and denoted μ_x and σ_x^2 respectively. The CDF of q_x is denoted by $F_x(q) : \mathbb{Q} \rightarrow [0, 1]$. The quantile function is denoted by $F_x^{-1}(h) : [0, 1] \rightarrow \mathbb{Q}$, and is defined by the generalized right-continuous function $F_x^{-1}(h) := \min\{q' \mid F_x(q') \geq h\}$.

We also denote the superquantile by $\bar{Q}_x(h) : [0, 1] \rightarrow \mathbb{Q}$, which is defined as $\bar{Q}_x(h) := \mathbb{E}[q_x \mid q_x \geq F_x^{-1}(h)]$, where $\mathbb{E}$ indicates the expectation. The superquantile gives the expected value of q_x conditional on it being greater than the quantile ($F_x^{-1}(h)$). Therefore the superquantile includes information about the entire tail of the distribution instead of just a single quantile value, and is thus more useful when considering the worst outcomes of distributions and to designers with risk-averse preferences. The quantile function is often known as the value at risk (VaR) and the superquantile as the conditional value at risk (CVaR).²⁴

B. Optimization Formulation

We propose an optimization formulation that searches design space in order to find a set of designs $\mathbb{S}$ that satisfy general non-domination criteria making them of interest to the designer, out of the set of all possible designs $\mathbb{Y}$. Therefore it is no longer straightforward to express the formulation of the optimization as a minimization problem, since an objective function(s) is no longer being strictly minimized. This mirrors the paradigm in which optimization is often used for engineering design: we do not wish to find (a) truly optimal design(s) in any sense, rather optimization is used as a way of systematically and efficiently exploring design space.

Defining a list of n dominance criteria, such that $\mathbf{x} \prec_k \mathbf{y}$ indicates $\mathbf{x}$ dominates $\mathbf{y}$ subject to criterion k , we are looking to obtain the intersection of the non-dominated sets defined by each individual criterion, given by Equation 1:

$$\mathbb{S} := \bigcap_{k=1}^n \{\mathbf{y} \in \mathbb{Y} \mid \text{there does not exist } \mathbf{x} \in \mathbb{Y} \text{ for which } \mathbf{x} \prec_k \mathbf{y}\}. \quad (1)$$

In this work we will consider at most $n = 2$, but this formulation is provided for the general case. With such a formulation there are various criteria a designer might wish to use.

C. Examples of Dominance Criteria

Statistical moments are a convenient metric to use in a OUU formulation because often they are easily attainable. For example, when propagating uncertainties using polynomial chaos,^{25,26} statistical moments are available analytically. Additionally, the central limit theorem provides rigorous analytical estimates of the error in sample based estimates of these moments.²⁷

If the distribution of q were Gaussian, then these two moments would be sufficient to fully describe the distribution, and the Pareto front of μ_x and σ_x would contain all designs of interest. For practical design problems this is rarely (if ever) the case, not only because the feasible set of q is not the same as the support of a Gaussian, i.e. $\mathbb{Q} \subset [-\infty, \infty]$, but particularly because the system models are non-linear. Therefore there are other criteria by which to compare designs; we use the following definitions of some relevant criteria:

- Design $\mathbf{x}$ *Pareto* dominates design $\mathbf{y}$ if

$$\mu_x < \mu_y \text{ AND } \sigma_x < \sigma_y, \quad (2)$$

indicating that design $\mathbf{x}$ has a better expected performance than $\mathbf{y}$ and is more robust. This gives the traditional “robust optimization” approach often suggested for aerospace design.

- Design $\mathbf{x}$ *zeroth order stochastically dominates* (ZSD) design $\mathbf{y}$ if

$$F_x^{-1}(1) < F_y^{-1}(0), \quad (3)$$

indicating that the worst possible outcome for design $\mathbf{x}$ is better than the best possible outcome for design $\mathbf{y}$.

- Design $\mathbf{x}$ *quantile* dominates design $\mathbf{y}$, given n_h values of interest $h_k \in (0, 1)$, if

$$\forall k \in \{1, \dots, n_h\}, F_x^{-1}(h_k) < F_y^{-1}(h_k). \quad (4)$$

This formulation was suggested in Ref.,²¹ where a multi-objective formulation trading off the quantile function at several values of interest was proposed. However, given a set of evenly spaced values h_k , in the limit of $n_h \rightarrow \infty$ this becomes the first order stochastic dominance criterion.

- Design $\mathbf{x}$ *first order stochastically dominates* (FSD) design $\mathbf{y}$ if

$$\forall h \in (0, 1), F_x^{-1}(h) < F_y^{-1}(h), \quad (5)$$

indicating that design $\mathbf{x}$ is more likely to achieve all smaller values than any given value of q than design $\mathbf{y}$.

- Design $\mathbf{x}$ *superquantile* dominates design $\mathbf{y}$, given n_h values of interest $h_k \in (0, 1)$, if

$$\forall k \in \{1, \dots, n_h\}, \bar{Q}_x(h_k) < \bar{Q}_y(h_k) \quad (6)$$

This is simply an extension of treating quantiles as separate objectives to treating superquantiles as separate objectives. However, in a similar fashion to the quantile case, as $n_h \rightarrow \infty$ we recover the second order stochastic dominance criterion.

- Design $\mathbf{x}$ *second order stochastically dominates* (SSD) design $\mathbf{y}$ if

$$\forall h \in (0, 1), \bar{Q}_x(h) < \bar{Q}_y(h), \quad (7)$$

indicating that for any given value of q' , the expected value of q_x conditional on it being worse than q' is better than that for q_y .

While quantile and superquantile dominance are also Pareto dominance criteria under different objectives, we use the term Pareto dominance to refer specifically to Pareto dominance using the first two statistical moments as objectives, since this is the approach that is commonly used and that we are improving upon. This list is not exhaustive, but outlines the orders of stochastic dominance that we are suggesting are useful alternatives to the traditional Pareto dominance criterion. Still, the appropriate criteria to use depend on the scenario for which the system is being designed as well as the designer’s preferences.

D. Selecting Dominance Criteria

The traditional Pareto dominance of statistical moments criterion has been used because in general a goal of a design optimization is to improved the expected performance of a system. However in many cases this can give highly non-robust designs, so considering variance (or standard deviation) is an acknowledgment that robustness is a desirable trait, and by using both moments as objectives attempts to obtain designs that have a good expected performance but which are also robust. Unfortunately if there is a large trade-off between μ_x and σ_x , then the highly robust designs can reduce the expected performance too much to be of real interest in many scenarios.

Often in engineering design, many uncertainties are only realized a single time prior to the system being operated. For example, uncertainties representing a disparity between simulated predictions and physical reality, or uncertainties in the performance of future technology used in a design. Under such uncertainties, FSD is relevant, since once a value of q_x has been realized it is subsequently fixed, and so a design that is more likely to achieve performance at least as good as any given value of q than another design can be considered to be superior in such scenarios. Using FSD as a criterion filters out designs with poor performance over their entire CDF, as illustrated by the hypothetical CDFs on Figure 2a. However only using FSD as a criterion can give designs with a small probability of good performance but which are highly non-robust, and so we propose using FSD in addition to Pareto dominance of statistical moments.

If a designer is primarily interested in obtaining good performance, but is averse to a risk of obtaining poor performance, SSD is a relevant criterion. Using SSD gives the set of designs that any designer with risk-averse preferences would choose, by filtering out non-robust designs with poor tail performance, as illustrated by the hypothetical CDFs and superquantiles on Figure 2b.

Figure 2: Using FSD to filter out designs with poor performance over the entire CDF, and using SSD to filter out designs with poor performance in the tail of the CDF. In both cases design B (red) is dominated by design A (blue).

In some scenarios a designer may in fact be interested in robustness over good performance, in which case variance as an objective is appropriate. For example, the system may be a component of a larger multi-disciplinary product, and so robustness (implying the performance is predictable) may be suitable in a multi-disciplinary multi-team design setting, even at a sacrifice to performance. We argue in many of these cases a designer would not be interested in designs that are guaranteed to give worse performance than alternatives, thus using ZSD as an additional criterion is a useful option.

Regardless of the design scenario, for the optimization problem to be well posed, the intersection of the non-dominated sets must be non empty for any given set of designs $\mathbb{Y}$. If this were not the case, the combination of the dominance criteria does not give a preorder over the set of possible designs and an optimizer will not know how to progress.

When combining stochastic dominance with statistical moment Pareto dominance, the intersection is non empty since a better mean is a necessary condition for any order of stochastic dominance, and so at least the minimum mean design will be in both the Pareto non-dominated set and any stochastically non-dominated set.^{11,28} It is worth noting that combining different orders of stochastic dominance is equivalent to only using the highest order criterion, since lower order stochastic dominance is a necessary condition for higher order stochastic dominance.

Once the relevant dominance criteria have been selected, it remains for the optimization problem of Equation 1 to be solved numerically given a computational budget. Since global optimization heuristics are well suited to this optimization formulation, we suggest numerical implementations of the criteria and modify NSGA-II and MOTS to obtain approximations to a general non-dominated set.

E. Implementation

To implement the dominance criteria numerically, estimates of metrics such as mean, variance, quantile values and superquantile values must be obtained. Since q at a given design $\mathbf{x}$ is a random variable, q_x , we have two primary options for obtaining metric estimates within an optimization. Firstly, we can use estimators of metrics obtained from different random samples of $\mathbf{u}$ at each iteration, and use an optimizer that is resistant to noise. In this case, we need to be able to estimate the error in comparing two designs via the selected non-dominance criteria to ensure it is sufficiently low for the results of an optimization to be meaningful. When using quantiles and superquantiles, this is a delicate issue since the central limit theorem cannot necessarily be used, but progress can still be made via bootstrapping.²⁹ Equipped with an error estimate, the concept of non-dominance can be extended to this noisy case in various ways, such as “probabilistic dominance” as suggested in Ref. 30, or “probability of wrong decision” as suggested in Ref. 31.

Alternatively, in order to eliminate the noise from a given optimization, the sample average approximation method can be adopted,³² whereby the samples of the underlying random inputs, $\mathbf{U}(\omega)$, are taken once prior to the optimization, giving $\mathbf{u}^j$, $j \in \{1, \dots, N\}$, and are used at every design point to generate the samples of q_x . This is the approach used in this work, partly so that we can use deterministic multi-objective optimizers, but also because it is common in aerospace design problems to propagate the influence of uncertainties on q via the use of surrogate models such as polynomial chaos or kriging.¹ In this case we can subsequently sample the surrogate a large number of times at small computational cost, so that the error in estimating the empirical CDF from sampling becomes small compared to the error due to the discrepancy between the surrogate and the true output, which itself is small if the surrogate is sufficiently accurate.

Thus in this work, the information that is given to the optimizer for each design is taken to be the quantity of interest q evaluated at N samples of the underlying random inputs $\mathbf{U}(\omega)$: $q_x^j := q(\mathbf{x}, \mathbf{u}^j)$, $j = 1, \dots, N$, where the samples $\mathbf{u}^j$ are fixed throughout an optimization. The values of q can be evaluated directly from the system model or from a surrogate model fitted to q as a function of $\mathbf{u}$ for a given design $\mathbf{x}$. These N samples of q_x are then used to obtain numerical estimates of the metrics relevant to the dominance criteria being used.

Numerical Dominance Criteria

For Pareto dominance, estimators of the mean and standard deviation can be computed directly using the standard sample average and sample standard deviation statistical estimators, which are given by:

$$\hat{\mu}_x = \frac{1}{N} \sum_{j=1}^N q_x^j, \quad (8)$$

$$\hat{\sigma}_x = \left(\frac{1}{N-1} \sum_{j=1}^N (q_x^j - \hat{\mu}_x)^2 \right)^{\frac{1}{2}}. \quad (9)$$

For ZSD, $\max\{q_x^j \mid j = 1, \dots, N\}$ is compared with $\min\{q_y^j \mid j = 1, \dots, N\}$. FSD is determined by comparing the quantile function over N evenly spaced values of h over $[0, 1]$, and so empirical estimators for $F_x^{-1}(i/N)$, $i = 1, \dots, N$ are required. The empirical CDF can be found by ordering the samples of q_x resulting in N pairs of values $(q_x^{i:N}, i/M)$, where $q_x^{i:N}$ is the i^{th} order statistic of the N samples of q_x which satisfies:

$$q_x^{i:N} : q_x^{i-1:N} \leq q_x^{i:N} \leq q_x^{i+1:N} \quad \forall i \in \{2, \dots, N-1\}. \quad (10)$$

Thus the estimator of $F_x^{-1}(i/N)$ is the i^{th} order statistic $q_x^{i:N}$.

Similarly, SSD is determined by comparing the superquantile function over N evenly spaced values of h over $[0, 1]$, so empirical estimators for $\bar{Q}_x^{-1}(i/N)$, $i = 1, \dots, N$ are required. Here we use the sample average of all samples that are higher than the quantile estimator:

$$\bar{q}_x^i = \frac{1}{N-i+1} \sum_{j=i}^N q_x^{j:N}, \quad (11)$$

where $\bar{q}_x^i$ as the empirical estimator of $\bar{Q}_x(i/N)$.

Therefore introducing additional dominance criteria does not change the amount of computational effort per design analyzed (measured as number of evaluations of q), since exactly the same information is being used to obtain statistical moments as to compare types of stochastic non-dominance: the samples q_x^j , $j = 1, \dots, N$.

Using these estimators to evaluate the dominance criteria, Algorithm 1 outlines how a general optimization algorithm/heuristic is used to obtain an approximation, $\hat{\mathbb{S}}$, to the non-dominated set $\mathbb{S}$ that satisfies Equation 1.

Algorithm 1 Using an optimizer to obtain an approximation, $\hat{\mathbb{S}}$, to a non-dominated set $\mathbb{S}$

```

 $\hat{\mathbb{S}} \leftarrow \{\}$ 
while optimizer_not_terminated do
   $\mathbf{x} \leftarrow$  Optimizer selects new design
  if  $\hat{\mathbb{S}}$  is empty then  $\hat{\mathbb{S}} \leftarrow \{\mathbf{x}\}$ 
  should_add = True
  for  $\mathbf{y} \in \hat{\mathbb{S}}$  do
    x_dominates_y  $\leftarrow$  False
    y_dominates_x  $\leftarrow$  False
    for  $k$  in number_of_criteria do
      if  $\mathbf{x} \prec_k \mathbf{y}$  then
        x_dominates_y  $\leftarrow$  True
      if  $\mathbf{y} \prec_k \mathbf{x}$  then
        y_dominates_x  $\leftarrow$  True
    if x_dominates_y then
       $\hat{\mathbb{S}} \leftarrow \hat{\mathbb{S}} \setminus \mathbf{y}$ 
    if y_dominates_x then
      should_add  $\leftarrow$  False
  if should_add then
     $\hat{\mathbb{S}} \leftarrow \hat{\mathbb{S}} \cup \mathbf{x}$ 
Return  $\hat{\mathbb{S}}$ 

```

In this work we modify both a multi-objective genetic algorithm and a multi-objective Tabu search for use in Algorithm 1.

Modifying Optimization Algorithms

The genetic algorithm (NSGA-II) consists of selection, crossover, mutation, and replacement (for details see Ref. 22). The replacement operator is the only one that requires modifying for use with general non-dominance criteria, and it consists of non-dominated sorting and a crowding distance operator.

In standard multi-objective optimization concerning Pareto dominance, the non-dominated sorting assigns each candidate a rank R corresponding to whether that candidate would be on the Pareto front were the candidates on the first $R - 1$ Pareto fronts removed. The non-dominated sorting can therefore straightforwardly use the general non-dominance criteria instead of Pareto dominance.

The crowding distance is obtained by first sorting the population according to each objective, assigning the solutions with the smallest and largest objective values an infinite distance, and then assigning all other solutions indexed by $j \in \{2, \dots, N - 1\}$ the following Euclidian measure of distance (for the case where μ and σ are the two objectives):

$$c_j = \frac{\mu_{j+1} - \mu_{j-1}}{\mu_{max} - \mu_{min}} + \frac{\sigma_{j+1} - \sigma_{j-1}}{\sigma_{max} - \sigma_{min}} \quad (12)$$

The crowding distance operator is less trivial to modify for general dominance criteria, since there is no direct equivalent of Euclidian distance when comparing two CDFs. In this work we therefore continue to use the crowding distance as defined in Equation 12 even if we are not using Pareto dominance as a criterion. Alternative measures of the differences between CDFs such as the Kolmogorov-Smirnov distance³³ could be used instead, but we do not envisage that this would have a significant impact on the performance of the optimizers.

The Tabu search (MOTS) consists of a local Hook and Jeeves move, a diversification operator, an intensification operator, and a step size reduction operator (for further details see Ref. 23). The local search relies on a comparison operator that determines the non-dominated designs out of neighboring designs, so again this can straightforwardly be modified to use general non-dominating criteria. The other operators in general do not need to be modified, however some intensification operators make use of a crowding distance similar to the NSGA-II optimizer, for which Equation 12 can still be used.

The implementations of these optimizers used in this work have been coded in Python, and can be found and downloaded from: www-edc.eng.cam.ac.uk/aerotoools/dominanceoptimizers. We refer to them as the genetic algorithm and the Tabu search optimizers.

III. Algebraic Test Problem

To illustrate the difference between non-dominated sets under the different criteria, and how the performance of optimizers can be improved by making use of these criteria, here we perform investigations using the algebraic test problem given below. This test problem has two design variables and two uncertainties, and the quantity of interest q is given by:

$$\begin{aligned} p_1 &= 2 (\| \mathbf{x} - [0, 0.7] \|^2) - 0.4 \\ p_2 &= 5 (\| \mathbf{x} - [0, -0.7] \|^2) \\ q_1 &= 5u_1u_2x_1 \\ q_2 &= -x_2u_2^2 \\ q_3 &= x_2u_2^3 \\ q &= 10 + 2(q_1 + q_2 + q_3 + \min(p_1, p_2)) \end{aligned}$$

over a design space where both x_1 and x_2 are bounded by $[-1, 1]$, and where u_1 and u_2 are uniformly distributed over $[-1, 1]$. The optimizers are implemented as outlined in Section II.E, using 200 fixed samples of $\mathbf{U}(\omega)$ to obtain 200 samples of q_x at each design point. The key optimization parameters for the two optimizers are given in Table 1.

Genetic Algorithm		Tabu Search	
Population Size	25	Intensification Steps	10
Crossover Rate	0.9	Diversification Steps	15
Mutation Rate	0.7	Step Size Reduction Steps	22
Mutation Standard Deviation	0.175	Step Size Reduction Ratio	0.6
		Initial Step Size	0.2

Table 1: Key optimization parameters used by the optimizers

A. Single Optimizations

Firstly, optimizations are run to find the non-dominated sets under the following criteria: Pareto dominance, FSD, SSD, and both FSD and Pareto dominance, given a computational budget where a maximum of 250 design points can be visited. For the resulting non-dominated sets found by the optimizers, Figure 3 gives their mean and standard deviation and CDFs, and Figure 4 gives their design variables.

The traditional Pareto front of moments on Figure 3 shows a significant trade-off between mean and standard deviation: from the minimum mean extreme to the minimum standard deviation extreme, the standard deviation is reduced by $\simeq 96\%$ at a penalty of $\simeq 26\%$ to the mean. While this represents a significant improvement in robustness, the large penalty to the mean means that many of the robust designs are first order stochastically dominated by designs further along the Pareto front, since their CDFs do not cross. The most robust designs are even zeroth order stochastically dominated, such that they are guaranteed to give a worse performance than other designs, as annotated on Figures 3b and 3d.

If only FSD is used, the curve of the Pareto front continues past the minimum mean design, and many designs that are Pareto dominated but that have a small probability of achieving the lowest values of q are non-dominated under FSD, as annotated on Figures 3a and 3c. Thus using both Pareto dominance and FSD filters out these highly non-robust designs.

The Pareto front is split into two discontinuous sections that give designs with differently shaped CDFs and, as seen on Figure 4, are located in different areas of design space. If just SSD is used, only a small section of the Pareto front is non-dominated - all the designs from the lower section of the Pareto front are dominated. This is because, while the designs on the lower section are robust, their CDF has a long right tail which penalizes the superquantile heavily, as annotated on Figure 3d. In contrast, the designs from the upper section of the Pareto front have no right tail and their variance is mainly from a long left tail which improves their superquantiles.

(a) Genetic algorithm moments

(b) Tabu search moments

(c) Genetic algorithm CDFs

(d) Tabu search CDFs

Figure 3: Moments and CDFs of non-dominated sets obtained from optimizations on the algebraic test problem.

These results highlight the differences in using the different dominance criteria as discussed in Section II.D. It is notable that there is significant overlap between the different non-dominated sets, and so, for example, by finding the Pareto front we also have access to SSD non-dominated designs. However, the primary motivation for utilizing the different criteria is that under a limited computational budget it is often infeasible for optimizers to approximate the entire Pareto front. Therefore restricting the non-dominated set can reduce the size of the design space that contains non-dominated designs and so can improve the efficiency of the optimizers searching for these designs.

B. Convergence Over Multiple Optimizations

To investigate how using different dominance criteria can improve the efficiency of the optimizers in approximating the non-dominated set, here multiple optimizations are performed using both the genetic algorithm

Figure 4: Design variables of non-dominated sets obtained from optimizations on the algebraic test problem.

and Tabu search optimizers. For this, we specifically compare using SSD with using Pareto dominance. First the genetic algorithm is run with a population of 250 for 25 generations obtaining the SSD non-dominated set. This is then taken as the “true” optimal set of solutions, $\mathbb{S}_{true}$, and the convergence of the optimizers under a stricter computational budget to this set is investigated.

Convergence is measured using the average distance, in terms of Euclidian distance in mean/standard deviation space, from each design in the “true” set to its closest point in the current set of optimal designs at a given point in an optimization, $\hat{\mathbb{S}}$. Indexing the points in the “true” optimal set with i and those in the current optimal set as j , this measure is given by:

$$\text{Avg. Min. Distance} = \frac{1}{|\mathbb{S}_{true}|} \sum_{i=1}^{|\mathbb{S}_{true}|} \min\{\sqrt{(\mu_j - \mu_i)^2 + (\sigma_j - \sigma_i)^2} \mid j \in 1, \dots, |\hat{\mathbb{S}}|\} \quad (13)$$

A total of 10 optimizations are run using both the modified NSGA-II and Tabu search algorithms to obtain the Pareto and SSD non-dominated sets. Figure 5 plots the convergence of the Avg. Min. Distance given by Equation 13 in terms of number of design points evaluated for the best 8 runs as dotted lines, and then the average of these runs as a solid line.

Figure 5 illustrates that when finding an approximation to the SSD non-dominated set, a smaller area of design space is of interest and so the optimizers spend more of the computational budget exploiting the behavior of designs in this area instead of exploring other areas of design space. Whilst these results are only specific to this particular problem, it clearly illustrates the mechanism that can be exploited to accelerate optimizations towards designs of most interest in general. To explore to what extent this mechanism can offer benefit in physical design problems, two such problems, for which different dominance criteria are appropriate, are explored in Sections IV and V. Section IV investigates using SSD as a criteria to improve optimizations for an acoustic horn design problem, whereas Section V investigates using both Pareto dominance and FSD to improve optimizations for a transonic airfoil design problem.

(a) Genetic algorithm convergence

(b) Tabu search convergence

Figure 5: Convergence of the optimizers to $\mathbb{S}_{true}$ when using the Pareto dominance criterion and when using both Pareto dominance and FSD as criteria on the algebraic test problem.

IV. Acoustic Horn Design Problem

As an example of a simple design problem for which there is a large trade-off of statistical moments on the Pareto front (having been considered previously³⁴), we consider the shape optimization of an acoustic horn under uncertainty, minimizing the reflection coefficient subject to an uncertain wavenumber. This problem is highly non-linear and the CDF of reflection coefficient can have long tails of poor performance, and so second order stochastic dominance (SSD) is an appropriate criterion to use. Therefore on this problem using SSD is compared to using Pareto dominance.

The flow field is solved using a finite element discretization of the domain using 35895 states. The horn shape is parameterized by the width at 6 evenly distributed points along its length from root to tip, with the width at the root fixed at 0.5m, and at the tip fixed at 3.0m. These widths and their bounds are detailed in Table 2.

No.	Notation	Lower Bound [m]	Upper Bound [m]
1	$x_{1,nom}$	0.68	1.01
2	$x_{2,nom}$	1.04	1.38
3	$x_{3,nom}$	1.41	1.73
4	$x_{4,nom}$	1.77	1.99
5	$x_{5,nom}$	2.04	2.43
6	$x_{6,nom}$	2.48	2.82

Table 2: Design variables for the acoustic horn design problem

The upper and lower wall impedances are set to 50, and the wavenumber is uniformly distributed over the range $[1.0, 1.8]$. The quantity of interest is the reflection coefficient (a measure of the horns efficiency):

$$q = \left| \int_{\Gamma_{inlet}} u \, d\Gamma - 1 \right|, \quad (14)$$

where Γ_{inlet} is the inlet surface, and u is the acoustic pressure field. The acoustic pressure field for geometry given by the upper and lower bounds on the design variables, along with the resulting reflection coefficient, is plotted on Figure 6. This illustrates how the quantity of interest - the reflection coefficient - is governed by a highly non-linear physical phenomenon and so its probability distribution is unlikely to be well described by the first two statistical moments.

Figure 6: Magnitude of the complex acoustic pressure field for the horn geometry given by the lower and upper bounds on the design variables for a wavenumber of $k = 1.4$.

A. Optimization Results

We run optimizations using both the genetic algorithm and Tabu search with a maximum number of design points that can be visited of 100, using 200 samples of the uncertain wavenumber to analyze each design as outlined in Section II.E. The same initial population / starting point is used for optimizations under the different dominance criteria to give a fair comparison, and the resulting non-dominated sets are given on Figure 7.

With both optimizers, the non-dominated set obtained when only using Pareto dominance gives a large trade-off between mean and standard deviation, as seen on Figures 7a and 7b. The CDFs of reflection coefficient for many of the designs in this set have a long tail, but from Figures 7c and 7d, we can see that shortening this tail to reduce the variance requires penalizing the performance over the rest of the CDF so much that the designs become second order stochastically dominated. Thus when using SSD, a smaller area of design space is of interest and the optimizers are able to obtain a better approximation to the true non-dominated set using the given computational budget. Consequently, many of the designs resulting from the SSD optimizations both Pareto dominate *and* stochastically dominate many designs resulting from the Pareto optimizations, suggesting that this set of designs is more useful to a designer.

These results reinforce the observations made in Section III: when faced with a restrictive computational budget, using prudent dominance criteria (in this case SSD) instead of relying on Pareto dominance of statistical moments can accelerate the convergence of the optimizers towards the designs of most interest.

(a) Genetic algorithm moments

(b) Tabu search moments

(c) Genetic algorithm CDFs

(d) Tabu search CDFs

Figure 7: Moments and CDFs of the non-dominated designs obtained for the acoustic horn design problem after 100 design points visited.

V. Transonic Airfoil Design Problem

Robust airfoil shape optimization is a practical aerospace design problem that has been considered multiple times in the literature. Existing studies show that a lift-to-drag ratio optimization subject to uncertain Mach number gives a Pareto front with a large trade off between statistical moments.^{4, 7, 8, 18, 35} From the results of such studies, we expect that many of the designs on the Pareto front for such a problem will be stochastically dominated, and so expect to see improved results using both Pareto dominance and FSD in our proposed formulation. To investigate whether this is the case, we tackle a transonic airfoil design problem, using the freely available SU2 CFD solver to perform the analysis, which has been validated for transonic flow.^{36, 37} We investigate using Pareto dominance and using both Pareto and first order stochastic dominance in optimizations of lift-to-drag ratio.

A. Problem Setup

This problem considers the design of an airfoil subject to a uncertain Mach number, M , centered on the nominal value of 0.71 and a fixed angle of attack of $\alpha = 2^\circ$, starting from an initial NACA0012 airfoil design. The Mach number distribution is taken as the beta distribution that is plotted in Figure 8.

Figure 8: Distribution of the uncertain Mach number for the airfoil design problem.

The airfoil geometry is parameterized using the 14 Hicks-Henne bump functions given in Table 3 along with their bounds (which define the limits of design space). These bump functions are used to deform the initial NACA0012 geometry and mesh using SU2's built in mesh deformation capability.

Table 3: Surface, location, upper and lower bounds of Hicks-Henne bump functions for the airfoil shape parametrization (as a fraction of chord length).

Surface	Location	Upper Bound	Lower Bound
U	0.05	15×10^{-3}	-10×10^{-3}
U	0.20	25×10^{-3}	-15×10^{-3}
U	0.40	25×10^{-3}	-20×10^{-3}
U	0.60	25×10^{-3}	-20×10^{-3}
U	0.75	20×10^{-3}	-5×10^{-3}
U	0.86	15×10^{-3}	-5×10^{-3}
U	0.93	10×10^{-3}	0×10^{-3}
L	0.05	15×10^{-3}	-10×10^{-3}
L	0.20	25×10^{-3}	-20×10^{-3}
L	0.40	25×10^{-3}	-25×10^{-3}
L	0.60	25×10^{-3}	-10×10^{-3}
L	0.75	20×10^{-3}	-5×10^{-3}
L	0.86	15×10^{-3}	0×10^{-3}
L	0.93	10×10^{-3}	0×10^{-3}

The design is subject to a geometry constraint where a torsion box of dimension 0.5×0.087 chord lengths that just fits inside the initial NACA0012 geometry must also fit inside any new design, representing structural requirements. If the box cannot fit inside the geometry the following penalty is added to the lift-to-drag ratio (otherwise $g_{geom} = 0$):

$$g_{geom} = 400M(1 + 100d_{fit}) - 200 \quad (15)$$

where d_{fit} is the distance by which the box protrudes from the geometry.

In addition, since RANS CFD is unable to capture instabilities relating to the shockwave-boundary layer interaction (i.e. buffet), to avoid the optimizer obtaining designs that have a good steady state performance over the Mach number range but that are unlikely to be used in reality due to instability, we impose a buffet-related constraint on the behavior of the flow over the trailing edge. We use three constraints based on the criteria in Ref. 38 that are designed to indirectly measure the onset of buffet. We penalise a design for the following: if the onset of trailing edge separation occurs upstream of 90% chord, if the shock foot separation bubble and trailing edge separation join up, or if the total separated length is greater than 20% chord. The shock foot separation bubble is identified if there is a separation region starting within 0.15 chord lengths of the shock, which is detected by a gradient in C_p of greater than 5. If any of these constraints are violated, the following penalty function is added to the lift-to-drag ratio (otherwise $g_{sep} = 0$):

$$g_{sep} = 200M(1 + 1000A_{friction}) - 100 \quad (16)$$

where $A_{friction}$ is the area of negative skin friction coefficient on a plot of distance along chord vs skin friction coefficient.

Therefore the quantity of interest (to be minimized) returned after analyzing the airfoil is given by:

$$q = -\frac{C_L}{C_D} + g_{geom} + g_{sep} \quad (17)$$

B. Polynomial Chaos Surrogate

For practical design cycle time, we cannot afford to run the CFD model hundreds of times per airfoil to propagate the uncertainty, and so we augment the optimization loop with a non-intrusive polynomial chaos (NIPC) surrogate model. This involves evaluating the CFD model at several quadrature points over the uncertain domain, and then constructing a series of orthogonal polynomials that approximate the response for a given airfoil.^{25, 39, 40}

Therefore we need to establish how high the order of this NIPC expansion needs to be in order to accurately capture the behavior of an airfoil over the Mach number range. Figure 9 shows the convergence in terms of expansion order of a NIPC surrogate model to L/D for the baseline NACA0012 airfoil.

From these results, we determine that both the statistical moments and response are sufficiently converged with an order 10 NIPC expansion (requiring a CFD solve at 11 quadrature points per design), and so this is what is used in the subsequent optimizations.

C. Optimizations

Both optimizers are given a computational budget of 500 designs (corresponding to $500 \times 11 = 5500$ CFD runs), and both are run using just Pareto dominance and using both Pareto dominance and FSD. Figure 10 gives the statistical moments of all the designs visited by the optimizations. Figure 11 gives the statistical moments and CDFs of the non-dominated designs resulting from the optimizations. Animations of the progression of the optimizers can also be found at: www-edc.eng.cam.ac.uk/aerotoools/dominanceoptimizers.

The difference between the set of designs resulting from using different dominance criteria is stark. With the genetic algorithm, most of the designs obtained by using both dominance criteria almost entirely Pareto dominate *and* stochastically dominate those obtained by just using Pareto dominance. Therefore regardless of the designer's preferences, the set obtained by using both criteria is a more useful outcome of the optimization than the set obtained by using just Pareto dominance. Figure 10 reveals that this is because the optimizer spent more time evaluating designs that are robust but with poor expected performance compared when only using Pareto dominance as a criterion.

With Tabu search, Figure 10 reveals how both optimizations proceeded identically for two local search steps of the algorithm, and then diverge and explore different areas of design space. Therefore in contrast to

(a) NIPC statistics convergence

(b) NIPC response convergence

Figure 9: Convergence of an NIPC expansion to lift/drag of the initial NACA0012 airfoil design.

(a) Genetic algorithm points

(b) Tabu search points

Figure 10: Feasible points visited by the two optimizers for the airfoil design problem. Points visited by both optimizations are plotted in purple.

the genetic algorithm, the resulting non-dominated sets give two separate sections on a Pareto front, with the Pareto optimization favouring robust designs and the Pareto + FSD optimization favouring designs with better expected performance. However, Figure 11 once again shows that many of the robust designs are not only first order stochastically dominated, but are guaranteed to give a worse lift-to-drag ratio (corresponding to ZSD) than those obtained by using both criteria.

To investigate these resulting designs in more detail, we examine three designs taken from the Tabu search non-dominated sets labelled A, B and C on Figure 11b. Figure 12 gives their airfoil shapes (centered around the torsion box), their surface pressure distributions at the highest Mach number over the uncertainty range ($M = 0.79$), their lift and drag polars, and their CDFs of lift-to-drag ratio.

Figure 11: Moments and CDFs of the non-dominated designs obtained for the airfoil design problem after 500 design points visited.

Airfoil A, which has the best mean but is the least robust, is the thinnest, and has a significant “rooftop” pressure distribution on the upper surface that ends in a strong shock at higher Mach numbers. This gives good performance over lower Mach numbers but also produces a significant drag rise that onsets at comparatively low Mach number. Design B fattens the rear half of the airfoil slightly compared to design A which reduces the strength of the “rooftop” pressure distribution but delays the onset of the drag rise. Design C is a fatter airfoil with less camber, leading to an lower lift to drag ratio at all Mach numbers but also significantly delaying the onset of the drag rise.

Thus using the different dominance criteria has resulted in physically different sets of airfoils. This confirms the observation that using different dominance criteria led to the Tabu search optimizer exploring different areas of design space, highlighting the mechanism that we have observed throughout these investigations. As a result, with both the genetic algorithm and Tabu search, an improved set of designs was obtained by using both Pareto dominance and FSD compared to using just Pareto dominance.

Figure 12: Comparison of selected airfoils from the Tabu search optimizations

VI. Conclusion

In this work we have proposed a formulation for optimization under uncertainty that can obtain an improved set of designs over the traditional robust Pareto front (that trades off the first two statistical moments of the quantity of interest) using the same computational budget. The formulation uses an arbitrary number of dominance criteria, such as stochastic dominance, and seeks an overall non-dominated set equivalent to the intersection of non-dominated sets defined by each criteria. By filtering out designs that exhibit undesirable behaviour under uncertainty, such as being guaranteed to give worse performance than other designs, the convergence of optimizers to designs of most interest can be accelerated. We outlined several dominance criteria and discussed the design scenarios in which they might be useful.

We demonstrated how existing multi-objective optimization algorithms can straightforwardly be modified to obtain approximations to a general non-dominated set. In the implementation we used a fixed set of samples of the uncertain input parameters throughout an optimization in order to solve it as a deterministic

problem. Future work should therefore investigate the error in using alternative dominance criteria when this fixed sample approach is not appropriate such that different samples of the uncertain input parameters are taken for each design. In this case the error due to noisy estimates of moments, quantiles and superquantiles should be quantified and accounted for in the optimization.

Optimizations on an algebraic test problem illustrated the difference between non-dominated sets under several criteria, and illustrated the mechanism by which optimizations can be accelerated using a prudent choice of dominance criteria. We applied our proposed formulation to an acoustic horn design problem using second order stochastic dominance, and to a transonic aerofoil design problem using both Pareto dominance and first order stochastic dominance. In both cases, by using these choices of dominance criteria, an improved set of designs compared to using just Pareto dominance of moments was obtained under the same computational budget.

Acknowledgments

This work is funded by the Engineering and Physical Sciences Research Council (EPSRC) UK, under grant number EP/L504920/1. The authors thank D.B.P. Huynh for his development of the acoustic horn finite element model and K.E. Willcox for providing the code for this model.

References

- ¹Keane, A. J. and Nair, P. B., *Computational Approaches for Aerospace Design: The Pursuit of Excellence*, Wiley, New York, 2005.
- 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.
- ³Beyer, 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.
- ⁴Ghisu, 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.
- ⁵Tachikawa, B. H., Schiavazzi, D., Arima, T., and Iaccarino, G., "Robust Optimization for Windmill Airfoil Design under Variable Wind Conditions," *Centre for Turbulence Research Preceedings of the Summer Program*, 2012.
- ⁶Keane, 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.
- ⁷Dodson, M. and Parks, G. T., "Robust Aerodynamic Design Optimization Using Polynomial Chaos," *Journal of Aircraft*, Vol. 46, No. 2, 2009, pp. 635–646.
- ⁸Lee, D. S., Periaux, J., Onate, E., Gonzalez, L. F., and Qin, N., "Active Transonic Aerofoil Design Optimization Using Robust Multiobjective Evolutionary Algorithms," *Journal of Aircraft*, Vol. 48, No. 3, 2011, pp. 1084–1094.
- ⁹Zhong, X., Ding, J., Li, W., and Zhang, Y., "Robust Airfoil Optimization with Multi-Objective Estimation of Distribution Algorithm," *Chinese Journal of Aeronautics*, Vol. 21, No. 4, 2008, pp. 289–295.
- ¹⁰Liu, Y., Bai, J., and Livne, E., "Robust Optimization of Variable-Camber Continuous Trailing-Edge Flap Static Aeroelastic Action," *AIAA Journal*, Vol. 55, No. 3, 2017, pp. 1–13.
- ¹¹Ogryczak, W. and Ruszczyński, A., "From Stochastic Dominance to Mean Risk Models: Semiderivations and Risk Measures," *European Journal of Operational Research*, Vol. 116, 1999, pp. 33–50.
- ¹²Shapiro, A., Dentcheva, D., and Ruszczyński, A., *Lectures on Stochastic Programming: Modeling and Theory*, SIAM, 2009.
- ¹³Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D., "Coherent Measures of Risk," *Mathematical Finance*, Vol. 9, No. 3, 1999, pp. 203–228.
- ¹⁴Luedtke, J., "New Formulations for Optimization Under Stochastic Dominance Constraints," *SIAM Journal on Optimization*, Vol. 19, No. 3, 2008, pp. 1433–1450.
- ¹⁵Dentcheva, D. and Ruszczyński, A., "Optimization with Stochastic Dominance Constraints," *SIAM Journal on Optimization*, Vol. 14, No. 2, 2003, pp. 548–566.
- ¹⁶Dentcheva, D. and Ruszczyński, A., "Portfolio Optimization with Stochastic Dominance Constraints," *Journal of Banking and Finance*, Vol. 30, No. 2, 2006, pp. 433–451.
- ¹⁷Petrone, G., Iaccarino, G., and Quagliarella, D., "Robustness Criteria In Optimization Under Uncertainty," *Evolutionary and Deterministic Methods for Design, Optimization and Control*, 2011.
- ¹⁸Seshadri, P., Constantine, P., Iaccarino, G., and Parks, G., "A Density-Matching Approach for Optimization Under Uncertainty," *Computer Methods in Applied Mechanics and Engineering*, 2016.
- ¹⁹Cook, L. W. and Jarrett, J. P., "Horsetail Matching: A Flexible Approach to Optimization Under Uncertainty," *Engineering Optimization*, 2017.
- ²⁰Cook, L. W., Jarrett, J. P., and Willcox, K. E., "Extending Horsetail Matching for Optimization Under Probabilistic, Interval and Mixed Uncertainties," *AIAA Journal*, 2017.
- ²¹Quagliarella, D., Petrone, G., and Iaccarino, G., "Optimization Under Uncertainty Using the Generalized Inverse Distribution Function," *Modeling, Simulation and Optimization for Science and Technology*, Vol. 34, 2014.

- ²²Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, Vol. 6, No. 2, 2002, pp. 182–197.
- ²³Jaeggi, D., Parks, G., Kipouros, T., and Clarkson, P., "The Development of a Multi-Objective Tabu Search Algorithm for Continuous Optimisation Problems," *European Journal of Operational Research*, Vol. 185, No. 3, mar 2008, pp. 1192–1212.
- ²⁴Alexander, S., Coleman, T. F., and Li, Y., "Minimizing CVaR and VaR For a Portfolio of Derivatives," *Journal of Banking and Finance*, Vol. 30, No. 2, 2006, pp. 583–605.
- ²⁵Ghanem, R., "Polynomial Chaos in Stochastic Finite Elements," *Journal of Applied Mechanics*, Vol. 31, No. 5, 1990, pp. 607–634.
- ²⁶Xiu, D. and Karniadakis, G. E. M., "The Wiener-Askey Polynomial Chaos for Stochastic Differential Equations," *SIAM Journal on Scientific Computing*, Vol. 24, No. 2, 2002, pp. 619–644.
- ²⁷Lemieux, C., *Monte Carlo and Quasi-Monte Carlo Sampling*, Springer, 2009.
- ²⁸Levy, H., "Stochastic Dominance and Expected Utility: Survey and Analysis," *Management Science*, Vol. 38, No. 4, 1992.
- ²⁹Quagliarella, D. and Iuliano, E., "Robust Design of a Supersonic Natural Laminar Flow Wing-Body," *IEEE Computational Intelligence Magazine*, 2017.
- ³⁰Teich, J., "Pareto-Front Exploration with Uncertain Objectives," *Evolutionary Multi-criterion Optimization*, 2001.
- ³¹Hughes, E., "Evolutionary Multi-objective Ranking with Uncertainty and Noise," *Evolutionary Multi-Criterion Optimization*, 2001.
- ³²Kleywegt, A. J., Shapiro, A., and Homem-de Mello, T., "The Sample Average Approximation Method for Stochastic Discrete Optimization," *SIAM Journal on Optimization*, Vol. 12, No. 2, 2002, pp. 479–502.
- ³³Stephens, M. A., "EDF Statistics for Goodness of Fit and Some Comparisons," *Journal of the American Statistical Association*, Vol. 69, No. 347, 1974, pp. 730–737.
- ³⁴Ng, L. W. T. and Willcox, K. E., "Multifidelity Approaches for Optimization Under Uncertainty," *International Journal for Numerical Methods in Engineering*, 2014.
- ³⁵Cook, L. W. and Jarrett, J. P., "Robust Airfoil Optimization and the Importance of Appropriately Representing Uncertainties," *AIAA Journal*, Vol. 55, No. 11, 2017, pp. 3925–3939.
- ³⁶Palacios, F., Colonno, M. R., Aranake, A. C., Campos, A., Copeland, S. R., Economon, T. D., Lonkar, A. K., Lukaczyk, T. W., Taylor, T. W. R., and Alonso, J. J., "Stanford University Unstructured (SU2): An Open-source Integrated Computational Environment for Multi-Physics Simulation and Design," *51st AIAA Aerospace Sciences Meeting*, 2013, pp. 1–60.
- ³⁷Palacios, F., Economon, T. D., Aranake, A. C., Copeland, S. R., Lonkar, A. K., Lukaczyk, T. W., Manosalvas, D. E., Naik, K. R., Padron, A. S., Tracey, B., Variyar, A., and Alonso, J. J., "Stanford University Unstructured (SU2): Open-source Analysis and Design Technology for Turbulent Flows," *AIAA Journal*, Vol. 51, 2014, pp. 1–19.
- ³⁸Eastwood, J. P. and Jarrett, J. P., "Toward Designing with Three-Dimensional Bumps for Lift/Drag Improvement and Buffet Alleviation," *AIAA Journal*, Vol. 50, No. 12, dec 2012, pp. 2882–2898.
- ³⁹Xiu, D., Lucor, D., Su, C. H., and Karniadakis, G. E., "Stochastic Modeling of Flow-Structure Interactions Using Generalized Polynomial Chaos," *Journal of Fluids Engineering*, Vol. 124, 2002, pp. 51.
- ⁴⁰Hosder, S., Walters, R. W., and Balch, M., "Point-Collocation Nonintrusive Polynomial Chaos Method for Stochastic Computational Fluid Dynamics," *AIAA Journal*, Vol. 48, No. 12, 2010, pp. 2721–2730.