CHAPTER

20 Optimization in probabilistic domains: an engineering approach

Panagiotis Tsirikoglou[a], **Konstantinos Kyprianidis**[c], **Anestis I. Kalfas**[d], **Francesco Contino**[b]

[a]*Limmat Scientific AG, Zurich, Switzerland*
[b]*Université Catholique de Louvain, Thermodynamics and Fluid Mechanics Group, Louvain, Belgium*
[c]*Mälardalens Högskola, Energy Engineering, Västerås, Sweden*
[d]*Aristotle University, Department of Mechanical Engineering, Thessaloniki, Greece*

CONTENTS

20.1 Introduction

Engineering design is facing a transition from the classic one-off designs to a more complete design pipeline that includes design space exploration and optimization processes. This transition is led by the continuous endeavor for added value in products and processes and the strong objective-oriented policy for sustainable growth, while the big advancements in algorithmic development, numerical modeling, and computational resources are key enablers. A characteristic example of the merits of such transition is provided on aircraft design, which constitutes the application field of focus of the current work. Thus, examining the design process and related research over the last years, optimization is identified in different stages of product or service development aiming to gain substantial competitive advantages. In the field of aircraft design, competitive advantages can potentially have the form of fuel savings, noise reduction, an increase in passengers safety and comfort, and improved fleet management and flight data harvesting. On top of the industry needs for more sophisticated

Nature-Inspired Computation and Swarm Intelligence. https://doi.org/10.1016/B978-0-12-819714-1.00031-2

products, policy in the form of directives and legislation of aircraft industry is pushed to comply with future strategic goals for sustainable development, environment protection, and human safety.

Essentially optimization is an automated decision making process that incorporates couples search and decision making strategies to an engineering optimization case. This case describes in detail the problem considered for optimization, and hence roughly contains the design variables, parameters, and outputs of the underlying numerical model and the objectives/constraints that need to be optimized. Therefore, optimization is an objective-oriented design process, where the derived designs are propagated and selected as "best" with respect to a number of well-defined objectives. Thus, the definition of "best" design(s) in an optimization problem cannot stand alone as the given solution to address every need, but it is rather complementary to the objectives defined in the optimization problem.

In parallel, it progressively became well understood that the nature of engineering models and their variables or parameters is rather uncertain than deterministic. The impact of those input uncertainties is significant, and often optimal designs obtained by deterministic optimization approaches deteriorate from the desired performance point or even fail to satisfy critical constraints of the real engineering system. To address this issue the design variables and parameters are now defined using a probabilistic distribution function and a relative or absolute range in order to effectively describe the uncertainty, while an uncertainty propagation technique quantifies their effect on a quantity of interest, such as objectives or constraints. The described process enables the probabilistic design optimization (PDO) to define new designs of desired qualities, but also designs that retain those qualities under the presence of the input variations. As expected the integration of uncertainty quantification to the optimization process comes at a price. Furthermore, the propagation of those uncertainties to the quantities of interest of the case considered for optimization dictates the sampling of a new space on top of the exploration of the design space from the optimizer, as mentioned before. Thus, it is understood that PDO intensifies the computational cost issues that already exist in deterministic optimization cases (Bellman, 1961).

Until now, motivation, advantages, and shortcomings of the deterministic and probabilistic optimization have been identified. Moreover, a great number of optimization frameworks in both domains are developed and benchmarked against artificial optimization problems, where optimal designs are known (Li et al., 2013; Quagliarella et al., 2019). Despite the fact that the same frameworks are used, integration to engineering design is not a trivial process. The necessary pipeline that facilitates the automated evaluation of the objectives and constraints is a challenging task, particularly when various disciplines are combined, e.g., thermal power–electrical coupled simulations (Sahoo et al., 2019).

Not different to the rest of the engineering fields, integration of a full PDO framework to aircraft design cases presents the same difficulties. A wide range of numerical models are available, varying from fast, low-fidelity models to time consuming, high-fidelity approaches that capture large amounts of physics, such as computational fluid dynamics (CFD) methods. Moreover, the definition of the optimization problems can strongly challenge the convergence of the optimization algorithm towards a global optimal design when a large number of design variables and uncertainties are considered and/or design space is heavily constrained. Therefore, feasibility of PDO applications to aircraft design is not satisfied, and great care has to be taken in order to increase the overall computational efficiency. Minimizing the number of calls of the expensive, original engineering model greatly contributes to the reduction of the computational demands. The main enablers of this minimization are mathematical models (Forrester and Keane, 2009), often known as surrogate models, that can mimic the response of the original expensive model for an initial limited training database of designs. Apart from surrogate evaluation, the capability of an optimizer to locate fast the global optimal design and the computational demands of the uncertainty propagation technique are crucial to the computational efficiency of the overall PDO framework. It is understood that merits in computational time can be obtained from development in several research fields. The longstanding experience and continuous efforts so far have increased the technology readiness level of the various techniques and led to several software solutions, such as DAKOTA (Adams et al., 2015), UQLab (Marelli and Sudret, 2014), and OpenMDAO (Gray et al., 2019).

In the current chapter, an overview of the computational pipeline of a PDO framework is first provided. Then, we focus our analysis on three aspects of the pipeline that relate to the integration of PDO applications in aircraft design and strongly affect their feasibility: problem definition, surrogate models, and global optimization schemes. In the last section, the current state of the art of the PDO applications to aircraft design are reviewed and discussed in relation with the aforementioned aspects of the PDO computational pipeline.

20.2 Probabilistic design optimization framework

The integration of PDO to engineering cases requires the development and implementation of a well-defined pipeline that automates and iterates through multiple searches and decision making. Fig. 20.1 illustrates an indicative implementation of such pipelines.

It should be highlighted that the flowchart illustrates the pipeline executed in one iteration of the PDO run. The total amount of iterations is directly related to the computational budget. A block-by-block description of the pipeline is first delivered.

- **Design space:** Initial input and all candidate designs obtained by the optimizer in each iteration are sampled here. The design space is based on the definition of all design variables and their ranges. Finally, any possible constraints on the input design variables segment the design space, and are thus described here.
- **Uncertain space:** For each design in design space, a set of new designs is defined here, in order to quantify the impact of uncertainties in objectives and constraints. The size of this set depends on the uncertainty propagation technique and the

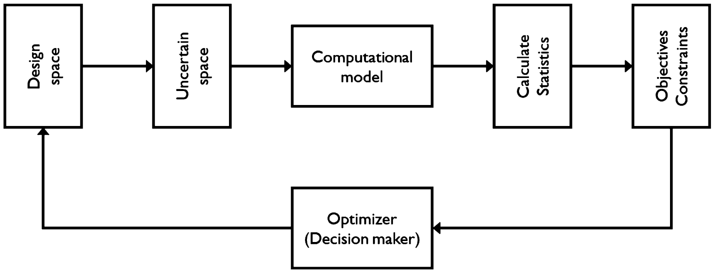

FIGURE 20.1

Computational pipeline of optimization in probabilistic domains. The pipeline describes one iteration of the optimization process. The overall optimization run constitutes a maximum number of iterations, defined by the computational budget.

dimensionality of the uncertain space. Finally, to perform this operation the probability distribution functions and ranges of all the uncertainties are employed.

- **Computational model:** The evaluation of the outputs of the engineering model is performed for the whole number of designs defined in design and uncertain space. The evaluation time is a key aspect that determines the overall computational budget and affects several choices regarding the structure of the PDO framework. To enable PDO for engineering models with relatively high evaluation time (a couple of minutes can be enough), the surrogate model totally or partially substitutes the original engineering model described above.
- **Calculate statistics:** The statistical measures of objectives and output constraints as defined in the optimization problem are calculated. The uncertainty propagation techniques sets the background methodology for this calculation.
- **Objective space:** The total set of objectives defined in the optimization problem form the objective space. Thus, all the calculated objectives and any possible output constraints are placed here.
- **Optimizer:** This block represents the selected optimization scheme. Two main operations are performed. Firstly, the candidate solutions are assessed and selected designs are propagated as the current optimal designs set. Moreover, the overall current best design, or set of designs (Pareto front) in the case of multiple objectives, is defined. Secondly, the search strategy is applied to current optimal designs set as described before, in order to further search the design space and provide new candidate solutions.

Please recall that the described actions are part of one iteration in the PDO framework. The number of iterations performed is dependent on the overall computational budget and the actual cost per iteration, i.e., how much original engineering model evaluation, nd how much time per evaluation are needed.

It is understood that the PDO framework consists of several methodologies that need to be effective and efficient to finally obtain a meaningful, optimal outcome in a realistic time frame. Moreover, its application to an engineering model requires a meticulous definition of the optimization problem, since the number of design variables, uncertainties, objectives, and constraints and their proper definition has a great impact on both the efficiency of the framework and the engineering impact of its optimal outcome.

Assuming that the optimization problem is defined properly and a robust engineering model is given, three methodologies of the PDO framework are important: the uncertainty propagation technique, the surrogate modeling, and the optimization scheme. Firstly, the propagation of the uncertainties is important for the accurate calculation of the statistical measures on the models' output, and the necessary amount of designs that are required to perform that operation. Moreover, the scale-up of the technique to a large number of uncertainties is of great importance for the integration of the PDO framework to more realistic engineering cases. A great deal of research is focused on developing such techniques (Abraham et al., 2017; Blatman and Sudret, 2011), aiming to accurately calculate the underlying partial differential equations of objectives and constraints using the lowest possible number of designs. Secondly, surrogate modeling can significantly extend the feasibility of PDO to more realistic, complex engineering cases. To achieve that, though, a proper surrogate management framework needs to be established in order to achieve good prediction accuracy and maintain it for cases that incorporate a large number of design variables and parameters. Thirdly, the selection of the optimization scheme affects the effectiveness of the search strategy and the quality of the decision making throughout the whole design process. Moreover, the class of the selected optimization scheme and its specific type determine the exploration and exploitation trade-off, and hence the amount of model evaluations needed until convergence to a local or global optimum solution.

In the current work we assume a robust engineering model and a state-of-the-art uncertainty propagation technique, and we focus on the following.

- **Problem definition:** The traits and effects of a proper problem definition are discussed.
- **Surrogate modeling:** The positioning of a surrogate model in a PDO framework and its managements is presented and thoroughly discussed.
- **Optimization scheme:** The selection of an optimization scheme is analyzed, based on the various trends of exploration/exploitation trade-off of the current state-of-the-art schemes.

20.2.1 Problem definition

Formulation of the deterministic optimization is the basis towards the proper definition of a PDO case. The conventional formulation of such a design problem uses a simple statement to link the objective(s) to the design variables, imposing input and output constraints:

$$\min_{x} \; J_m = f_m(x), \quad m = 1, 2, \ldots, M, \tag{20.1}$$

subject to

$$g_e(x) = 0, \quad e = 1, 2, \ldots, E, \quad h_i(x) \geq 0, \quad i = 1, 2, \ldots, I,$$

where the objective J_m is evaluated based on the outputs of the engineering model f_m seeking the global optimum design vector x^*. The search for this optimal design is subject to equality, g_e, and inequality, h_i, constraints that can affect the shape of the objective and design space, as described in the previous section. At this stage, the proper definition of design variables and its ranges and the selection of meaningful objective(s) are essential to lay the basis where PDO will be built upon. Last but not least, the proper definition of the constraints guarantees that the optimal outcome of the designs reflect the limitations of the actual technology.

Since a solid deterministic basis is formed, a problem statement for the final PDO case is needed. This additional layer constitutes the definition of the input uncertainties and the statistical formulations of the objectives and output constraints. These statistical measures quantify the effect of the uncertainties on the original objective J_m, hence becoming the new objectives of the optimization problem in the probabilistic domain. That said, the transformed new objectives need to address both the performance of the model and its sensitivity with respect to the defined input uncertainties. Two main approaches dominate the engineering design in probabilistic domains: robustness and reliability.

The term robustness is used to express and quantify the variance of each deterministic objective J_m, defined in Eq. (20.1), with respect to the input uncertainties. In this approach, each deterministic objective is transformed to its expectation and variance measures. Fig. 20.2 provides a visual example of this transformation, for a simple single-objective optimization problem that has one uncertain design variable.

To express that in an optimization case form, assume that the objectives, J_m, in Eq. (20.1) equal to one ($m = 1$, J_1), and hence we deal with a single-objective problem. As explained in the typical approach of robustness-driven design, the transition from deterministic to PDO doubles the number of objectives, since it is necessary to account here for both the expectation and the variance measures of the objective. The following equation describes the aforementioned rationale:

$$\min_{x} \; E[J_1], Var[J_1], \tag{20.2}$$

subject to

$$g_e(x) = 0, \quad e = 1, 2, \ldots, E,$$

$$h_i(x) \geq 0, \quad i = 1, 2, \ldots, I,$$

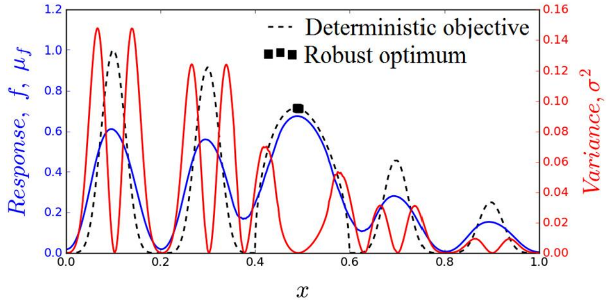

FIGURE 20.2

Transformation of a deterministic objective to an expectation measure and its variance under the presence of input uncertainties. The original deterministic objective is represented with the dashed line. As indicated, the new robust optimum is shifted to a more flat area, where variance is minimized.

where $E[J_1]$ and $Var[J_1]$ are the expectation and variance measures, respectively. Thus, the resulting optimization problem is biobjective. First, it should be identified if the relationship of the objectives is conflicting or not. For example, in Fig. 20.2 the resulting objectives in the probabilistic space are not conflicting, and hence one solution satisfies both. In the case of conflicting objectives, the extrapolation is straightforward, but note that higher-dimensional objective spaces provoke a few more issues in the computational pipeline of the PDO. The principal issue is that the presence of more than one conflicting objective changes the interpretation of "optimal" for the candidate designs through the course of optimization. In that case, the comparison of two candidate designs is not limited to the better/worse bipole, but it is rather described by the Pareto dominance relations (Deb et al., 2002). Thus, the optimal outcome is not a single (local or global) best design that could be captured for a given budget, but a set of designs in the form of an optimal Pareto front.

To avoid the increase in the number of the objectives and to utilize optimizers designed for single-objective decision making, collapsing of the various objectives to one weighted objective vector is followed. This collapsing is essentially the weighted summation of all the objective values in the objective vector. Despite the effectiveness of the method in reducing the size of the objective space, the optimal outcome consists of a single design, and hence the rest of design contained in the Pareto front is neglected. Repeated optimization runs with different assignment of weights in the collapsed objective vector can provide more designs that correspond to the true Pareto front. However, the accurate mapping between the weight values and the part of the Pareto front captured is not yet achieved.

As described, robust design optimization is the proper process to identify designs with good aspects and low variance at the same time. However, the operation of many engineering systems is characterized by various constraints, some of which prove critical. Therefore, reliability of the engineering systems under the presence of uncertainties becomes a primary design concern, and hence a different formulation of the optimization problem needs to be constructed. Assuming again a single-objective version ($m = 1, J_1$) of the classic deterministic optimization case (Eq. (20.1)) the reliability-based optimization problem is

$$\min_x E[J_1] \quad \text{or} \quad \min_x E[J_1] + wVar[J_1], \tag{20.3}$$

subject to

$$P[g_e(x) = 0] \leq P_o, \quad e = 1, 2, \ldots, E,$$

$$P[h_i(x) \geq 0] \leq P_o, \quad i = 1, 2, \ldots, I,$$

where $E[J_1]$ and $Var[J_1]$ are the expectation and variance measures, respectively, w is an assigned weight value, and P represents the failure probability of the respective constraints, which has to be lower than a certain threshold noted as P_o. The main differentiating point is the use of the probability failure in the constraints of the underlying model. In this way, a minimum level of reliability under uncertainties, controlled by the P_o threshold, is maintained throughout the PDO process. The accurate calculation of the probability failure is a challenging task, since it requires the evaluation of more designs. Thus, continuous development of efficient schemes for the calculation of the probability of failure are detrimental for the feasibility for such schemes.

20.2.2 Surrogate model

Time is the universal constraint applied in any aspect of life and finally life itself. As expected, the available computational time is the main limitation in probabilistic design and optimization as well. There are several components of the computational pipeline (see Fig. 20.1) that control the overall budget requirements, such as the uncertainty propagation method, the exploration/exploitation trade-off of the optimizer, and the execution of the engineering model. The latter is critical to the feasibility of the probabilistic design and optimization process, particularly in the engineering field, where the calculation time per sample is relatively high or a high number of design variables and parameters are usually considered.

To quantify this effect, an example from simple aerodynamics is used, where the panel method (Drela, 1989) and Reynolds-averaged Navier–Stokes (RANS) equations are utilized to calculate lift and drag coefficients of an airfoil. The RANS method, which incorporates a larger amount of physics, is more than 10 times slower than the fast, but low-physics, panel method. Moreover, the absolute number regarding calculation time is 10 mins for the RANS calculation (on 24 cores) and 0.5 mins (single-core) for the panel method. Considering the absolute numbers, the RANS method is not prohibitive and someone can assume that the probabilistic design and optimization should be easily applied. However, if the evaluation of a full-factorial design of experiment (DoE) (Garud et al., 2017) is considered, the difference in computational time demands becomes significant as the number of design variables and

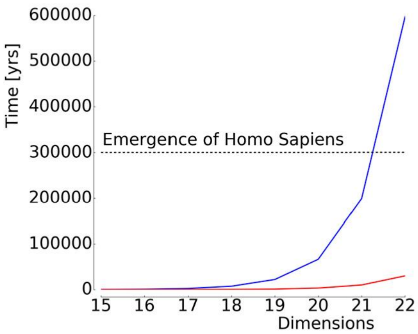

FIGURE 20.3

An illustrative example of the curse of dimensionality in engineering design, where an exponential increase in computational time is identified with respect to the number of design variables.

parameters is considered. Fig. 20.3 quantifies this difference for various numbers of design variables.

As can be seen, the increase in the number of variables, or search space in general, causes an exponential increase in the computational time needed to perform a full-factorial assessment of the model. This increase is not limited to the calculation of the full-factorial DoE, but it extends to operations such as optimization in deterministic and probabilistic domains, since searching strategies of the optimizers have to anticipate same-size or larger design spaces. In the case of PDO, the computational time demands are generally higher, due to the need of additional calculations to extract the statistics for each candidate design considered. Fig. 20.3 shows that a lower computational time per design delays the increase of the overall calculation time with respect to the design variables, thus making feasible probabilistic optimization for higher-dimensional design problems. Following the given example and considering the low complexity of the example given, it is apparent that the number of evaluations of the engineering model is crucial to the feasibility of the optimization process and the quality of the optimal outcome.

To enable the use of PDO in engineering design and extend its feasibility to more complex models or larger design spaces, approximations of the original engineering models are used. The aim of such approximation techniques, known as surrogate models, is to make sufficiently accurate predictions using the lowest number of original model evaluations possible. The significant computational advantages originating from their use led to the development of different kinds of surrogate models, such as kriging (Forrester and Keane, 2009) and its variants (Kleijnen, 2017), artificial neural networks (Cheng et al., 2016), polynomial regression (Forrester and Keane, 2009), and support vector regression (Smola and Schölkopf, 2004). The effectiveness and efficiency of the overall optimization scheme are controlled by two factors: the type of the surrogate models used and the structure of the scheme.

The first factor is related to the capabilities of the various models to derive quality predictions under certain conditions, e.g., linear or nonlinear original model responses. Chatterjee et al. (2019) performed a comparative assessment of different models indicating that anchored ANOVA, decomposition ANOVA, and polynomial chaos expansions (PCEs) (Wiener, 1938; Xiu and Karniadakis, 2002) are promising surrogates of the original engineering model, particularly for complex nonlinear responses.

The second factor is related to the structure of the overall optimization scheme, as regards the function and the management of the surrogate models within. Drawing experience from surrogate-based optimization schemes in deterministic domains, the management of surrogates is mainly divided in two cases. The first and simplest one is the a priori construction of the training database and the subsequent use of the surrogate models as global approximators over the whole design space. This management scheme is simple to implement, although the complete substitution of the original engineering model by the surrogate requires high prediction accuracy in the whole design space to derive optimal design close to reality. That is indeed feasible to achieve in lower-dimensional design spaces, and thus a lot of initial studies on surrogate-based optimization made use of this surrogate management structure. However, to obtain impactful designs, a high number of design variables and complex, high-physics models are needed, hence making the creation of a global approximator surrogate model a difficult task, both in terms of prediction accuracy and time demands. To tackle this rather strong limitation, adaptive formation of the training database of the surrogate model is suggested in the seminal study of Jones et al. (1998) and further developed in several studies (Shan and Wang, 2010). The basic principle behind all the adaptive management methods is the progressive build-up of the surrogate training database with designs that are of interest both for the course of the optimization run and the improvement of the prediction accuracy (Liu et al., 2018). That essentially translates to better prediction in design areas of interest.

While surrogate models as global approximators exist also in the probabilistic optimization field, adaptive formulation of the training database is a key aspect here. Moreover, the need to predict both the response of the original model and the behavior of its statistical measures adds a second level of prediction. To address this issue, one iteration of probabilistic optimization is split to the optimization and uncertainty loops, where different surrogate models are defined and used (Chaudhuri et al., 2019). The term “different” refers mainly to the designs that constitute their training databases and the type of data considered for prediction. Fig. 20.4 illustrates optimization and uncertainty loops in a typical iteration of a probabilistic optimization scheme (left) and the positioning of the surrogate models in such an iteration.

A common characteristic in some implementations of this multilevel surrogate modeling for probabilistic optimization is the combination of the design and uncertain space in one combined space (Arsenyev et al., 2015). This combined space allows the formulation of a training database that can describe the response of the original model with respect to all the parameters that induce a smaller (uncertainties) or larger (design variables) change to the inputs. The formulation of this design database al-

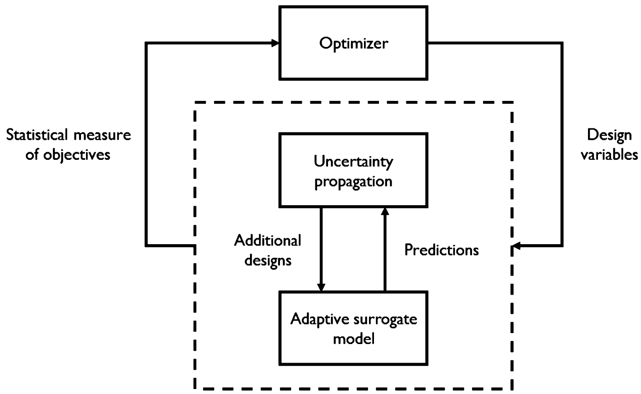

FIGURE 20.4

Positioning of an adaptive surrogate model in a sequential interaction of optimization and uncertainty loops. The flowchart is based on the work of Chaudhuri et al. (2019).

lows then the creation of the first-level surrogate and its sampling in order to obtain response data. The former, first-level surrogate is then coupled to the uncertainty propagation technique that needs to sample the surrogate model on a local scale in order to transform the input uncertainties to the statistical measure of preference. Based on these data, which link the design variables to the selected statistical measure, a training database is formulated, and the second-level surrogate model is fitted. The latter is then coupled to the search strategy of the optimizer, since it explicitly links the design variables to the objective considered for optimization.

A closer examination of the overall structure of the aforementioned scheme highlights again the approach regarding the formulation and treatment of the surrogate models. In particular, the training databases associated to the first-level surrogate model greatly affect the prediction accuracy of the surrogates in both levels. That said, one-off construction of that training database requires an a priori selection of designs that are well spread in the whole combined space (Arsenyev et al., 2015). Moreover, the size of the training database is increasing in order to guarantee a lower bound on the prediction accuracy of the first level surrogate, e.g., $(25\ 30) \cdot N_d$. Following the deterministic optimization paradigm, such techniques are limited to lower-dimensional spaces, and hence they are usually rendered insufficient for the particularly high-dimensional problems of PDO. Therefore, adaptive formulation allows for a smaller, initial training database, and the progressive definition of designs that are of interest for both the optimizer and the surrogate model (Liu et al., 2018).

Finally, the use of surrogate models within the context of probabilistic optimization raises the question of good prediction performance at both a local and global scale of the design space. Two aspects of the aforementioned optimization schemes are of interest here: combined space and adaptivity. Firstly, the combined space enables the definition of a DoE that samples sufficiently well the variations at both local and global scale. This advantage comes from the fact that the combined space has the shape of a hypercube, where some of its edges represent the range of the uncertainties and the rest the range of the design variables. As expected, a large number of uncertainties increase the hypervolume of the combined space, thus necessitating a large initial training database to achieve satisfactory prediction accuracy. To alleviate that, an initial sensitivity analysis on an uncertainty level is suggested. The analysis allows the reduction of the size of the combined space by adding only the most influential uncertainties.

Secondly, adaptive construction of the training database supports the good predictions at both local and global scale as the optimization run progresses. The reason behind that improvement is the gradual convergence of the optimization search strategy towards a specific area of the design space. This convergence reduces the Euclidean distance of the designs that are propagated to the training input database, thus allowing the surrogate model to predict more accurately a limited area of the design space. Despite the progressive improvement of the performance, the size of the initial training database and space filling properties of the designs have to be carefully defined to avoid any deterioration in the overall prediction performance of the defined surrogate models.

20.2.3 Optimization scheme

The effect of an optimization scheme is crucial to the quality of the optimal designs. That effect is reflected mainly in the ability of the optimizer to explore the design space and capture the optimal design, and the required budget to perform the aforementioned tasks. Those two effects are represented as the exploration/exploitation trade-off. In the current section, the exploration/exploitation trade-off is discussed for different classes of optimizers with a focus on global, nature-inspired optimization schemes. It should be highlighted that the goal is not to produce an explicit and detailed analysis of the existent optimization schemes. We rather aim to examine the overall developments in the optimization field with respect to their exploration/exploitation trade-off, hence projecting their potential to the PDO field.

The exploration/exploitation trade-off is considered as one of the key performance indicators of every optimization scheme. The interpretation of this indicator relies on the fact that every optimizer is essentially a decision maker, with an underlying search strategy that produces a series of candidate designs. This underlying searching strategy can be aggressive, thus producing designs that exploit the maximum improvement of the objectives under a narrow range of options. In contrast, other search strategies can be totally explorative, thus obtaining a variety of new candidate designs, though without considering the improvement of objective values. Those two aspects are conflicting in the search strategies that are developed so far, complying with the common perception that someone cannot explore more options if aggressive decision making is performed. Therefore, assuming that exploration and exploitation are objectives that need to be maximized in the current and future developed optimization schemes, their different combination should form a Pareto front. Fig. 20.5 demonstrates a qualitative interpretation of the exploration/exploitation trade-off in the form of a Pareto front as described above.

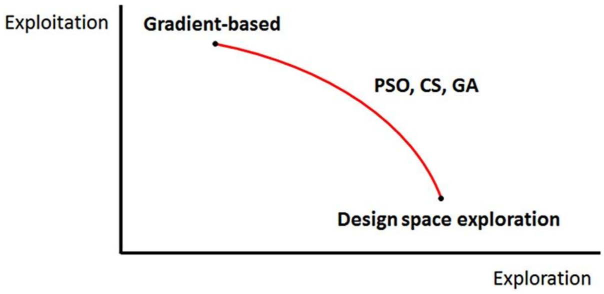

FIGURE 20.5

Qualitative Pareto front of exploration and exploitation as desired qualities of an optimization scheme. In the upper left fully exploitative gradient-based schemes are located, while in the lower right, extreme pure-design space exploration by space filling DoE is placed. Finally, all the gradient-free, nature-inspired optimizers lie in the middle part of the Pareto.

A brief analysis of the current developments in the optimization field is provided here, based on their location in the depicted Pareto. Starting from the upper left, the highly exploitative gradient-based schemes (Nocedal and Wright, 2006) are located. As the name states, gradient-based optimizers use gradient information to identify local areas of the design space that maximize the improvement of the defined objectives. Different mathematical approaches are employed in order to determine the search direction in each iteration (Nocedal and Wright, 2006) and calculate the gradients (Martins et al., 2003; Mader and Martins, 2012), thus modifying the overall search and computation efficiency. The main requirements for such optimization schemes are a starting design and the gradient information, while configuration of the algorithm is not an issue. Despite their significant computational efficiency and their sought mathematical background, gradient-based optimizers are strictly local search algorithms. Thus, their ability to locate the global optimum design is heavily dependent on the starting design, when complex engineer model responses are considered. To intensify the exploration and scale it up to a more global level, multistart, gradient-based optimization schemes have been developed (Chernukhin and Zingg, 2013). The definition of multiple new starting points allows the algorithm to perform the highly exploitative local search in several parts of the design space. Aiming to achieve the maximum of exploration of the design space, the multiple starting points are usually part of a space filling DoE, such as Latin hypercube sampling (Garud et al., 2017) or Sobol (Garud et al., 2017). This group of optimizers are located again in the upper left part of the Pareto front (see Fig. 20.5) but not in the extreme edge as the pure gradient-based optimizers.

Moving from the upper left to the middle and the lower right parts of the Pareto, the exploitation skills are more balanced to the exploration capabilities of the optimization schemes. In this area, the gradient-free class of optimizers is located. Gradient-free algorithms are a big family of stochastic optimizers that operate using only the value of the objectives to assess candidate designs and guide their search strategy. The simplicity and robustness of their search strategy allows them to perform under noncontinuous and complex responses of the original engineering model. The inputs of the algorithm consist of a set of initial designs and the configuration of their parameters. The latter is of great importance for the performance of the algorithm, while the definitions of the configuration set is a nontrivial task, due to the intrinsic randomness of their search strategies.

Several implementations of the gradient-free optimization paradigm exist in the literature (Boussaïd et al., 2013) aiming to capitalize the good search capabilities, while the development of this type of optimizers remains an animated topic of research. The main differentiation point between those implementations is the core research strategy often inspired by search patterns or processes that exist in nature (Kennedy and Eberhart, 1995; Michalewicz, 1995; Storn and Price, 1997; Yang and Deb, 2010; Yang, 2009). Among others, the original implementation of particle swarm optimization (PSO) (Kennedy and Eberhart, 1995) is a classic example of combination of exploration and exploitation skills. The optimizer exploits the food search mechanisms of birds in order to build an optimized search strategy. The particles (generalizing the term from bird) represent the candidate designs that the optimizer captures throughout the course of optimization. The critical part of this search mechanism is the definition of the velocity for each particle,

$$v_{i,d}^{t+1} = v_{i,d}^{t} + c_1 r_1^t (P_{i,d}^t - X_{i,d}^t) + c_2 r_2^t (P_{g,d}^t - X_{i,d}^t), \tag{20.4}$$

where $X_{i,d}^t$ is the position of particle i in the design space during generation t and $P_{i,d}^t$ and $P_{g,d}^t$ are the local and global best position for the dimension d, respectively. Moreover, c_1 and c_2 are constants regulating two important terms in the updated velocity formula: cognitive and social. These two terms represent the two different mechanisms that a particle learns. The cognitive mechanism enables particle i to adjust its velocity towards the best position encountered until now, while the social mechanism allows the particle to drift its velocity towards the best solution captured by the whole swarm of particles. Finally, r_1^t and r_2^t represent random numbers uniformly distributed in $[0, 1]$. The randomization added to the search step enhances the exploration of the design space, while it helps to avoid the entrapment in local optima.

The numerous developments (Bonyadi and Michalewicz, 2017; Poli et al., 2007; Harrison et al., 2018) based on this initial simple idea are a strong indication of the capabilities of the PSO scheme, but also for the nature-inspired, global optimizers. Evidently, a quick examination of the research outcomes in the optimization development field demonstrates several new optimizers and variants of the most powerful ones (Tilahun et al., 2019; Al-Dabbagh et al., 2018; Jayabarathi et al., 2018). Despite the large production of optimization schemes, a relatively small fraction is used in expensive engineering design problems, due to their high computational demands and the lack of integration of novel optimizers to complex, i.e., multilevel, surrogate evaluation schemes. Using again the example of the Pareto front (see Fig. 20.5), efforts needs to be made in order to build upon the current developments and push the Pareto front to the direction where exploration and exploitation are both enhanced.

To this end, gradient-free and gradient-based schemes are coupled to optimization frameworks that aim to benefit from the advantages of both. The so-called hybridization employs a gradient-based optimizer and a switching criterion related to the convergence of the designs, in order to activate the gradient-based optimization scheme and increase the convergence speed (Bos, 1998). Despite the simplicity of the idea, it provided results that prove its benefits in real engineering problems (Chernukhin and Zingg, 2013; Vicini and Quagliarella, 1999; Bos, 1998). The merits of such optimization frameworks can be further extended by capitalizing the recent advancements in both optimization algorithmic development and surrogate modeling, aiming to create an optimization scheme with enhanced exploration/exploitation trade-off and maximized computational efficiency for PDO of expensive engineering cases, such as aircraft design.

20.3 Probabilistic optimization in aircraft design

In this section the advancements in aircraft design are examined from the PDO standpoint. Moreover, the different aspects of the computational pipeline as discussed above are linked to the different aspects or levels of aircraft design.

An aircraft as a finalized product is overly complicated, containing a very high number of components and several systems in place. Therefore, on a research level several engineering design frameworks are used to investigate single components, systems, and sets of systems at the aircraft level. To further analyze the overall research outcome, Fig. 20.6 illustrates the different levels of aircraft analysis and design.

Starting from the top, the aircraft level investigates the behavior and dynamics of the whole aircraft assuming that the output of the defined model encounters a number of phenomena originating from the different systems and components. The intractable complexity of the process necessitates the use of several assumptions in order to build the engineering model under consideration. The assumption making is essential to the reliability of the model and the engineering impact of the results. As expected, the fidelity of the computational approaches followed is low, due to limitations in the development of such large-scale engineering models and the computational resources.

The aircraft design as a whole, consists of several systems responsible for different functions, such as propulsion, electrical, landing, airframe, and more. The reduction in underlying complexity allows for an increase in the fidelity of the computational approaches. Despite the reduced levels, simplification through assumptions and lower-physics modeling is still necessary to render feasible studies of this scale. Finally, aiming to achieve more complete and thoroughly investigated designs, coupling of the various systems is becoming more and more necessary. Therefore, integrated systems design is proposed as one of the suitable approaches to increase the engineering impact of the designs produced.

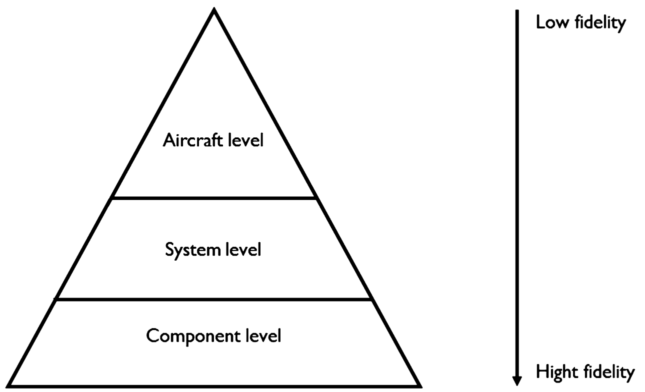

FIGURE 20.6

Segmentation of aircraft design: aircraft level, system level, and component level. As moving towards the integration of all components and systems in one entity, underlying models are simplified, i.e., levels of fidelity are decreasing.

Moving to the bottom of the design pyramid (see Fig. 20.6), components such as compressors, turbines, combustors, landing gear, electrical wiring, and many more are the basis of the whole design approach in aircraft vehicles. The relatively minimal complexity allows for the use of detailed designs and high-fidelity models such as CFD (Skinner and Zare-Behtash, 2018) approaches. Despite the increased level of details encountered in the design process, interactions with other components are largely neglected.

20.3.1 Applications on aircraft and system level

As described in the introduction section, engineering design in general is transitioning from a solely one-off design approach to a complete pipeline process that contains design space exploration, uncertainty quantification, and optimization in deterministic and probabilistic domains. In the current section, we examine the current status and outcomes of this design phase transformation regarding the first two top-level approaches in aircraft design, namely, the aircraft and the system level. Those approaches are treated together due to some similarities of the underlying models considered for optimization and the definition of the problem.

One-off design and assessment processes are the basis of aircraft development, since a nominal complete design is necessary to be obtained at least in the early design stage. Evidently, a lot of research studies are dedicated to the development of new design ideas (Drela, 2011a; Hall and Crichton, 2005) aiming to provide a complete view of a new candidate design. Building on this, new numerical schemes for design evaluation and optimization, and evolution of the computational framework are the enablers of research studies on the exploration and optimization of the initial one-off designs.

Aiming to demonstrate the described transformation, the example of the D8 transport aircraft configuration (Drela, 2011a) is discussed. The development of the D8 aircraft is supported within the N+3 initiative of NASA, aiming to shape future aircraft vehicles with significantly reduced fuel consumption. The initial development of the D8 aircraft is followed by further design examination, such as by Yutko et al. (2017), where boundary layer ingestion is introduced and the initial structure of the airframe is refined. Further than the one-off designs, Drela (2011b) demonstrated that deterministic optimization can produce more optimized variants of the initial aircraft, considering a number of different systems and assessing the various designs on a mission level analysis. A critical factor of the successful optimization run is the development of a set of models, namely, the Transport Aircraft System OPTimization (TASOPT) (Drela, 2010), which can assess different aircraft designs at feasible computational budgets. Finally, based on the strong basis of the available initial designs and efficient engineering models, probabilistic optimization of the D8 aircraft was performed by Ng and Willcox (2016). Considering the significant increase in computational costs, due to the need of a statistical estimator for every candidate design, an information reuse method is developed and coupled with the classic Monte Carlo (MC) method, achieving reductions of 90% compared with the original MC. As regards the optimal design outcome, a new aircraft configuration is captured that achieves 84% of predefined performance criteria, compared with an initial 22%.

Keeping aside the example of the complete design pipeline of the D8 aircraft, more probabilistic optimization studies on an aircraft design level are identified. Jaeger et al. (2013) introduced uncertainties in a short-range aircraft conceptual design optimization problem. Similar to the design of the D8 example, the availability of simple models for the aircraft design is one of the main enablers of this study (Birman and Druot, 2011). The modeling framework facilitates a range of models regarding the environmental impact, fuselage, wings and tail, propulsion system, and landing gear able to provide assessments of conceptual designs at the aircraft level. The performed optimization obtained robust aircraft designs that exhibit a 90%–95% chance of achieving the specified geometry and performance constraints. In a further attempt to integrate PDO to the conceptual design stage, Clark et al. (2019) used surrogate models as simplified, low-fidelity models to derive robust configurations of a generic fighter aircraft under mission uncertainties. The idea of using surrogate models in that design stage is particularly supported by a nondeterministic implementation of kriging (Bae et al., 2019), providing a reliable approximation tool for future use within such optimization frameworks. In all studies considered, the set of engineering models for the calculation of the quantities of interest are simplified, hence inducing some epistemic uncertainties. To assess the impact of those uncertainties, Molina-Cristóbal et al. (2014) introduced a novel epistemic uncertainty propagation technique coupled to a black-box modeling approach. The work targets the gas turbine system and its interactions with the airframe, thus focusing more one the integrated system design than the whole aircraft. Robust engine configurations that reduce the effect of the epistemic model uncertainties are obtained.

In previous sections, three aspects of the PDO pipeline were discussed: problem definition, evaluation of the engineering model, and optimization schemes. In the last part of this section a status of the applications as described above in relation to those three aspects is provided.

Firstly, probabilistic optimization coupled to the aircraft- and system-level design strongly relates to the problem definition. Despite the fact that the objective in such studies usually refers to some weight indicator on a mission level, i.e., maximum take-off weight, problem definition needs careful treatment regarding the constraint definitions. The latter are rather important due to the fact that constraints in the conceptual design phase represent critical technology limitations. Therefore, the complete definition of the probabilistic optimization problem significantly affects the engineering impact and feasibility of the underlying study.

Secondly, the evaluation of the engineering model is a relatively trivial task, since fast, simplified models are used at this level of design. Thus, the use of surrogate models, as global or adaptive approximators, valuable as it is considered, does not critically alter the feasibility boundaries of such applications.

Thirdly, the selection of the optimization scheme is slightly biased towards the gradient-based schemes. This preference stems from the strong mathematical background of those schemes, associated to their efficient handling of heavily constrained design spaces. Moreover, the relatively low evaluation cost of the engineering model enables the calculation of the necessary gradient information even with more computational time consuming techniques. This bias though does not exclude the use of derivative-free schemes. Particularly gradient-free optimizers designed for constrained spaces, such as COBYLA (Powell, 1994), come also into play when large search spaces (design or combined) with possible complex responses are formed. The low evaluation cost of the original engineering model again can support the increased computational demands of such optimization schemes. As expected, the availability of relatively fast engineering models simplifies the computational pipeline and allows for some flexibility in the choices of the optimization scheme. However, careful selection is necessary, since even low absolute evaluation time, e.g., several minutes, can result in computationally intractable optimization runs (see Fig. 20.3).

20.3.2 Applications at the component level

The literature analysis on the design optimization approaches applied to the aircraft and system levels indicates that PDO is gradually becoming part of the design, even in early stages. To complete the literature analysis, the status of the probabilistic design and optimization methods applied to the design of components and subcomponents of aircraft vehicles is examined.

Fundamental differences of component design with respect to aircraft and system design are identified in the response type and level of fidelity of the underlying engineering model and in the problem definition. The latter, in the case of component design, incorporates design specifications, variables, and parameters from a single technology domain, thus neglecting possible interactions. This does not mean of course that such optimization problems are simple to solve.

A vast number of research studies are available in the literature regarding the standalone design and deterministic optimization of aircraft components and subcomponents. The popularity of the subject stems from the large number of different components and subcomponents and the availability of different fidelity levels in the numerical approaches. Those two characteristics significantly increase the potential number of cases considered for design optimization, while they allow for a large variance in the complexity of the final case. Among others, airframe components such as wings and airfoils and gas turbine components such as turbine blades and compressors are thoroughly investigated and well documented (Skinner and Zare-Behtash, 2018; Du et al., 2017; Amrit et al., 2017; Chen and Agarwal, 2014). Due to this rich basis of design cases, penetration of deterministic optimization increases along with the advancements in different aspects of optimization pipeline, such as optimizers and surrogate modeling.

Based on the background described above, the animated research around uncertainty quantification and the development of more efficient uncertainty propagation techniques enables the realization of optimization in probabilistic domains. One of the first emerging and most popular applications is the PDO of an airfoil (Choi and Kwon, 2014; Wu et al., 2018; Rumpfkeil, 2012). The relatively low complexity of the overall optimization case and the continuous development of the lower- and higher-fidelity numerical approaches linked to the evaluation of the objectives allowed the fruitful production of research outcomes. Despite the generally lower evaluation times per design, the use of surrogate models is still necessary in order to increase the size of the combined space, thus creating a more comprehensive and detailed optimal design. Several studies employ a multilevel prediction scheme, as described in Section 20.2.2. As discussed, adaptivity and the definition of multiple levels in prediction improve the performance of the prediction scheme at a global and local scale. On that issue, Rumpfkeil et al. (2017) introduced a different technique, where clustering of the available training data is followed by the definition of multiple local surrogates. The different approximations made at the local level are probabilistically combined in an agglomerated final estimation. By extending these features to a multifidelity approach, good prediction accuracy is achieved for a limited size of the initial training database. Finally, the extensive study of this specific design problem led to the definition of benchmarks that use a specific airfoil geometry (Quagliarella et al., 2019), aiming to assess in a systematic and effective way novel PDO frameworks. Therefore, it is understood that the maturity level of methods related to uncertainty quantification is increasing.

In the current study, airfoils are considered as a partial design problem (sort of a subcomponent) of other components in the engine and the airframe systems. Due to the same reasons as the airfoil design problem, the wing, a component of the airframe system, is particularly well studied. Extending to PDO, wing-related applications are delivered (Liang et al., 2011), proving the readiness level of such design optimization problems as well. Due to this maturity, extensions of the former design problem to a multidisciplinary setting are identified. Jacome and Elham (2017) further optimize the wing geometry of a popular civil aircraft, using coupled aerodynamics-structural computation over a surrogate-assisted mission analysis. The successful optimization of the wing shape for given deterministic and probabilistic formulations of the objectives is performed in a cost-effective manner, while past flight data are used to define realistic uncertainty ranges. The latter is of great importance since a realistic problem definition significantly increases the engineering impact of the studies. In the same manner, but in a different component setting, Kamenik et al. (2018) use laser scans to model the manufacturing uncertainties of a high-pressure turbine blade. The defined uncertainties then serve as inputs in a probabilistic optimization scheme, aiming to derive robust shape designs. Despite the importance of realistic uncertainty ranges (see Section 20.2.1), a handful of research works combines the efforts of quantifying them with the optimization process.

Comparing to the application of probabilistic optimization at the aircraft and system levels, component-level design illustrates a significantly higher number of applications due to the higher number of possible cases, and their varying complexity. Following the same procedure as in the previous section, the application at the component level is discussed in relation with the three aspects of the computational pipeline of the probabilistic optimization, highlighted in the previous section: problem definition, evaluation of the engineering model, and optimization schemes.

Problem definition is a critical part in every optimization problem, both in deterministic and probabilistic domains. Different from the aircraft and system levels, the objectives here significantly vary, due to the focus on specific and not similar components. From the constraints standpoint, a varying level of complexity is also identified. In reality engineering search spaces (design and combined) are usually constrained, thus necessitating a more complex problem definition, if maximization of engineering impact is desired.

In a reverse trend, the engineering model is evaluated, through surrogate models are critical to the feasibility in most of the studies. This strong dependence stems from the generally increased amount of physics captured, pursuing an in-depth investigation of the isolated component. Kriging and its variants (Kleijnen, 2017) are key enablers of surrogate evaluation, while the research focuses on the construction of adaptive, multilevel surrogate prediction frameworks that can address the need of high prediction accuracy at both local and global levels.

Finally, a relatively strong bias towards the local gradient-based algorithms is identified. The main reason behind that choice is the generally increased computational demands for the evaluation of one design, e.g., turbine blade design cases. Therefore, the increased available budget per design limits the number of original model evaluations. Therefore, local, gradient-based optimizers are selected to further evolve the components considered for optimization. Moreover, adjoint formulations in shape optimization problems cheaply provide gradient information even in high-dimensional search spaces. To combine the necessary strong exploitation skills with more exploration, thus more alternate and possibly impactful designs, the use of hybrid scheme is suggested.

20.4 Conclusions

PDO applications in engineering, and particularly in aircraft design, are the results of an ongoing transition from one-off design processes to more complete design frameworks in the probabilistic domains. Longstanding research efforts focusing on a wide range of numerical techniques and algorithms support the development of such frameworks and their applications mainly through the maximization of the computational efficiency.

In this continuous pursuit of computational efficiency, two critical aspects of the PDO pipeline were identified and examined: the surrogate modeling and the selection of the optimization scheme. As regards the first one, global approximation using a priori training databases cannot be of use in the current status of engineering design problems, due to the higher-dimensional search spaces of the probabilistic optimization, i.e., for the same accuracy training databases of increasing size are needed. Moreover, PDO raised the additional demand of prediction accuracy at both global and local levels. A dominant solution captured was the multilevel, adaptive surrogate evaluation, creating effective prediction schemes at significantly reduced costs. However, the good performance of the surrogate at both local and global levels at the first iterations of the PDO process needs to be more investigated.

The maximization of the overall computational efficiency is strongly supported by an optimization scheme with strong exploration and exploitation skills. To establish a wide assessment, gradient-based and gradient-free optimizers were discussed within the context of the exploration/exploitation trade-off. The ongoing debate between the gradient-based and gradient-free algorithms, as the right way-to-go in engineering optimization, will be answered by optimization schemes that effectively combine the global exploration of the gradient-free algorithms and the local exploitation of the gradient-based schemes. To this end, building from the source optimization frameworks can facilitate good global exploration skills from nature-inspired algorithms and exploitation from the gradient-based ones.

Finally, the aircraft design field was examined with respect to the status of the PDO applications. Many interesting methods were identified, in relation to the surrogate evaluation and optimization scheme. However, one main characteristic was highlighted as a strong skill: knowledge of the PDO problem. The knowledge of the problem considered for optimization, which directly reflects the quality of the problem definition, has great impact on the derivation of meaningful results. Therefore, more studies that can define the ranges and types of uncertainties are encouraged, aiming to obtain a general picture regarding the fidelity of the technologies involved. Moreover, the continuous efforts to understand the aspects of different problems will enable safe generalizations and classifications, allowing the development of optimization frameworks operating on a more case-dependent basis.

References

Abraham, S., Raisee, M., Ghorbaniasl, G., Contino, F., Lacor, C., 2017. A robust and efficient stepwise regression method for building sparse polynomial chaos expansions. Journal of Computational Physics 332, 461–474.

Adams, B., Bauman, L., Bohnhoff, W., Dalbey, K., Ebeida, M., Eddy, J., Eldred, M., Hough, P., Hu, K., Jakeman, J., Stephens, J., Swiler, L., Vigil, D., Wildey, T., 2015. Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 6.0 user's manual. Sandia Technical Report SAND2014-4633.

Al-Dabbagh, R.D., Neri, F., Idris, N., Bab, M.S., 2018. Algorithmic design issues in adaptive differential evolution schemes: review and taxonomy. Swarm and Evolutionary Computation 43, 284–311.

Amrit, A., Du, X., Thelen, A.S., Leifsson, L.T., Koziel, S., 2017. Aerodynamic Design of the RAE 2822 in Transonic Viscous Flow: Single- and Multi-Point Optimization Studies.

Arsenyev, I., Duddeck, F., Fischersworring-Bunk, A., 2015. Surrogate-Based Robust Shape Optimization for Vane Clusters.

Bae, H., Clark, D.L., Forster, E.E., 2019. Nondeterministic kriging for engineering design exploration. AIAA Journal 57 (4), 1659–1670.

Bellman, R., 1961. Adaptive Control Processes: A Guided Tour. Princeton University Press.

Birman, J., Druot, T., 2011. Robustness assessment of a margin-setting process in preliminary aircraft sizing. In: 12e congrès annuel de la société Française de recherche opérationnelle et d'aide à la décision.

Blatman, G., Sudret, B., 2011. Adaptive sparse polynomial chaos expansion based on least angle regression. Journal of Computational Physics 230 (6), 2345–2367.

Bonyadi, M.R., Michalewicz, Z., 2017. Particle Swarm Optimization for Single Objective Continuous Space Problems: A Review.

Bos, A., 1998. Aircraft conceptual design by genetic/gradient-guided optimization. Engineering Applications of Artificial Intelligence 11 (3), 377–382.

Boussaïd, I., Lepagnot, J., Siarry, P., 2013. A survey on optimization metaheuristics. Information Sciences 237, 82–117.

Chatterjee, T., Chakraborty, S., Chowdhury, R., 2019. A critical review of surrogate assisted robust design optimization. Archives of Computational Methods in Engineering 26 (1), 245–274.

Chaudhuri, A., Marques, A.N., Lam, R., Willcox, K.E., 2019. Reusing Information for Multifidelity Active Learning in Reliability-Based Design Optimization.

Chen, X., Agarwal, R.K., 2014. Shape optimization of airfoils in transonic flow using a multi-objective genetic algorithm. Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering 228 (9), 1654–1667.

Cheng, T., Wen, P., Li, Y., 2016. Research status of artificial neural network and its application assumption in aviation. In: 2016 12th International Conference on Computational Intelligence and Security (CIS), pp. 407–410.

Chernukhin, O., Zingg, D.W., 2013. Multimodality and global optimization in aerodynamic design. AIAA Journal 51 (6), 1342–1354.

Choi, S., Kwon, H.I., 2014. Robust Design Optimization Using a Trended Kriging Surrogate Model and Applications to Unsteady Flows.

Clark, D.L., Allison, D.L., Bae, H., Forster, E.E., 2019. Effectiveness-based design of an aircraft considering mission uncertainties. Journal of Aircraft 56 (5), 1961–1972.

Deb, K., Pratap, A., Agarwal, S., Meyarivan, T., 2002. A fast and elitist multiobjective genetic algorithm: Nsga-ii. IEEE Transactions on Evolutionary Computation 6 (2), 182–197.

Drela, M., 1989. Xfoil: an analysis and design system for low Reynolds number airfoils. In: Mueller, T. (Ed.), Low Reynolds Number Aerodynamics. In: Lecture Notes in Engineering, vol. 54. Springer, Berlin.

Drela, M., 2010. Tasopt 2.00–transport aircraft system optimization. MIT 3.

Drela, M., 2011a. Development of the D8 Transport Configuration.

Drela, M., 2011b. Simultaneous Optimization of the Airframe, Powerplant, and Operation of Transport Aircraft. Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology.

Du, X., Leifsson, L., Koziel, S., Bekasiewicz, A., 2017. Airfoil design under uncertainty using non-intrusive polynomial chaos theory and utility functions, International Conference on Computational Science, ICCS 2017, Zurich, Switzerland, 12–14 June 2017. Procedia Computer Science 108, 1493–1499.

Forrester, A.I., Keane, A.J., 2009. Recent advances in surrogate-based optimization. Progress in Aerospace Sciences 45 (1), 50–79.

Garud, S.S., Karimi, I.A., Kraft, M., 2017. Design of computer experiments: a review. Computers & Chemical Engineering 106, 71–95.

Gray, J.S., Hwang, J.T., Martins, J.R.R.A., Moore, K.T., Naylor, B.A., 2019. OpenMDAO: an open-source framework for multidisciplinary design, analysis, and optimization. Structural and Multidisciplinary Optimization 59, 1075–1104.

Hall, Cesare A., Crichton, Daniel, 2005. Engine and instalations configuration for a silent aircraft. In: 17th ISABE Conference.

Harrison, K.R., Engelbrecht, A.P., Ombuki-Berman, B.M., 2018. Self-adaptive particle swarm optimization: a review and analysis of convergence. Swarm Intelligence 12 (3), 187–226.

Jacome, L.B., Elham, A., 2017. Wing aerostructural optimization under uncertain aircraft range and payload weight. Journal of Aircraft 54 (3), 1109–1120.

Jaeger, L., Gogu, C., Segonds, S., Bes, C., 2013. Aircraft multidisciplinary design optimization under both model and design variables uncertainty. Journal of Aircraft 50 (2), 528–538.

Jayabarathi, T., Raghunathan, T., Gandomi, A., 2018. The bat algorithm, variants and some practical engineering applications: a review. In: Nature-Inspired Algorithms and Applied Optimization. Springer, pp. 313–330.

Jones, D.R., Schonlau, M., Welch, W.J., 1998. Efficient global optimization of expensive black-box functions. Journal of Global Optimization 13 (4), 455–492.

Kamenik, J., Voutchkov, I., Toal, D.J.J., Keane, A.J., Högner, L., Meyer, M., Bates, R., 2018. Robust turbine blade optimization in the face of real geometric variations. Journal of Propulsion and Power 34 (6), 1479–1493.

Kennedy, J., Eberhart, R., 1995. Particle swarm optimization. In: Proceedings of the 1995 IEEE International Conference on Neural Networks, pp. 1942–1948.

Kleijnen, J.P., 2017. Regression and kriging metamodels with their experimental designs in simulation: a review. European Journal of Operational Research 256 (1), 1–16.

Li, X., Tang, K., Omidvar, M., Yang, Z., Qin, K., 2013. Benchmark functions for the CEC'2013 special session and competition on large-scale global optimization. Gene 7 (33).

Liang, Y., Cheng, X.Q., Li, Z.N., Xiang, J.W., 2011. Robust multi-objective wing design optimization via CFD approximation model. Engineering Applications of Computational Fluid Mechanics 5 (2), 286–300.

Liu, H., Ong, Y.S., Cai, J., 2018. A survey of adaptive sampling for global metamodeling in support of simulation-based complex engineering design. Structural and Multidisciplinary Optimization 57 (1), 393–416.

Mader, C.A., Martins, J.R.R.A., 2012. Derivatives for time-spectral computational fluid dynamics using an automatic differentiation adjoint. AIAA Journal 50 (12), 2809–2819.

Marelli, S., Sudret, B., 2014. UQLab: a framework for uncertainty quantification in MATLAB. In: Proceedings of the 2nd International Conference on Vulnerability, Risk Analysis and Management (ICVRAM2014). Liverpool, United Kingdom.

Martins, J.R.R.A., Sturdza, P., Alonso, J.J., 2003. The complex-step derivative approximation. ACM Transactions on Mathematical Software 29 (3), 245–262.

Michalewicz, Z., 1995. Genetic algorithms numerical optimization and constraints. In: Proceedings of the 6th International Conference on Genetic Algorithms, pp. 151–158.

Molina-Cristóbal, A., Nunez, M., Guenov, M.D., Laudan, T., Druot, T., 2014. Black-box model epistemic uncertainty at early design stage. An aircraft power-plant integration study. In: 29th Congress of the International Council of the Aeronautical Sciences.

Ng, L.W.T., Willcox, K.E., 2016. Monte Carlo information-reuse approach to aircraft conceptual design optimization under uncertainty. Journal of Aircraft 53 (2), 427–438.

Nocedal, J., Wright, S.J., 2006. Numerical Optimization. Springer.
Poli, R., Kennedy, J., Blackwell, T., 2007. Particle swarm optimization. Swarm Intelligence 1 (1), 33–57.
Powell, M.J.D., 1994. A direct search optimization method that models the objective and constraint functions by linear interpolation. In: Gomez, S., Hennart, J-P. (Eds.), Advances in Optimization and Numerical Analysis. Kluwer Academic, Dordrecht, pp. 51–67.
Quagliarella, D., Serani, A., Diez, M., Pisaroni, M., Leyland, P., Montagliani, L., Iemma, U., Gaul, N.J., Shin, J., Wunsch, D., Hirsch, C., Choi, K., Stern, F., 2019. Benchmarking Uncertainty Quantification Methods Using the NACA 2412 Airfoil with Geometrical and Operational Uncertainties.
Rumpfkeil, M., 2012. Optimization Under Uncertainty Using Gradients, Hessians, and Surrogate Models.
Rumpfkeil, M.P., Hanazaki, K., Beran, P.S., 2017. Construction of Multi-Fidelity Locally Optimized Surrogate Models for Uncertainty Quantification.
Sahoo, S., Zhao, X., Kyprianidis, K.G., Kalfas, A.I., 2019. Performance assessment of an integrated parallel hybrid-electric propulsion system aircraft. Paper No. GT2019-91459. American Society of Mechanical Engineers.
Shan, S., Wang, G.G., 2010. Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions. Structural and Multidisciplinary Optimization 41 (2), 219–241.
Skinner, S., Zare-Behtash, H., 2018. State–of–the–art in aerodynamic shape optimisation methods. Applied Soft Computing 62, 933–962.
Smola, A.J., Schölkopf, B., 2004. A tutorial on support vector regression. Statistics and Computing 14 (3), 199–222.
Storn, R., Price, K., 1997. Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. Journal of Global Optimization 11 (4), 341–359.
Tilahun, S.L., Ngnotchouye, J.M.T., Hamadneh, N.N., 2019. Continuous versions of firefly algorithm: a review. Artificial Intelligence Review 51 (3), 445–492.
Vicini, A., Quagliarella, D., 1999. Airfoil and wing design through hybrid optimization strategies. AIAA Journal 37 (5), 634–641.
Wiener, N., 1938. The homogeneous chaos. American Journal of Mathematics 60 (4), 897–936.
Wu, X., Zhang, W., Song, S., 2018. Robust aerodynamic shape design based on an adaptive stochastic optimization framework. Structural and Multidisciplinary Optimization 57 (2), 639–651.
Xiu, D., Karniadakis, G.E., 2002. The Wiener–Askey polynomial chaos for stochastic differential equations. SIAM Journal on Scientific Computing 24 (2), 619–644.
Yang, X.S., 2009. Firefly algorithms for multimodal optimization. Lecture Notes in Computer Science 5792, 169–178.
Yang, X.S., Deb, S., 2010. Engineering optimization by cuckoo search. International Journal of Mathematical Modelling and Numerical Optimisation 1 (4).
Yutko, B.M., Titchener, N., Courtin, C., Lieu, M., Wirsing, L., Tylko, J., Chambers, J.T., Roberts, T.W., Church, C.S., 2017. Conceptual Design of a D8 Commercial Aircraft.