

Review Article

State-of-the-art in aerodynamic shape optimisation methods



S.N. Skinner*, H. Zare-Behtash

University of Glasgow, School of Engineering, Glasgow G12 8QQ, UK

ARTICLE INFO

Article history:

Received 18 July 2016

Received in revised form

13 September 2017

Accepted 16 September 2017

Available online 29 September 2017

Keywords:

Optimisation

Gradient-based optimisation

Gradient-free optimisation

Aerodynamics

Hybrid optimisation

Assisted optimisation

ABSTRACT

Aerodynamic optimisation has become an indispensable component for any aerodynamic design over the past 60 years, with applications to aircraft, cars, trains, bridges, wind turbines, internal pipe flows, and cavities, among others, and is thus relevant in many facets of technology. With advancements in computational power, automated design optimisation procedures have become more competent, however, there is an ambiguity and bias throughout the literature with regards to relative performance of optimisation architectures and employed algorithms. This paper provides a well-balanced critical review of the dominant optimisation approaches that have been integrated with aerodynamic theory for the purpose of shape optimisation. A total of 229 papers, published in more than 120 journals and conference proceedings, have been classified into 6 different optimisation algorithm approaches. The material cited includes some of the most well-established authors and publications in the field of aerodynamic optimisation. This paper aims to eliminate bias toward certain algorithms by analysing the limitations, drawbacks, and the benefits of the most utilised optimisation approaches. This review provides comprehensive but straightforward insight for non-specialists and reference detailing the current state for specialist practitioners.

© 2017 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Contents

1. Introduction	934
2. General considerations of aerodynamically orientated optimisation	934
2.1. Basic problem formulation	934
2.2. Design variables and geometric representation	934
2.3. Constraint handling	937
2.4. Problem discretisation	938
3. Application of gradient-based methods to aerodynamic optimisations	939
4. Application of gradient-free methods to aerodynamic optimisations	941
4.1. Genetic algorithms	942
4.2. Particle swarm optimisation	946
4.3. Simulated annealing	949
5. Application of hybridised algorithms to aerodynamic optimisations	950
6. Application of surrogate modelling to aerodynamic optimisation	954
7. Conclusion	957
Acknowledgements	958
References	958

Abbreviations: ACARE, Advisory Council for Aeronautics Research in Europe; ADODG, Aerodynamic Design Optimisation Discussion Group; ALPSO, Augmented Lagrange Particle Swarm Optimisation; ARMOGA, adaptive range multi-objective genetic algorithm; CFD, computational fluid dynamics; CRM, common research model; EI, expected improvement; ES, evolutionary strategy; FFD, free-form deformation; GBM, gradient-based method; GDEA, Genetic Diversity Evolutionary Algorithm; GA, genetic algorithm; HM, hybrid method; LHS, latin hypercube sampling; MDO, multi-disciplinary design optimisation; MIGA, multi-island genetic algorithm; MO, multi-objective; MS, multi-start; NN, neural network; NSGA, non-dominated sorting-based genetic algorithm; NURBS, non-uniform rational basis spline; PDE, partial differential equation; PSO, particle swarm optimisation; RANS, Reynolds-averaged Navier-Stokes; RBF, radial basis function; SA, simulated annealing; SIMPSA, simplex simulated annealing; SLSQP, sequential least squares programming; SNOPT, sparse nonlinear optimiser; SQP, sequential programming; sGA, population structured genetic algorithm; μ GA, micro-genetic algorithm.

* Corresponding author.

E-mail address: s.skinner.1@research.gla.ac.uk (S.N. Skinner).

<https://doi.org/10.1016/j.asoc.2017.09.030>

1568-4946/© 2017 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Aerodynamic shape optimisation has become an indispensable component for any effective and robust aerodynamic design. Between now and 2030, there will be an estimated global demand for approximately 27,000 new passenger aircraft potentially worth up to £ 2.3 trillion. These aircraft must comply with strategic research agenda developed by the Advisory Council for Aeronautics Research in Europe (ACARE) which aim to enforce strict emission targets – CO_2 emissions per passenger kilometre to be reduced by 75%, NO_x emissions by 90% and perceived noise by 65%, all relative to the year 2000 [1]. This can only be achieved by marrying novel light weight and flexible materials, e.g. composites, with a highly optimised aerodynamic aircraft configurations; the airframe contribution should be in the order of 20–25% for fuel consumption reduction. The impact of aviation on the environment is now a main driving factor affecting the aerodynamics of future aircraft [2].

Optimisation is the process of obtaining the most suitable solution to a given problem, while for a specific problem only a single solution may exist, and for other problems there may exist multiple potential solutions. Thus, optimisation is the process of finding the ‘best’ solution, where ‘best’ implies that the solution is not the exact solution but is sufficiently superior [3]. Most current design optimisation approaches are heavily dependant on user training and experience requiring an array of specialised optimisation tools and compact shape parametrisation. This constitutes a major obstacle to robustness and reliability [4]. Another persistent difficulty in aerodynamic optimisation is the ability to define an analysis method that is capable of operating as many time as required (often thousands of times) and integrating it appropriately with an optimisation strategy. The methods employed must execute with realistic run times, dependant on computational resource, but must also be sophisticated enough to capture enough information to analyse local geometry that feeds into a globally optimal system [5].

Many optimisation problems, especially those involved with large design spaces with coupled variables, inherently fall into the category of multi-disciplinary design optimisation (MDO), which in turn require a multi-objective (MO) compromise to be an effective design. The main motivation for applying MDO is that the performance of a real system is driven not only by the performance of individual disciplines but also by their coupled interactions. It is no longer acceptable to consider the aerodynamic analysis alone, its far reaching coupled effects to other disciplines must also be taken into account if a truly optimal design is to be reached. Studies, such as that by Werter and Breuker [6], have shown that aeroelastic tailoring through aerostructural optimisation offer clear performance advantages. A survey of developments in multi-disciplinary optimisation for aerospace applications is presented by Sobieszczanski-Sobieski and Haftka [7]. Furthermore, Wunderlich [8] provide a summary of multi-disciplinary considerations.

Forums for comparing aerodynamic optimisation, such as the AIAA Drag Prediction Workshop [9,10], are well established and enable different research groups (such as the MDO laboratory led by Professor Martins at the University of Michigan) to validate their codes utilising common research models (CRM). Additionally, researchers in the aerodynamic design optimisation community have developed the AIAA Aerodynamic Design Optimisation Discussion Group (ADODG). The ADODG define benchmark optimisation problems enabling research groups to test and compare optimisation codes for various problems; a broad scope of flow conditions and geometries are considered. Several studies utilising common research benchmarks are suggested [11–15]. Furthermore, new comers into the field of optimisation are directed to Rios and Sahinidis [16], and Amaran et al. [17] offering comprehensive

reviews of accessible generalised optimisation algorithms and their applications examining and contrasting relative performance.

This study reviewed a total of 304 papers considering more than 120 peer reviewed journals and several well-established conference proceedings; from which 229 papers have been selected to provide a comprehensive and well-balanced review of optimisation algorithms that are dominant in the field of aerodynamic optimisation. This review begins with a brief examination of general considerations required for aerodynamic optimisations. The fundamental operation, to-date research application, and challenges encountered relating to the optimisation strategies employed are analysed. The review discusses studies from several areas for aerodynamic design optimisation also covering associated multi-disciplinary analysis.

2. General considerations of aerodynamically orientated optimisation

2.1. Basic problem formulation

The field of optimisation is expansive, and the choice of a suitable algorithm is highly problem dependant [18]. A general optimisation problem can be presented mathematically as [19]:

$$\begin{aligned} & \text{Minimise } \left\{ \begin{array}{l} F(\mathbf{X}) \quad \text{Objective function} \\ \text{With } < \text{respect } < \text{to } \left\{ \begin{array}{l} \mathbf{X} \quad \text{Design } < \text{variables} \end{array} \right. \end{array} \right. \\ & \text{Subject } < \text{to } \left\{ \begin{array}{l} g_i(\mathbf{X}) \leq 0 \quad i = 1, l \quad \text{Inequality } < \text{constraints} \\ h_j(\mathbf{X}) = 0 \quad j = 1, l \quad \text{Equality } < \text{constraints} \\ X_k^l \leq X_k \leq X_k^u \quad k = 1, N_{DV} \quad \text{Parameterised } < \text{constraints} \end{array} \right. \end{aligned}$$

$$\text{where } < \mathbf{X} = \left\{ \begin{array}{l} X_1 \\ X_2 \\ \vdots \\ X_{N_{DV}} \end{array} \right\}$$

Most optimisation methods use an iterative procedure. The initial set \mathbf{X} design variables, which in the context of aerodynamic optimisation this is referred to as the baseline configuration, and is updated until a minimum of $F(\mathbf{X})$ is identified or the optimisation process runs out of allocated time/iterations. In the initial set-up of the optimisation problem consideration must be given to: 1) the level of information fidelity required from the flow solver, dependant on the type of problem; 2) scope of parameterised design space; 3) types of design variables, e.g. discrete and/or continuous; 4) single or multi-objective optimisation; 5) constraints handling; 6) properties of the design space, e.g. number of local optima, discontinuities. It is important to note that no optimisation procedure guarantees the global optima of the objective function $F(\mathbf{X})$ will be found: the process may only converge towards a locally optimal solution. Typically in this situation there are three possibilities: 1) restart the optimisation process to investigate if the same solution is found; 2) approach the design problem with a different optimisation methodology to compare solution quality at a high computational expense; or 3) accept the optimum found knowing that while it is superior to the baseline configuration it may not be the optimal solution.

2.2. Design variables and geometric representation

In general, it is necessary to implement geometric parametrisations in such a way that reduce the complexity and cost of the optimisation process but do not restrict communication

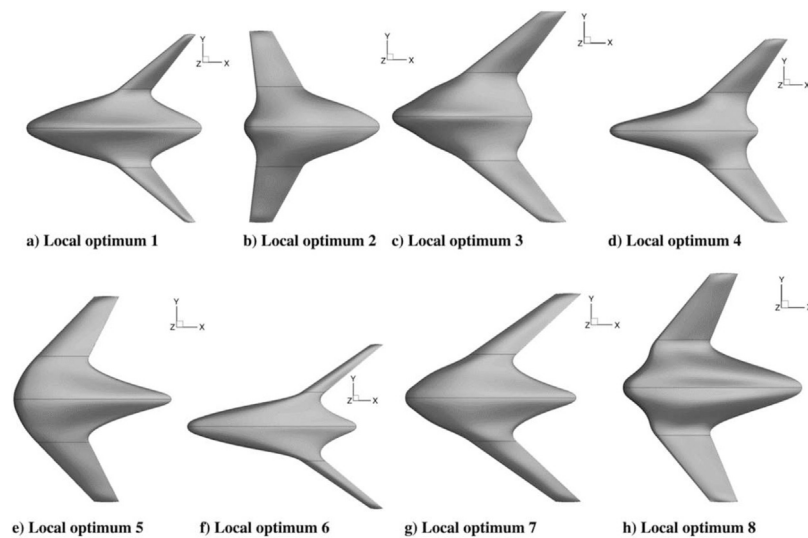


Fig. 1. Local optima found in a blended wing optimization [23].

of variables or the degree to which aerodynamic performance is optimised. Parametrisation aims to balance the fundamental compromise between computational speed of the optimisation run-time, favouring a tight parametrisation. Zhang et al. [20] show that the defined dimensionality of a problem for shape optimisations can restrict the optimal design. Using too few variables may prove certain potential improvements impossible. Conversely, if too many design variables are used, particularly if variables are strongly coupled, the search landscape can become intractably complex to navigate. Increasing the dimensionality of a given problem excessively leads to a paradox, first addressed by Sobieszczanski-Sobieski [21,22], in which increasing the number of design variables leads to a decrease in the number of variables that can be manipulated as a direct result of increased coupling. It is often desirable to limit the allowable design variables to avoid geometries that cannot be evaluated with sufficient accuracy by the flow solver: due to meshing limitations for example. Furthermore, this can help to avoid geometries that are unacceptable in terms of some criteria, or similarly, restrict the optimisation to geometries that are necessary for other criteria. Regardless of the user defined parametrisation, the final design is most definitely suboptimal – often limited by parametrisation [4].

Chernukin and Zingg [23] conducted one of the few studies on how the number of variables used, and the related modality, can affect aerodynamic designs and highlight that distinguishing between multi-modality and poor optimiser convergence can prove problematic. By increasing the dimensionality of a design space it can be expected, but not guaranteed, to increase the modality of the search space. Initialising 224 random starting geometrics with 368 design variables they demonstrated the presence of at least 8 local optima for a blended wing optimisation, shown in Fig. 1. All 8 solutions satisfied both optimality and feasibility tolerances with the objective value varying by approximately 5% between the local optima. The planform shapes are distinct and so demonstrate that geometric variation is significant between local optima which share similar performance characteristics.

Furthermore, the method of geometric parametrisation used to communicate a set of variables plays an important role in identifying optimal aerodynamics. It determines what shapes and topologies can be represented, and how many design variables are necessary for sufficient representation of the geometry. Thus, parametrisation dictates particular geometric requirements and has a strong influence on the design landscape. Therefore it can-

not be precluded that different geometric parametrisations will increase or decrease the degree of modality, linearity, or discontinuity observed. Additionally, a complex geometry parametrisation may impose distinct computational costs. Representations of a geometry can be broken down into a number of categories but in a more broad sense they can be considered to be constructive, deformative, or volume based.

Constructive models include functions which define basic body shapes, spline methods (such as Bezier splines, basis splines (B-splines), non-uniform rational basis spline (NURBS)) and partial differential equations. Jansen et al. [24] used a medium-fidelity aerostructural panel code to perform optimisation of a conceptual wing configurations, shown in Fig. 2. The basic wing topology was defined through a series of globally enforced geometric variables to manipulate a series of wing sections. Parametrising the entire geometry in this way typically allows for global shape control with few basic variables. This method is well suited to low-fidelity aerodynamic models if a wide allowable design scope is necessary – no need for mesh deformations.

Spline-based geometric parametrisations are used to represent two- or three- dimensional surfaces and are typically used in conjunction with higher-fidelity flow solvers, such as Euler and Navier Stokes solvers, with the control points being the design variables. Bezier splines are most efficient to evaluate requiring few variables and have been used for efficient aerofoil definition by Peigin and Epstein [25]. Modification of any single control point defining a Bezier spline will modify the entire curve and thus is inherently effective for global shape definition, but has very limited local control. B-splines address this issue of local control allowing single control point modifications to modify small portions of the overall curve. This allows for more complex aerofoil definitions, as demonstrated by Koziel et al. [26], and can enable the use of hinged control surfaces to an otherwise rigid body. NURBS increase the local deformation control over surface definitions further in order to have more complex geometric shapes such as fairings or wing-fuselage junctions. Vecchia and Nicolosi [27] and Hashimoto et al. [28] adopt NURBS to parametrise the entire aircraft configuration in order to reduce drag of the vehicle through steam-lining fillets and fairings. Fig. 3 shows an example of NURBS control points re-defining the surface over the upper section of the fuselage/wing juncture.

Geometry definition through the use of partial differential equations (PDEs) are not as commonly used as well-established spline-based methods but are just as versatile for geometry surface




Fig. 2. Global variables parametrising wing topology [171].

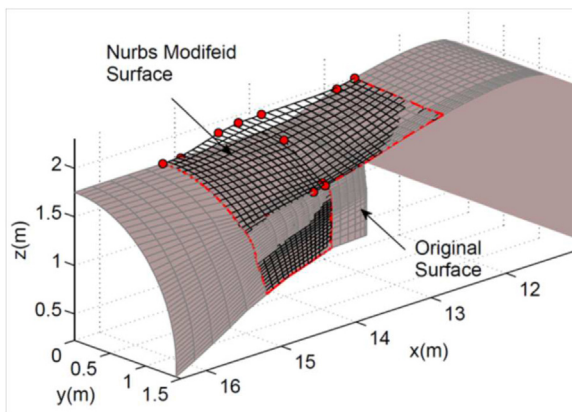


Fig. 3. NURBS surfaces parametrising surface blend on fuselage [27].

definition. Athanasopoulos et al. [29] show that for equivalently complex surface construction PDEs require fewer design variables, resulting in a more compact design space. Due to the small set of design parameters required by the PDE method the computational cost associated with the optimisation of a given aerodynamic surfaces can be reduced [30]. In a PDE-based method the parameters are boundary values to the PDE, hence the relationship between the value of the design parameter and the geometry can be unclear making methodical surface deformations tedious. This is likely why the aerodynamic definition of a body in an optimisation scheme does not use PDE representation even though it may initially seem a more appropriate method. Comparatively, spline-based methods are conceptually simpler and will provide a more direct relationship between design parameters and the resulting geometry and thus allow better control over the range of geometries that can be generated.

If optimisation establishes performance metrics from computational fluid dynamics (CFD), the simplest methods for body surface definitions are deformative ones. In deformative methods the mesh points on the surface of the body are directly treated as design

variables [31], and their position can be perturbed by the optimiser in order to generate new shapes. These approaches have the significant advantage that any geometry the mesh generation algorithm is capable of can be evaluated, however it is likely to require many hundreds of design variables; deformations are therefore usually limited to single-degree-of-freedom deformations. A common method used for aerodynamic optimisation is the free-form deformation (FFD) approach which is useful if the geometry manipulations are particularly complex; FFD is covered in depth by Kenway and Martins [32]. This approach embeds the solid geometry within a FFD hull volume (volumes are typically trivariate analogues of Bezier splines, B-splines or NURBS), which are parametrised by a series of control points as shown in Fig. 4. These control points deform the volume which translate to geometric changes of the solid geometry rather than redefining the whole geometry itself which can give a relatively more efficient set of design variables. A key assertion of the FFD approach, when applied within a CFD environment, is that a geometry has constant topology throughout the optimisation process; [12] this is typical of high-fidelity optimisations where the initial geometry considered is sufficiently close to the optimal solution. Fig. 4 shows the FFD hull volume enclosing a wing with 720 geometric control points used by Lyu et al. [12] which control shape deformation in the vertical (z) axis. The initial random wing deformation and associated optimised wing cross-sections at select locations are also shown. A similar method is based on radial basis function (RBF) interpolation which defines data sets of design variables and their global relationships. Fincham and Friswell [33] use radial basis functions to optimise morphing aerofoils and report that they provide a means to deform both aerodynamic and structural meshes and interpolate performance metrics between two non-coincident meshes.

Volumetric-based body representation have been used for optimisation but rarely in the field of aerodynamics, a recent review of the applicability of volumetric parametrisation for aerodynamic optimisation is given by Hall et al. [34]. The authors point out the limitations of volumetric representations stating that black-box optimisers cannot be used and even gradient-based methods can often be impractical.




Fig. 4. Free-form deformation parametrising wing with 720 control points. Selected wing cross-sections are highlighted to show initial perturbed (red), and final optimised (blue) wing cross-sections. Associated C_p distributions shown [12]. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

2.3. Constraint handling

Constraint handling in aerodynamic, and indeed any industrial optimisation problem, plays a consequential role in the quality and robustness of an optimised solution within the defined design space. Geometric parametrisation itself poses a constrained optimisation problem since, in addition to minimising the objective $F(\mathbf{X})$, the design variables must satisfy some geometric constraints. Constraint management techniques found in literature which have been classified by Koziel and Michalewicz [35] and Siem and Innocente [36] as: 1) strategies that preserve only feasible solutions with no constraint violations: infeasible solutions are deleted; 2) strategies that allow feasible and infeasible solutions to co-exist in a population, however penalty functions penalise the infeasible solutions (constraint based reasoning); 3) strategies that create feasible solutions only; 4) strategies that artificially modify solutions to boundary constraints if boundaries are exceeded; and 5) strategies that repair/modify infeasible solutions.

Most commonly optimisations apply weighted penalties to the objective function if the constraint(s) are violated. The reason for this is that penalty functions are often deemed to ease the optimisation process, and bring the advantage of transforming constrained problems into unconstrained one by directly enforcing the penalties directly to the objective function. With this method Pareto-optimal solutions with good diversity and reliable convergence for many algorithms can be obtained easily when the number of constraints are small; fewer than 20 constraints. It becomes more difficult to reach Pareto-optimal solutions efficiently as the number of constraints increase, and the number of analyses of objectives and constraints quickly becomes prohibitively expensive for many applications. This is because the selection pressure decreases due to the reduced region in which feasible solutions exist [37].

Kato et al. [38] suggest that in certain circumstances Pareto-optimal solutions may exist in-between regions of solution feasibility and infeasibility. This is illustrated in Fig. 5, where it is seen that feasible and infeasible solutions could be evaluated in parallel to guide the optimisation search direction towards feasible design spaces. This is intuitively true for single discipline aerodynamic optimisation problems where often small modifications to design variables can largely impact the performance rendering designs infeasible. Algorithm understanding of infeasible solutions can help in the betterment of feasible solutions though algorithm learning/training and constraint based reasoning. Robinson et al. [39], comparing the performance of alternative trust-region con-

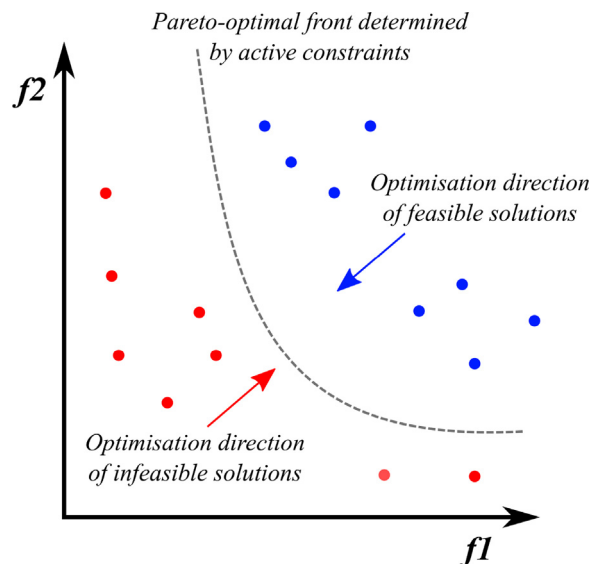


Fig. 5. Concept of using parallel evaluation strategy of feasible and infeasible solutions to guide optimisation direction in a GA.

straint handling methods, showed that reapplying knowledge of constraint information to a variable complexity wing design optimisation problem reduced high-fidelity function calls by 58% and additionally compare the performance to alternative constraint managed techniques

Elsewhere, Gemma and Mastroddi [40] demonstrated that for the multi-disciplinary, multi-objective aircraft optimisations the objective space of feasible and infeasible design candidates are likely to share no such definitive boundary as shown in Fig. 6. With the adoption of flutter constraints, structural constraints, and mission constraints solutions defined as infeasible under certain conditions would otherwise be accepted, hence forming complex Pareto fronts. Interdisciplinary considerations such as this help to develop and balance conflicting constraints. For example, structural properties which may be considered feasible, but are perhaps heavier than necessary will inflict aeroelastic instabilities at lower frequencies.

In the aerospace industry alone there are several devoted open-source aerodynamic optimisation algorithms with built-in constraint handling capability. For example COBYLA [41–43], DIRECT [44–47], NOMAD [48], and HAVOC [49,50] each offer

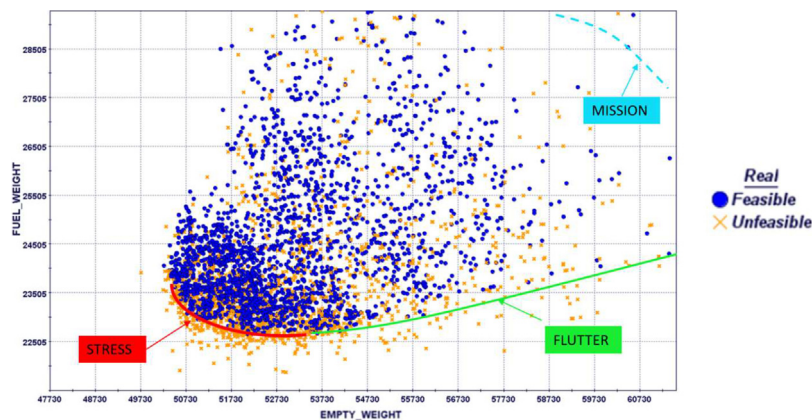


Fig. 6. Overlapping boundaries for feasible and infeasible solutions objective space [40].

derivative-free optimisation algorithms capable of handling constraints explicitly; each adaptable to the users aerodynamic solver whether commercial or in-house. Some studies [51–53] have also adopted MATLAB's optimisation tool-box for successful optimisation constraint management.

2.4. Problem discretisation

Problem discretisation is achieved through surface and volume meshing/panelling of the fluid domain in order to formulate a discretised representation of the computational domain for which applied physics solvers provide a numerical solution too.

A persistent challenge in the application of numerical simulations (whether high- or low-fidelity) is the accurate evaluation of the optimisation objective function which is strongly dependant upon the body/domain discretisation. Given that multiple iterations are required within an individual optimisation process, it is necessary to maintain a balance between efficiency and accuracy; particularly if one employs multi-start or multi-point optimisation. For this reason, many aerodynamic shape CFD-based optimisations use a fixed mesh constructed around some near optimal baseline geometry. Zhang et al. [54] perform aeroelastic optimisation with flutter constraints, which they avoid convergence issues associated with dynamic mesh deformation by employing a fixed over-set mesh strategy. However, the task of creating a mesh *a priori* which remains appropriate for the whole design search space is often infeasible. Certain parameters, such as induced drag, are particularly sensitive to the grid size, hence aerodynamic properties of a body must be shown to be grid independent for the entirety of the design space for the mesh employed. If the aerodynamic properties of a body are not grid independent it is very easy for the mesh to become inappropriate as geometry and/or flow conditions change during the optimisation process. This can lead to ill-defined optimisation information and render the optimisation process useless as mesh discretisation errors are exploited leading to false designs [55].

Developments in robust CFD mesh adaptation/refinement and deformation have helped make optimisation procedures more reliable and have made it possible to expand the scope of design parameters [56,57]. This enabling a more flexible design space. Adaptive meshes help in the progressive optimisation of a design, particularly when compressible aerodynamics is considered. Li and Hartmann [58] compare aerofoil optimisations for a fixed mesh and an adaptively refined mesh. They demonstrate that the adaptive meshing, reducing discretisation error, was not only computationally cheaper but capable of finding better solutions. The fixed mesh was unable to sufficiently evaluate shock discontinuities at the surface of the body. Nemic and Aftosmis [55], in the optimisation

of a supersonic inverse-design problem, also show that progressive mesh adaptation reduces the required computational effort during early design iterations and improves objective function convergence. Adaptive meshing can however introduce issues with optimiser convergence if poorly implemented. These issues may include unnecessary design evaluations for the course mesh, but also fine mesh iterations to undo/reverse false progress [55].

An alternative to adaptive meshing is used by Lyu et al. [12] called multi-level optimisation. This involves the systematic increase in mesh fidelity as certain stages of optimality is achieved; the mesh stages used by Lyu et al. [12] are shown in Fig. 7. These mesh stages would be developed before the optimisation and would allow the course grid to do much of the optimisation reducing the overall computational demand. When moving to the next mesh level for higher-fidelity optimisation, care must be taken as the updated mesh will compute different performance metrics from the previous mesh.

For aeroelastic problems in which the geometry is expected to structurally deform, or the definition of a perturbed volume mesh is needed, mesh deformations must be considered. Mesh perturbing algorithms are common practice however special attention must be given to how node translations and rotations effect the orthogonality of boundary layer mesh; this will drastically change results obtained and may give rise to a situations of false optimisation. Chiba et al. [59] highlight that independent modification to the surface and volume meshes, which are dictated by changes in the surface geometry, can result in surface mesh distortion. Fig. 8 shows an example of surface mesh distortion around the leading edge of the wing. The authors claim this to be a primary reason for restrictive optimisation and therefore concluded that the mesh distortions highly limit the reachable design space consequently resulting in either sub-optimal or pseudo-optimal results. Elsewhere mesh deformations have been employed successfully for high-fidelity aerodynamic [60–62] and aerostructural design [13,63], morphing wing structures [64], and ducted propeller optimisation [65].

For reduced, lower-fidelity models, such as vortex lattice methods the discretisation of the surface is much simpler, however if the body geometry is allowed to vary dynamically throughout the optimisation the spatial resolution of the panelling should be adjusted accordingly to maintain consistent numerical accuracy. Ning and Kroo [66] achieve this by discretising the wing span length by the panel size. This forced the dimensions of the wing to vary discontinuously. Jansen et al. [24] allowed the span to vary continuously by dividing the wing into independent segments and then assessing their contribution to the overall continuous span. Similar methods were used by Skinner and Zare-Behtash. [67]. Methods used by Vecchia and Nicolosi [27] fix the panels over the body and so




Fig. 7. Multi-level optimisation stages for the systematic increase of solver fidelity as different stages of solution optimality is achieved [12].




Fig. 8. Example of mesh distortion at the wing leading edge [59].

do not allow significant geometric span changes; shed wake panels through the aerodynamic analysis also have a fixed size.

3. Application of gradient-based methods to aerodynamic optimisations

Gradient-based optimisation is a calculus-based point-by-point technique that relies on the gradient (derivative) information of the objective function with respect to a number of independent variables. The nature in which gradient-based methods (GBM) operate make them well suited to finding locally optimal solutions but may struggle to find the global optimal [68]. With gradient-based algorithms an understanding of the design space is assumed, as an appropriately pre-conceived starting design point must be given. Kenway and Martins [13] point out that with increasingly higher fidelity aerodynamic optimisations, a more refined initial design should be used so that the optimisation does not diverge too far from the baseline. If large changes in topology are expected lower fidelity panel codes, such as that described by Vecchia and Nicolosi [27], can facilitate useful optimisation procedures. Typically, the higher the fidelity analysis used the more compact the design variables will need to be to allow effective optimisation with a gradient based optimiser.

Gradient-based optimisation is, in its most basic form, a two step iterative process which can be summarised mathematically as:

$$\mathbf{X}^{new} = \mathbf{X}^{old} + h\nabla f \quad (1)$$

where ∇f is the gradient of function $F(\mathbf{X})$, and \mathbf{X} is a vector of the design variables. The first step is to identify a search direction (gradient), ∇f , in which to move. The second step is to perform a one-dimensional line search to determine a distance/step size h along ∇f that achieves an adequate reduction of some cost function, i.e. define how far to move in the search direction until no more progress can be made [69]. A schematic diagram illustrating

the operation of a gradient-based optimisation is shown in Fig. 9. In-depth benchmarking of gradient based algorithms for aerodynamic problems has been conducted by Secanell and Suleman [70] and Lyu et al. [71].

Gradient based algorithms are extensively used in aerospace optimisation as they exhibit low computational demands when handling many hundreds of design variables – this makes them well suited for optimising shapes based on deformative geometric parametrisations [12,64]. Significant difficulties arise if they are not applied within a restricted set of functions with well defined slope values due to a dependency upon the existence of derivative information via some sensitivity analysis. There are several different methods for sensitivity analysis for which four general classes can be distinguished: 1) finite-difference methods; 2) complex-step derivative approximation; 3) automatic/algorithmic differentiation; and 4) analytic methods. It is important to understand their relative merits since none are a clear choice for all classes of problem. Comparative studies on the numerical sensitivity analysis for aerodynamic optimisation has been conducted by Martins et al. [72] and Peter and Dwight [73], while Martins and Hwang [74] offer a detailed discussion for computing derivatives within multi-disciplinary computational models.

The computational expense of evaluating gradients using finite-difference or the complex-step method provide a simple and flexible means of estimating gradient information, but are considered excessive with respect to hundreds of variables [75]. These approaches preserve discipline feasibility, but they are costly and can be unreliable. Finite-differencing, while not used to provide gradients for the optimisation itself have been used by Kenway and Martins [13] and Skinner and Zare-Behtash [67] to provide gradients for stability derivative constraints. For a restricted number of design variables the complex-step method is suitable for sensitivity analysis as demonstrated by Kenway and Martins [76]. They employ the complex step method to provided gradients for the

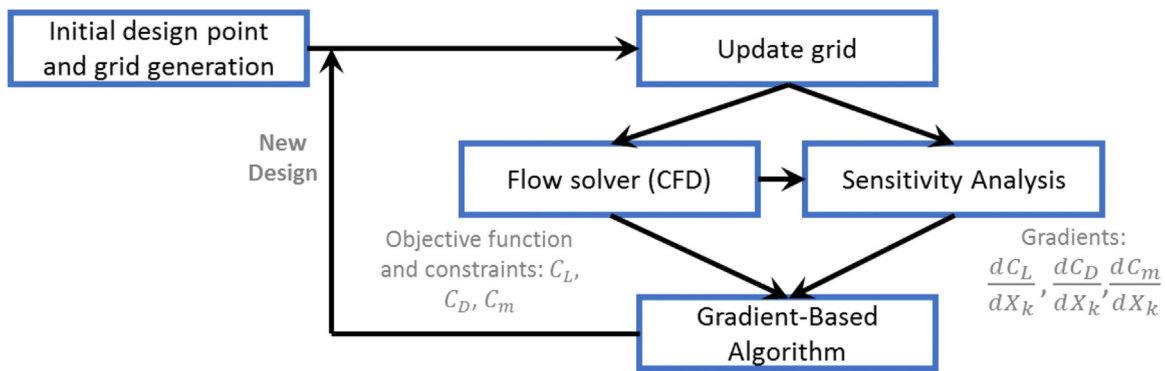


Fig. 9. Schematic diagram of a gradient-based aerodynamic optimisation process.

Sparse Nonlinear Optimiser (SNOPT) algorithm, originally developed by Gill et al. [77], in the constrained optimisation of wind turbine blades. It is commented on that to increase the dimensionality of the problem an analytic sensitivity analysis would have to be adopted.

Finite-differencing or complex-step methods employed for providing sensitivity analysis for low-fidelity codes can be considered appropriate due to low computational demand. Ning and Kroo [66] optimise a series of wing topologies investigating fundamental wing design trade-offs for which sensitivity analysis of the objective and constraints were approximated by finite-differencing. Results provided by the sequential quadratic programming method show robust and quick convergence able to determine relative gradients between approximated area-dependant weight, effects of critical structural loading, and stall speed constraints.

In the presence of several hundred design variables and constraints the analysis code will require a particularly long time to evaluate sensitivities. Automatic/algorithmic differentiation or analytic derivative calculations (direct or adjoint) can be used to avoid multi-discipline analysis evaluations. Pironneau [78], pioneered the adjoint method in fluid dynamics, showing that the cost of computing sensitivity information was almost completely independent of the number of design variables, and hence the overall cost of optimisation is roughly linearly proportional to the number of design variables. Lyu et al. [71] more recently demonstrated that both SNOPT and sequential least squares programming [79,80] (SLSQP) gradient-based algorithms with analytically derived adjoint gradients require far fewer total functions calls when compared to using finite-differencing for high-fidelity large scale aerodynamic optimisations. The adjoint form of the sensitivity information is particularly efficient for aerodynamic optimisation applications as the number of cost functions (outputs) is small, while the number of design variables (inputs) is relatively larger.

The discrete adjoint method (as opposed to continuous adjoint method) is generally favoured in aerospace-based optimisation as it ensures that sensitivities are exact with respect to the discretised objective function [81,82]. The implementation of the adjoint method for the governing equations of the flow analysis can often be difficult to derive and require direct manipulation; adjoint methods require much more involved detailed knowledge of the computational domain. One way to approach this difficulty is to use automatic/algorithmic differentiation, which is a method based on the systematic application of the differentiation chain rule to the source code to compute the partial derivatives required by the adjoint method. Mader et al. [83] developed a discrete adjoint method for Euler equations using automatic differentiation, later followed by Lyu et al. [84] who extended and developed this adjoint implementation to Reynolds-averaged Navier-Stokes

(RANS) equations and introduced simplifications to the automatic differentiation approach. Methods developed have shown robust an efficient application to high-fidelity optimisation [12,63].

Hicken and Zingg [85] adopted similar methods for the high-fidelity aerodynamic optimisation of non-planar wings addressing the non-linearity of wake shape and how it can impact the induced drag. Several non-planar geometries, inherently creating non-planar wake-wing interactions, are optimised using discrete adjoint sensitivities and the SNOPT algorithm. This work illustrates the drawbacks in static-wake assumptions, demonstrating that higher-order effects must be included for accurate induced drag prediction and hence for meaningful optimisations. This work was followed by Gagnon and Hicken [86] for the aerodynamic optimisation of un-conventional aircraft configurations; adapted optimisation results are shown in Fig. 10. Here, for the aerodynamic metrics, the gradients are evaluated analytically by using the discrete-adjoint variables while other gradients are provided by the complex-step method. This work is notable as it enables used axial deformation combined with free-form deformation to achieve both local and global geometric manipulations, the effects of this can be seen in Fig. 10, allowing span, sweep, dihedral, taper, twist and aerofoil sectional shape changes. These global geometric variables are generally not considered in high-fidelity simulation. The cost of allowing such geometric variation away from the baseline under high-fidelity optimisation limited how many variables could be considered in any one optimisation process. The authors observed limited optimisation in some wing configurations because of this.

Aeroelastic optimisation requires the coupling of aerodynamic and structural models for most effective sensitivity analysis in optimisation routines. Even small changes in aerodynamic shape can have a large influence on aerodynamic performance with various flow conditions resulting in multiple shapes. Wing flexibility impacts not only the static flying shape but also its dynamics, resulting in aeroelastic phenomenon such as flutter and aileron reversal. Based on this principle, to enable high-fidelity aerostructural optimisation while encompassing hundreds of design variables, Martins et al. [87] proposed the use of a coupled adjoint method to compute sensitivities with respect to both the aerodynamic shape and the structural sizing. Kenway et al. [13,88] subsequently made several developments and demonstrated that the computation of coupled aeroelastic gradient calculations were scalable to thousands of design variables and millions of degrees of freedom, and since applied it to the aerostructural optimisation of high aspect ratio wings with different structural properties [89]. More recently, Burdette et al. [90,64] applied the coupled discrete adjoint method with the sparse non-linear optimiser SNOPT for wing morphology optimisation. This approach was capable of handling over a thousand design variables and constraints. The coupled adjoint method is also applicable to lower-fidelity models in studies

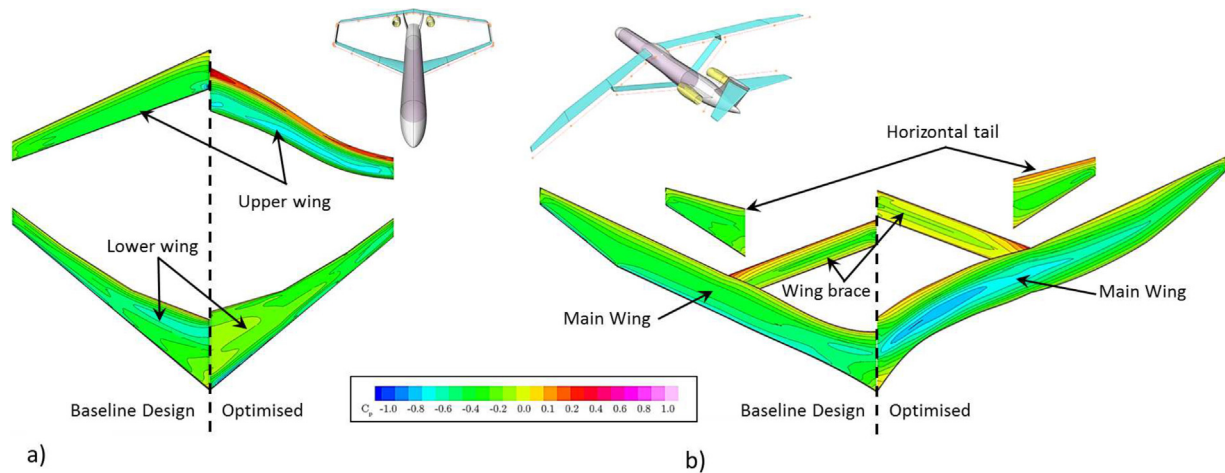


Fig. 10. Upper surface pressure coefficient contours over initial and optimised un-conventional aircraft configurations [86].

like that by Elahm and Tooren [91], where they used a vortex lattice method and finite element analysis tool capable of accurately mimicking high-fidelity accuracy at a greatly reduced computational cost.

The coupling of design constraints makes the optimiser additionally capable of considering more sophisticated criteria. Mader and Martins [92] included flight dynamics into the coupled adjoint sensitivity and explored the use of static and dynamic stability constraints. Their result showed that coupling stability constraint sensitivities into the adjoint formulation had a significant impact on optimal wing shape. Elsewhere, structural dynamics were considered by Zhang et al. [54] who used a coupled-adjoint formulation to include flutter constraints. The flutter constraints used the coupled aerodynamic/structural solver to suppress flutter onset by identifying dominant modes and adjusting variables such as the wing stiffness.

Grossman et al. [93] investigated using modular sensitivity analysis for aerostructural sequential optimisation of a sailplane. They showed that coupled aerostructural optimisation gave higher performance designs than those identified by sequential optimisation of aerodynamics followed by structural optimisation. Subsequently, Grossman et al. [94] optimised the performance of a subsonic wing configuration showing that while modular sensitivity analysis for sequential optimisation reduced the total number of function calls and sensitivity calculations, the wing performance gain was limited. When performing sequential optimisation the optimiser does not have sufficient information necessary for aeroelastic tailoring. This limitation of sequential optimisation is further explained by Chittick and Martins [95].

A significant drawback of all gradient-based algorithms is the requirement for continuity and low-modality throughout the design space otherwise the algorithm may become sub-optimally trapped. The challenge is that an aerodynamic shape analysis throughout a geometrically varying search space will encounter both non-continuous topological and local flow changes, each providing local optima [96,71,97]. Gradient-dependant algorithms' robustness significantly decreases in the presence of discontinuity and lack of convergence, usually related to turbulence modelling, making the objective function noisy [98]. Kenway [99] encountered such a problem with aerodynamic shape optimisation with a separation-based constraint formulation to mitigate buffet-onset behaviour at a series of operating conditions. The discontinuity from the 'separation sensor' function arose from monitoring the wing local surface for separated flow; this resulting in locally negative skin friction coefficients. To address this issue blending functions were to be implemented to smooth the discontinuity,

smearing the separation sensor value around the separated flow region.

Kenway and Martins [99], among several others [100,101,64,13], have used multi-point optimisation strategies in order to consider several operating conditions simultaneously. For more realistic and robust design it is crucial to take into account more than one operating condition, especially off-design conditions, which form additional multi-objective requirements into the optimisation. Fig. 11 shows the results for both single-point and multi-point optimisation of the Common Research Model configuration presented by Kenway and Martins; [99] results are shown for the nominal operating condition. The single-point optimisation achieved an 8.6 drag count reduction and the shock-wave over the upper surface of the wing is almost entirely eliminated. Drag divergence curves in this work show the nature of the single-point optimisation presenting a significant dip in the drag at the design condition, but the performance is significantly deteriorated at off-design conditions relative to the baseline condition. The multi-point optimisation, accounting for 3 design conditions, found that drag at the nominal operating condition increased by 2.8 counts and produced double shocks on the upper surface of the wing visible in Fig. 11. However, at the sacrifice of performance at the nominal operating condition, off-design conditions for the multi-point optimisation design was found to perform substantially better over the entire range of Mach numbers. Though biasing the optimisation toward certain operating conditions the authors show that multi-point optimisation with all conditions near the on-design condition is not sufficient for an overall robust design when considering operational envelopes.

4. Application of gradient-free methods to aerodynamic optimisations

The principal source of difficulty in the application of gradient-based optimisers is the requirement for having a non-discontinuous and mathematically predictable design space. Non-gradient based methods can prove more complex to implement than GBM, but they do not require continuity or predictability over the design space, and usually increase the likelihood of finding a global optimum [102].

Methods of optimisation known as metaheuristics can offer robust methods of finding a solution, and increases the likelihood of converging onto a solution at the global optimum. These gradient-free methods are known to be capable of engaging with numerically noisy optimisation problems that be difficult for GBM. This is because metaheuristics operate from a completely different




Fig. 11. High performance low drag solutions found for single design point at nominal operating conditions. For multi-point optimisation performance at the nominal operating condition is sacrificed [99].

paradigms usually based on the some naturally occurring phenomenon [103]. Unlike gradient methods, derivatives of the cost functions are not necessary, allowing metaheuristics to easily cope with non-continuous or numerically noisy cost functions. Furthermore, no pre-defined baseline design or knowledge of the design space is required and gradient-free methods typically optimise several solutions in parallel.

4.1. Genetic algorithms

Evolutionary computation is founded upon biological evolutionary theory of which there are four historical paradigms motivating activity within the field: Evolutionary Programming, Evolutionary Strategies (ES), Genetic Algorithms (GAs), and Genetic Programming [103]. The basic differences between these lie in the nature of the representation schemes, the reproduction operators, and selection methods. Focus will be placed upon genetic algorithms as they have successfully been applied to a wide-base of aerodynamic design optimisations due to their ease of use, broad applicability, and global perspective of the search domain.

GAs are a population-based optimisation technique based on the Darwinian theory of survival of the fittest: a primary aspect of evolution [104]. These algorithms are often praised for their ability to explore and exploit solutions simultaneously due to their inherent multi-start capability. They are easily capable of constructing insightful design trade-off relationships, referred to as Pareto fronts, between objectives as shown in Fig. 12. Fig. 12 shows the resulting Pareto optimal solutions found by Yamazaki et al. [11] in a winglet design problem illustrating the trade-off between pure drag (induced + wave + profile) and root bending moment.

GAs are well-suited to complex optimisation tasks as they can easily use both discrete and continuous variables, and can easily handle non-linear, non-convex, and non-continuous objective functions [51]. Chiba et al. [105] suggest that GAs have four distinct advantages which encourage their use in aerodynamic/aerostructural optimisations: 1) GAs have the ability to find multiple optimal solutions and design trade-offs; 2) GAs process information in parallel, optimising from multiple points within the design space; 3) high-fidelity CFD codes can be adapted to GAs

without any modifications; and 4) GAs are insensitive to numerical noise that may be present in the computation. The main drawbacks associated with these algorithms are high computational cost, poor constraints handling abilities, requirement for problem specific tuning and limitations in how many variables are feasible to handle. Studies have shown that GAs are very fast at identifying regions of optimality within a design space but demonstrate slow convergence as they moves nearer optimal solutions [18]. Some studies have tried to build on the classical GA to enhance its applications to aerodynamic optimisation [106].

GA optimisation is, in its most basic form, an iterative process which can be summarised as [107,103]:

1. **Random generation** of individuals to form the initial population.
2. **Evaluation** of the fitness/survivability of each individual in the population to the given environment. This would be done with the aerodynamic solver.
3. **Selection** of individuals to take part in genetic operations.
4. Apply genetic operations which mimic **reproduction** to define a new population.
5. **Iterate** over steps 2–4 are over multiple generations until some convergence criterion is met.

Fig. 13 illustrates this iterative process. Key influencers, each posing characteristic difficulties, in the construction of GAs suitable for aerodynamic optimisation problems include: 1) GA population size; 2) selection methods; 3) genetic operations (namely crossover [103,108] and mutation [103,109–111]); 4) genotype-phenotype mapping; 6) sufficient design constraints for adequate problem definition; and 7) computational resource. Insufficient selection of these factors can delay, if not prohibit, the performance of a GA in finding optimal solutions with the desired precision. Many variants of GAs exist in many innovative and abstract forms due to the array of different operators possible within the ‘selection’ and ‘reproduction’ stages.

Population size is very much related to the complexity of the problem and the number of design variables considered; understanding how the population size influences a particular problem

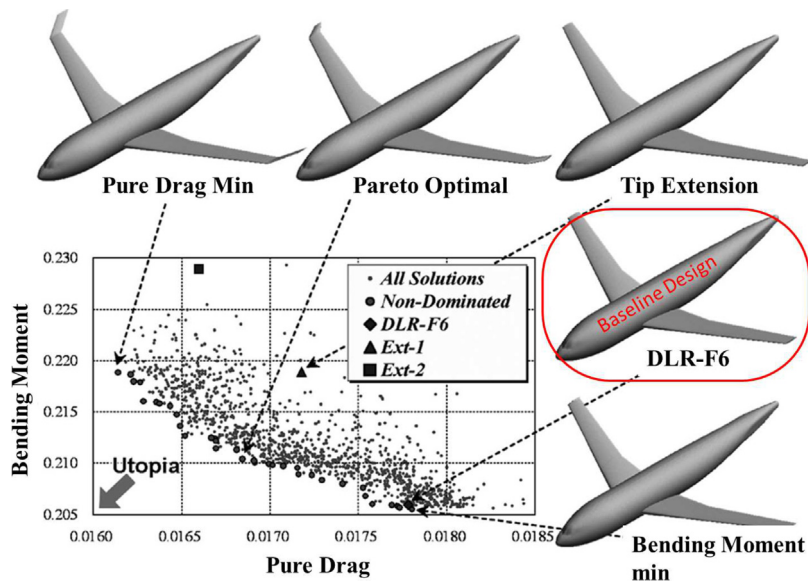


Fig. 12. Yamazaki et al. [11] non-dominated solution information for a winglet design problem using a genetic algorithm.

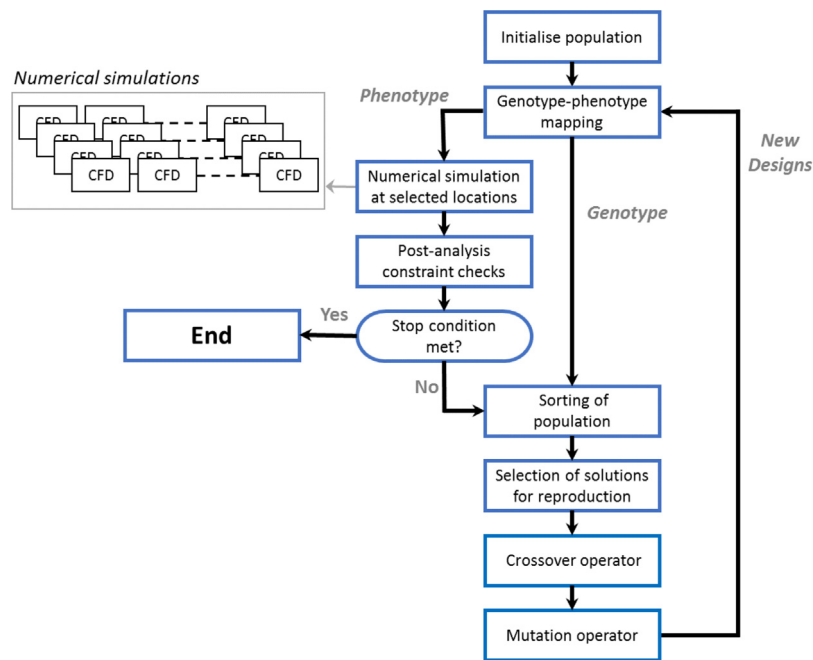


Fig. 13. Schematic diagram of typical genetic algorithm structure for aerodynamic optimisation.

is not a trivial task. A good selection of population size will improve both computation time and solution quality. It is generally argued that small populations can lead to premature convergence and poor solution optimisation and that larger populations may unnecessarily expend computational resources. Various generalised guidelines exist regarding appropriate population size and methods for tuning can be found in the literature [112–114]. Pandey et al. [115] present a detailed comparative review of approaches to prevent premature convergence based on several different factors affecting the GAs behaviour: initial population; population diversity; fitness function; search space scope Vs. selection pressure; problem difficulty Vs. number of individuals.

Selection methods and genetic operators exist in different forms and hold no strict rules on how they are implemented and is often down to user preferences. The most common selec-

tion methods for optimisations include: elitism; [111,97] roulette wheel; [107,103] and tournament selection [116,117,107]. These are somewhat independent of the genotype-phenotype mapping simply presenting alternative ways to select candidate solutions. After selecting candidate solutions the genetic operators breakdown and re-combination of schema (crossover operators) and perturb (mutation operations) random candidate variables. Genetic operations are dependant on genotype-phenotype mapping. The exact combination of genetic operations are often omitted from relevant literature however common examples of crossover include: single-point crossover; [97,110] two point crossover; [109–111] multi-point crossover; [118] gene-lottery; [119] uniform crossover; [110] and blended crossover (BLX- α) [120]. Mutation methods, and GAs as a whole, are covered in detail by Deb [120].

Perhaps one of the most common GA variants across all disciplines is the non-dominated sorting-based multi-objective evolutionary algorithm-II (NSGA-II) [121]. It is capable of both single and multi-objective optimisation, typically enforcing constraints through tournament selection and elitism. Lyu et al. [71] show that NSGA-II, for RANS-based optimisation, with a population of 24 for 200 generations struggled to handle more than eight variables incurring great computational expense and failed to meet the optimisation tolerance. Kim et al. [122] found NSGA-II to be sufficient for the aerodynamic and aeroacoustic optimisations of an axial-flow fan however only employed two design variables. Droandi and Gibertini [123] successfully applied the NSGA-II to the aerodynamic optimisation of tilt-rotor aircraft blade using a compressible Navier-Stokes solver. They define nine sections along the blade span for 27 variables and consider both single-point (helicopter mode or aeroplane mode operation) and multi-point (combined helicopter mode and aeroplane mode operations) optimisation. This study suggests NSGA-II can be expected to perform well in terms of Pareto optimisation tolerances due to the non-dominated organisation of the problem. This helps move the population toward Pareto optimality without suffering from convexity within the problem.

Genotype-phenotype mapping is a major influencer in a GAs performance. The genotype space presents the genetic blueprints or the DNA of the geometric search space which dictate the rules and traits of candidate solutions. The genotype is then translated through genotype-phenotype mapping to express a solution's attributes which add to or detract from its performance relative to the aerodynamic objective function. These are the solutions observable traits, referred to as the solutions phenotype, which feed into the numerical analysis.

Traditional GAs use binary encoding [124] to describe candidate solution genotype. In binary encoding, the number of exploitable schema is maximized [125] and the algorithm is able to converge quickly to a solution [126], but can yield low-quality solutions when applied with many variables resulting. In binary encoded GAs, string length must be assigned *a priori* giving the algorithm discretised precision; the higher the precision required, the longer the string length and considerably slows the algorithm and memory requirements [127,120]. Avenues for decreasing the length of large binary chromosomes are covered by Mcgookin [111].

A particular drawback in binary encoding is the existence of Hamming cliffs in the genotype search space [128]. This is a limitation caused by the discretisation of variables in genotype space. It requires the GA to simultaneously change the genotype bit representation in a very small and precise manner to achieve a more optimal solution. This is obviously an issue in consideration of aerodynamic optimisation as often small changes in variables can have large impact on performance, but its inevitable presence is often ignored. The probability that genetic operations will achieve this is unlikely while also the binary code does not preserve the locality of points in the phenotype space [129]. Additionally binary GAs often suffer from bias towards superior solutions in the early population stages due to genetic drift. Genetic drift will encourage premature convergence to sub-optimal solutions and is discussed in detail by Lim [130]. Work towards maintaining a heightened population genetic diversity ensures avoiding premature convergence [131], where diversity is the volume of dissimilarity between individuals within a given population. Non-random mating has been shown to maintain genetic diversity [132]. Massaro and Benini [133] discuss a variety of relevant diversity preserving techniques and compare several binary encoded GA variants.

Benini [134], Zhao et al. [135] and Skinner and Zare-Behtash [67] present different binary encoded GAs for aerodynamic design optimisation each maintaining diversity differently. Benini [134] use the Genetic Diversity Evolutionary Algorithm (GDEA) developed by

Toffolo and Benini [136] in the optimisation of a transonic compressor in CFD. This algorithm uses a diversity preserving mechanism which treats the genetic diversity as an optimisation objective to establish a criterion for fitness assignments. The dissimilarity between two individuals would be measured mathematically in phenotype space by Euclidean distance; this then increases the chance of aerodynamically superior but geometrically different solutions moving into genetic operations. This enhanced exploitation and exploration within the design space involving 23 design variables optimising a population of 20 individuals for 100 generations with a run time of 2000 processing hours, with four-processor parallel computing.

Skinner and Zare-Behtash [67] developed a population structured genetic algorithm (sGA) with dynamically structured binary genotypes enabling alternative non-planar wing configurations of to exist in parallel. The sGA facilitated gene comparability by means of a genetic hierarchy that enables diverse configurations to be maintained simultaneously through the inclusion or exclusion of certain characteristics [137]. Additionally, the application of an adaptive mutation rate was used to balance global and local searching ranges of the design space. This enhanced the search in the initial stages of the optimisation and then reduced exploration in the later stages as the optimisation process matured. This enabled the sGA to identify solutions with higher potential early in the optimisation process, and in later stages focus on the betterment of solutions identified taking emphasis away from exploring for entirely new solutions. The design space involving 28 design variables optimising a population of 100 individuals for 600 generations with a run time of 3–4 processing hours using a mid-fidelity aerodynamic analysis code.

Zhao et al. [135] used a Multi-Island GA (MIGA)[138,139] to optimise laminar flow regions over an aerofoil. The MIGA operates by dividing the population into sub-populations isolating them in different regions of the design space. Solutions thus develop into independent groups under different parametrisation and are systematically allowed to migrate between 'islands' and interact after every few generations [139]. A population of 50 solutions divided over 5 sub-populations were evaluated over 50 generations with solution migration required every 4 generations. The high frequency of migrations required to maintain diversity and the low sub populations suggest that the algorithm was very prone to genetic drift.

Real-number encoding is more commonly used to resolve limitations of binary encoding, and is widely confirmed to be more efficient than binary [140–143]. In real-number encoding the genotype and phenotype spaces share an identical topological structure [120], and has the advantage of dynamic coding and floating point representation to tackle large design spaces that require a continuous design space that would otherwise be discretised [144]. Real-encoding gives a robust search while keeping the string length small. In addition, real-encoding eliminates the need to define search boundaries for the algorithm as the distribution of solution candidates continuously moves in the direction of more promising regions within the design space [145]. If there are strict design variable limits then well-defined constraints which limit the algorithm variable boundaries must be used.

It is worthwhile indicating that ES are essentially the same as real-encoded GAs with the crossover operation removed and a less-random mutation methodology applied to influence parent solutions. Hence, ES are a perturbation based optimisation algorithm without the schema maintaining operations of a GA. ES are often deemed better for handling vast design spaces as demonstrated by Jones [146] and have been adopted in some instances for aerodynamic optimisations. Olhofer et al. [147], in the optimisation of turbine blades, indicate that ES can utilise smaller populations than GAs and achieve faster convergence thus reducing the compu-

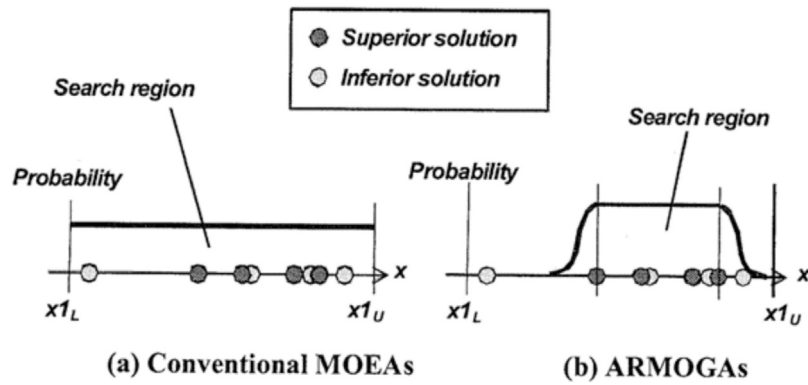


Fig. 14. Sketch of search regions [149].

tation time. However, it is indicated that subsequent optimisations are required to gradually increase the dimensionality of the problem to allow the ES to cope with a complex design spaces. Periaux et al. [148] employ a similar ES method in which they layer the optimisation algorithm's the population size, the mutation rate and CFD mesh refinement. It was found that the ES algorithm employed was unable to cope with multiple objectives or multi-disciplinary design problems.

Hashimoto et al. [28] explored the maximisation of fuselage/wing lift generation for high-wing aircraft configurations using RANS with design variables. They employ a real-encoded, adaptive range multi-objective genetic algorithm (ARMOGA)[149] to reduce the computational burden and the total number of function evaluations needed. ARMOGA was developed for integration with computationally demanding CFD-based optimisation to find multiple Pareto-optimal solutions more efficiently than conventional multi-objective GAs, such as NSGA-II [121]. This is due to their ability to perform concentrated searches, adapting their search regions based on population statistics as shown in Fig. 14. In comparison, the search region of traditional GAs remain constant. When tackling large aerodynamic optimisations, particularly when using CFD, it is detrimental to the optimiser performance to generate many infeasible solutions. Modifying the search range of the algorithm makes it possible to estimate candidate solutions efficiently and prevent wasting computational resource on the generation and evaluation of poor solutions [149]. A clear conflict here is the preservation of population diversity. Fig. 14 highlights that the population becomes confined into a reduced search space. This lowers computational expense, however it encourages a lack of genetic diversity. To counter this the ARMOGA uses population re-initialisation every few generations to randomly increase diversity. This preservation technique was adapted from micro-genetic algorithm (μ GA)[150] which was designed to maintain small populations. The ultimate problem with this method is that even though there is a constant infusion of new schema at regular intervals, increasing the dissimilarity between aerodynamic configurations, facets of the schema are not allowed to propagate and develop; this can become a fundamental limitation.

Chiba et al. [59] also employed the ARMOGA and present one of the few optimisation studies which consider a GA for high-fidelity multi-disciplinary optimisation. In order to meet the computational demands in a reasonable time two super computers were employed. The ARMOGA used 35 design variables with 5 constraints to optimise an aeroelastic wing by varying the wing aerofoil shape, twist and dihedral distributions while maintaining wing topology. A population of 8 individuals for 19 generations with a population refresh rate and search region range adjustment every 5 generations. It is reported that the algorithm did not converged but sufficient optimality was reached. A single generation composed of

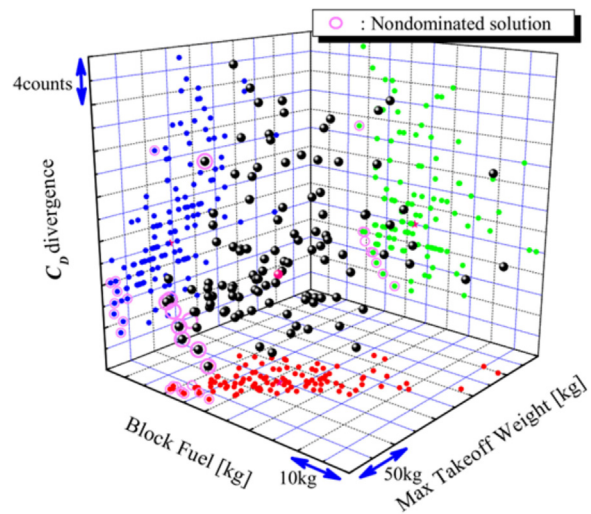


Fig. 15. All solutions identified by ARMOGA in three-dimensional space of all objective functions [59].

approximately 70 Euler computations to perform static aeroelastic deformation, and 90 RANS computations for the aerodynamic evaluation of the deformed wing; equating to roughly 880 hours CPU time. Furthermore, multi-point optimisation for 3 flight conditions were considered which is notably unique for optimisation studies of this type. Solutions evaluated are shown in Fig. 15, projecting all solutions evaluated in a three-dimensional space between the design objectives specified: minimisation of the block fuel, maximise take-off weight and maximise drag divergence. It is clear that the non-dominated solutions did not comprise a Pareto front and so no globally optimal trade-off was found.

Sasaki et al. [151] applied the ARMOGA to a transonic and supersonic wing design, optimising 72 design variables describing the wing shape and planform subject to four-objectives, and geometric constraints. The GA used a population of 64 due to the definition of a large search space. It was found that by tightening the definition of the thickness distribution constraints, more realistic wing solutions were identified. This highlights the capability of GAs in taking advantage of loose problem definitions, especially in multi-disciplinary optimisations. Sasaki et al. [152] point out that the GA solutions optimised for a cruising supersonic wing, with 105 variables and 15 constraints, satisfied all defined constraints but essentially over optimised the wing aerodynamics, thus compromising the structural integrity of the wing. This highlights the need for precise definitions of the problem via constraint handling to carefully reach optimised results. Optimal solutions are often found near, or on constraint boundaries, therefore, more thorough

constraint definition will result in a more efficient algorithm and superior results.

Finally, a resource based problem, almost exclusively a problem for heuristic optimisers and of particular importance to GAs, is the need to ensure careful distribution of computational resources; this is reflected by the use of parallelism of GA in all applications. Typically aerodynamic evaluations are executed in parallel with the grid generations and associated flow calculations distributed over available processing elements [153,145]. The selection of topology connecting the processors is sensitive and will significantly effect the calculation efficiency. Wang et al. [153] present a study on the performance of GAs operating in parallel for the optimisation of aerofoils with high-lift devices and convergent-divergent nozzles. Tse and Chan [154] and Holst and Pulliam [155] provide aerodynamic optimisations where the order 10^3 function evaluations were required to reach convergence making the employment of parallel computing mandatory.

4.2. Particle swarm optimisation

Particle Swarm Optimisation (PSO) is a stochastic population based (swarm) optimisation technique developed by Kennedy and Eberhart [156,157] in the 1990s. Mathematically, swarming is the collective decentralised motion of a large number of self-propelled entities as a collective animal behaviour and is exhibited by many living creatures such as birds, fish, and insects. Studies have praised the algorithm for easy implementation to aerodynamic solvers, its computationally inexpensive memory requirements, and simple mathematical operators [158]. PSO is developed from an agent orientated paradigm, meaning that particles are semi-autonomous agents capable of communicating and updating their status. This autonomy and ability for particles to sense their surrounding environment makes PSO an attractive gradient-free optimiser. However, the PSO is inherently unconstrained which can make handling variables, and enforcing design constraints difficult [51]. Zhang et al. [159] have carried out a comprehensive survey of several PSO variants outside of aerodynamic optimisation.

A swarm of particles represents a group of potential solutions moving in an n -dimensional space of design variables [160]. In the multi-dimensional design space, particles are assigned positions, representing the candidate solution, and velocity components which determines how the solution variables are updated. PSO algorithms, in their most basic form, can be summarised as [161,160]:

1. **Random generation** of particles in space to form the initial population with position and velocity vectors.
2. **Evaluation** of the fitness of each particle in the given environment.
3. **Update** each particle's velocity and position within the design space, adapting with regards to its personal best position (p_{best}) and the global best swarm position (g_{best}).
4. **Iterate** over steps 2 to 3 until some convergence criterion is met.

Particles typically update their search direction based on their current position (\mathbf{x}_k^i), the particles best known position (\mathbf{p}^i), and swarm best known position (\mathbf{p}^g). Fig. 16 illustrates this iterative process. The particles are typically manipulated according to Eq. (2), originally defined by Kennedy and Eberhart [157], with three weighting factors: the inertia weighting factor, w , the cognitive acceleration factor, c_1 , and the social acceleration factor, c_2 , where R is uniform random numbers from interval $[0, 1]$. The original

PSO algorithm uses constant values of 1, 2, and 2 for w , c_1 , and c_2 respectively [157].

$$\mathbf{v}_{k+1}^i = \underbrace{w\mathbf{v}_k^i}_{\text{Current < motion}} + \underbrace{c_1 < R(\mathbf{p}^i - \mathbf{x}_k^i)}_{\text{Particle < memory < influence}} + \underbrace{c_2 < R(\mathbf{p}^g - \mathbf{x}_k^i)}_{\text{Swarm < influence}} \quad (2)$$

The example for a particle searching with PSO is illustrated in Fig. 17, where \mathbf{x}_k^i is the position of particle i at iteration k , \mathbf{v}_k^i is the current velocity of particle i at iteration k , \mathbf{v}_{k+1}^i is its search direction, leading to an updated position \mathbf{x}_{k+1}^i of the particle i for iteration $k+1$. Movement of a particle is governed by interactions of inertial current motion, cognitive influence and social influence as shown. With a strong emphasis on the social parameter, with a low to zero cognitive parameter, the algorithm will converge quickly to an initial best solution. When there is more emphasis on the cognitive parameter, with a low or zero social parameter, no global information is shared and the algorithm will not converge properly. From these interactions the velocity vector \mathbf{v}_{k+1}^i is defined so that the particle moves to new position \mathbf{x}_{k+1}^i .

The basic algorithm described above has an undesirable property when $\mathbf{x}^i = \mathbf{p}^i = \mathbf{p}^g$ for any particle i . If this scenario arises, the velocity update Eq. (2) reduces to $w\mathbf{v}^i$. Therefore, when the velocity is close to zero, all particles will stop moving and converge prematurely. Key influences on the PSO algorithm performance in aerodynamic optimisation problems include: 1) swarm size and initialisation; 2) appropriate weighting of the inertial weight, and cognitive/social acceleration factors; 3) constraint-handling 4) computational resource.

Swarm size is directly related to the scale of the problem; with increasing shape dimensionality, more particles will be required to adequately search the design space for optimality. Shi and Eberhart [162] performed comprehensive studies on the effects of swarm size on the convergence properties and solution quality of the PSO algorithm. They concluded, in general purpose application, that swarm size had little effect on overall performance as long as the swarm falls within a reasonable range; this range typically being 30 to 40 particles. Unfortunately, most aerodynamic optimisation studies provide few details on swarm size. Additionally, poor performance may be attributed to the initial population distribution. If the population does not adequately cover a search region efficiently, it may not be possible to locate regions of optimality. Grosan et al. [163] discuss population behaviours and distribution in detail for general-purpose PSO application.

Many aerodynamic optimisation studies directly employ the classical Eq. (2) with no modifications [164,165,24,166,61,167]. Hence, PSO is often considered as a general-purpose optimiser, which is capable of handling different types of variables and functions with no adaptation necessary. Ouissa et al. [167] use the classical PSO algorithm, with 20 particles for 100 iterations, to optimise the energy capture by a wind turbine through real-time control of the blade pitch. They found that the original PSO algorithm was capable of dealing with strong non-linearity occurring from dynamically varying wind-speeds. Therefore, reflecting the capacity for adaptability and fast convergence for single objective, low search space dimensionality optimisation. The inertial weighting factor, w , was set to 0.73 which is a common factor among aerodynamic optimisation studies [168].

Venter and Sobieszczanski-Sobieski [164], also employing the classic PSO algorithm, showed that it can simultaneously handle continuous and discrete multi-disciplinary bi-level optimisations coping with severe numerical noise. The algorithm's weighting are modified however ($w = 1.4$, $c_1 = 1.5$, $c_2 = 2.5$) in order to bias convergence towards the globally best solution. No problem specific tuning was performed as emphasis of the work was placed on general application of the PSO algorithm in a multi-disciplinary

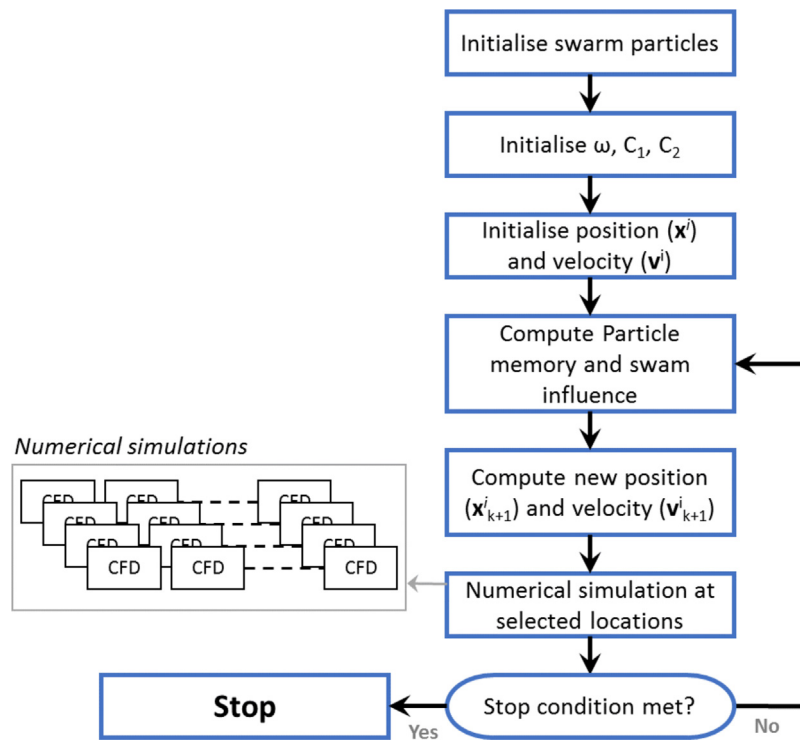


Fig. 16. Schematic diagram of typical particle swarm optimisation structure.

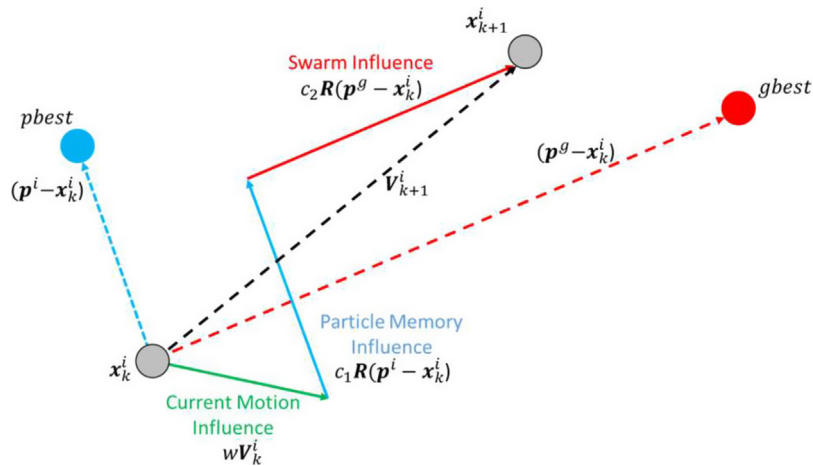


Fig. 17. Depiction of the velocity and position updates in particle swarm optimisation.

aerodynamic environment; a swarm size of 100 particles was used. The algorithm was given an unconstrained problem which aimed to maximise the range of the wing utilising 210 geometric and structural design variables. Although the problem presented by the authors was unconstrained the problem still required variable bounds. The approach used by the authors was to restrict the velocity vector of a particle forcing it to move back towards the feasible design space. Dealing with constraints in this way is discussed in depth by Perez and Behdinan [169]. Forcing particles to meet some condition may begin to artificially drive a solution. Alternatively, penalty-based constraints are known to work well with gradient-free optimisers but can have a significant influence on performance and can lead to numerical ill-conditioning.

Performing a statistical analysis of the PSO solutions Venter and Sobieszczyński-Sobieski [164] show that the standard deviation over ten independent optimisations was 0.17% relative to the mean objective function which improved by 6.43%. This robustness did

however come at a high computational cost, requiring an average of 9660 function calls per optimisation. It is identified that this high cost is due to the PSO algorithm's un-tuned weightings and the fact that the objective space contains severe numerical noise. Studies focused solely on tuning PSO performance such as that by Pant et al. [170] highlight that poor selection of the weighting parameters can lead to premature, very slow, or no convergence at all; the literature reflects that there is no standard way to balance cognitive or social weightings and problem specific tuning is necessary.

Jansen et al. [24] used a medium-fidelity aerostructural panel code approach and an augmented Lagrange multiplier particle swarm optimisation (ALPSO) algorithm to show that a winglet is optimal when span constraints are present. The algorithm was additionally capable of accurately quantifying optimal configurations and raked wing and box-wing under different conditions. When the algorithm compensated for the effects of viscous drag a C-wing configuration was optimal, however with the inclusion of

structural considerations the C-wing and box-wing configurations were found to add more structural weight than their respective drag reduction potential could compensate for. The ALPSO parameters used, defined by Jansen [171], a swarm size of 40 for 12 iterations using: $w = 2$; $c_1 = 2.5$; and, $c_2 = 0.5$. Despite the seemingly good performance, the optimiser was found to converge prematurely forcing several optimisations to be considered for each problem. The ALPSO algorithm was tuned to also reject infeasible solutions which was shown to help move the algorithm towards more feasible designs. The same algorithm was also employed by Haghighat et al. [166] with a similar physics solver for the design optimisation of an active load alleviation control system for flexible wings. The performance was not commented on by the authors. Lyu et al. [71] use the ALPSO algorithm for the lift-constrained drag minimisation of the NASA CRM under RANS equations. Using a swarm size of eight for eight wing twist variables they find that the optimiser is sufficiently capable however ill-suited to the problem and so incurred excessive computational expense.

The social and cognitive weightings can also be defined by constriction models, as implemented by Azab and Ollivier-Gooch [168]. In this constrained aerodynamic optimisation of aerofoil design, it was found that this approach was susceptible premature convergence but is overcome by randomly re-scattering particles if premature convergence was detected. In addition the use of scalar quantities for the acceleration factors in aerodynamic optimisations was expanded upon by introducing non-linear mathematical programming and additional weightings [172].

Using a fixed inertial weight can have detrimental effects on the convergence behaviour of the algorithm. Praveen and Duvigneau [165] present a dynamically decreasing inertia weighting factor (from 0.9 to 0.4) for the optimisation of supersonic and transonic wing designs using high-fidelity CFD. The dynamic update of the inertial weighting for the betterment of the PSO algorithm for shape optimisation was proposed by Fourie and Groenwold [172]. This ensured a more exploratory search for optimal configurations is followed by the promotion of convergence towards the best known solutions. It is common practice in aerospace focused optimisations that these weightings are not constant scalars, and instead are implemented as some vector of weights. Appropriate inertia weighting balances global and local searching to either increase solution diversity and search the whole design space or alternatively encourage intensified searches in local regions. The idea being to terminate the PSO algorithm with a more local search.

Wang et al. [158] and Nejat et al. [173] each present CFD-based aerofoil optimisations adopting dynamically varying weightings. Both studies argue that scalar implementation is more susceptible to becoming trapped in a single search direction towards a local optima. The inertial weighting linearly decreases in the same manner described for Praveen and Duvigneau [165]. The cognitive and social weighting, for both studies, vary randomly between 1.5 and 2.5. Wang et al. [158] found this to give the algorithm good global search ability while maintaining good convergence characteristics for 19 design variables achieving an average population drag reduction of 25%. The algorithm also seems to have been able to successfully identify regions within this design space that corresponded to supercritical and high-lift aerofoils with supercritical characteristics.

Nejat et al. [173] used dynamically varying weights to mitigate premature convergence. They modified the classic PSO algorithm to increase stability and address premature clustering in the early stages of the search process which PSO studies have found to be a common issue among the algorithm variants [174,175]. Modifications they employ include: 1) addition of a constriction coefficient to limit the maximum velocity of particles in space (divergence); 2) non-dominated sorting; 3) neighbour density estimation to encourage exploration of less crowded regions; and 4) addition of




Fig. 18. Nacelle design variables used by Li et al. [61].

a collision operator to the velocity update equation (Eq. (2)). Detection of less-crowded regions of the design space coupled with the collision operation introduces areas of attraction and repulsion to the decentralised search. Nejat et al. [173] found that this gave high diversification in early stages and avoided premature convergence; however, as the tendency for convergence increased the regions of repulsion had to decrease providing sufficient diversity at different stages of swarm optimality. The algorithm was shown to perform well when optimising an aerofoil stall characteristics based on four conflicting objectives, able to find optimal trade-off characteristics for the aerofoil maintaining suitable performance metrics. Other forms of repulsive PSO algorithms are present in aerodynamic optimisation problems, but are generally criticised for slowing the algorithm performance and poor ability for local search. Guglieri [126], for example, found that a repulsive PSO algorithm in the optimisation of helicopter rotor blades gave sufficient optimisation tolerances but performance was very slow likely due to solution diversity. Alternative methods to mitigate premature convergence may include craziness operators [176,165].

In contrast to repulsive PSO algorithms that have been used in aerodynamic design applications in an attempt to prevent swarm clustering, there are PSO algorithms which encourage swarm clustering. Li et al. [61], among others, [168,177] used clustering/grouping to improve the PSO exploration and exploitation of design variables and/or objective space. Using CFD-based optimisation, Li et al. [61] perform the optimisation of the engine nacelle/pylon position on an aircraft in order to minimise interference drag subject to vertical and horizontal position, shown in Fig. 18, and 12 constraints restricting maximum movement due to mesh limitations. The method used here to form sub-groups within a design space uses the classic PSO algorithm and simply develops different groupings with a shared objective due to the arrangement of unequal weightings. Therefore, the swarm develops simultaneously in various ways but maintains swarm social learning. By this means it ensures a global optimisation ability of the entire group while accounting for local search ability. The algorithm was able to communicate design features between the sub-groupings manoeuvring the engine nacelle in such a way to reduce the local velocity over the pylon, thus reducing the strength of the interference drag experienced reducing the aircraft drag coefficient by 1.1%.

A different method of swarm clustering is presented by Hu et al. [177] in optimising the air supply to maintain air quality inside the aircraft cabin. Sub-swarms maintain equal weightings, and alter-

natively are assigned different objectives. This essentially defines single-objectives sub-swarms with bias to different search regions which, through social interaction, form a multi-objective swarm. This represents a strong design strategy for the optimisation of problems where the relative importance of many conflicting objectives is unclear [161].

4.3. Simulated annealing

Simulated Annealing (SA) is a stochastic point-by-point optimisation algorithm method based on the physical cooling process of molten metal [178–181]. The physical process of annealing provides the inspiration for finding optimal solutions to combinatorial problems based on statistical mechanics. SA is loosely related to the Hill Climbing [178] optimisation algorithm in that both algorithms perform a point-by-point local search around the current candidate solution. The main significant difference is that Hill Climbing will only accept a new solution if it is better than the current one: SA incorporates the possibility of accepting poorer candidate solutions, thus meaning the algorithm can escape local optima. This procedure of accepting poorer solutions is referred to as the Metropolis Procedure/Criterion [179], and it acts to mimic the probability of atoms jumping between higher and lower energy levels as the material cools [182].

The main iterative steps which the SA algorithm operates is as follows [183,184]:

1. **Initial selection** of candidate solution and a high initial temperature (material state).
2. **Second selection** of a candidate solution is generated at random in the vicinity of the initial point.
3. **Comparison** is made calculating the difference between the aforementioned two solutions.
4. In the next iterations, another solution is created at random in the local neighbourhood of the current solution, and the **Metropolis Criterion** either accepts or rejects the new solution.
5. **Annealing schedule** is updated.
6. **Iterate** over steps 2 to 5 until convergence criteria is met.

In the literature, there are various suggestions on how to implement SA, however classical SA is expressed as Eq. (3) [111]. Typically, the temperature is decreased from the initial temperature, T_0 , to the current temperature, T_n , based on some reduction constant, γ^n , assigned by the annealing schedule at iteration n . Fig. 19 illustrates this procedure.

$$T_n = \gamma^n T_0 \quad (3)$$

The Metropolis Criterion performs a probabilistic check so that sub-optimal solutions are allowed to replace parameters that have better cost values. By analogy of the process of annealing, the probability (P) of the new candidate solution's cost (C_{new}) relative to the previous solution's cost ($C_{previous}$) is determined using Boltzmann's equation [179]:

$$P = \exp\left(\frac{C_{previous} - C_{new}}{T_n}\right) \quad (4)$$

P is then compared to a randomly generated number in the range $[0, 1]$ to determine acceptance or rejection of the new solution. A more in-depth discussion is provided by Rafferty [109]. Eglese [184] emphasises that application of the SA algorithm is sensitive to: 1) the initial temperature (T_0); 2) the annealing schedule/temperature function, to determine reduction constant γ at each iteration; 3) number of iterations to be carried out at each temperature; and 4) a stopping criterion for any application. The initial temperature is considered a significant control parameter such that initially,

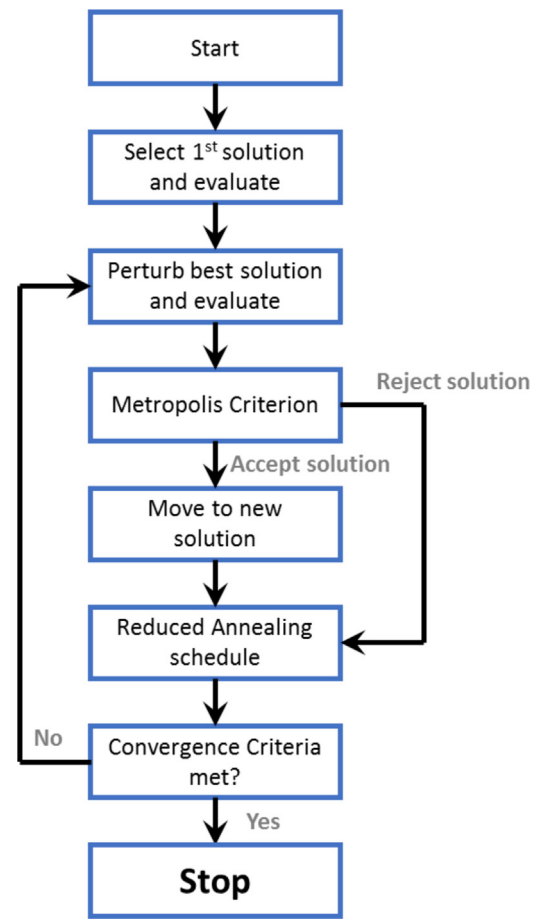


Fig. 19. Simulated annealing flowchart.

when the system is molten and the temperature is high, search perturbations are large [109]. As more iterations are performed, the temperature is decreased according to the so-called annealing/cooling schedule every n iterations; this is the so-called cooling cycle.

Tiow et al. [185] used the SA algorithm to optimise two separate turbomachinery cascades with the application of inverse design and CFD. The constraints, based on existing knowledge of the design space, were implemented to restrict the search to more optimal regions within the design space to increase the algorithm performance. Additional performance was incurred by coupling the CFD solver with a database of previously generated solutions to accelerate subsequent calculations. The algorithm used a fixed annealing schedule of $\gamma = 0.75$ which operated in conjunction with the Metropolis Criterion every 10 iterations for a total optimisation run of 300 iterations. Iterations used a purely random walk in all search directions and step sizes were taken from uncorrelated uniform distributions [186]. The initial temperature is set low at 0.5°C ; typical initialisation temperatures are much higher as low temperature initialisation can restrict the search of all design variables equally [187]. The low starting temperature suggests that parameters were finely tuned to the problem in order to remove unnecessary calculation [188]. Increasing the initial temperature above the optimal starting temperature would only increase computational expense without betterment of the solution. The objective of the optimisation was to minimise entropy losses through the cascade; this would provide a highly non-linear search space as the cost function, related to aerodynamic losses, is a function of blade shape therefore presenting many local optima. The blade profiles were parametrised with 58 control variables.

Loss reductions on the order of 20% were achieved in both test cases reflecting the ability of the algorithm to find superior results.

Leylek et al. [189], use the classical SA algorithm for a bat-inspired wing morphing aerodynamic optimisation in an attempt to expand the wing flight envelope; details of the SA algorithm are presented by Manzo et al. [190]. The algorithm is reported to have taken between 500 to 2000 iterations (equating to 3 to 12 hours total run time) terminated by a convergence criterion for which 10 successive function evaluations must deviate less than 1% from one another, thus saving unnecessary computation time. The algorithm determined the starting temperature by taking 30 random samples from each variable in order to assess the average cost increase associated with poorer solutions which is then divided by the initial Metropolis Criterion, 0.95. This ensured that the initialised temperature held an element of randomness and was very high, allowing the algorithm to avoid local optimal in early iterations. The annealing schedule used a reduction constant of 0.85, which implemented a temperature decay every 160 iterations. The authors point out that all 5 design variables describing 3 wing topologies with varying aerofoil section, shown in Fig. 20, were presented in a discretised space which was found to decrease the accuracy of the algorithm but also decrease the search space size and the total run time. Fig. 20 illustrates the capability of the wing morphology even though a discretised design space is used. The algorithm, using a panel code for aerodynamic analysis, was able to adequately capture the wing active camber balancing 4 performance metrics, each one posing conflicting performance characteristics. Comparing the three different bat-wing topologies the algorithm isolated and exploited the wing shape characteristics such as aspect ratio, camber distribution and wing tip deformation to improve the planform performance dependant on the performance metric in interest. This indicates that the algorithm was able to use bat-like morphology optimisation of a wing to provide significant performance gains and adaptability in the flight envelope.

The classical SA strategy lacks the inherent ability and mechanisms for co-operation between individual points due to its stochastic point-by-point behaviour. Thus, as the search space dimensionality increases, the classical SA algorithm will become less efficient in searching for the optimal solution. Motivated by this, Liu [191] modify the search behaviour of the classical SA algorithm by incorporating fractional factorial analysis based on Taguchi orthogonal arrays. This enhanced numerical convergence and accuracy of the optimiser for higher-order problems while addressing the algorithm's poor performance in highly-dimensional design spaces. Fig. 21a illustrates the search behaviour for the classical SA and the Taguchi-SA methods; Fig. 21b and c show the resulting search paths determined by the different search behaviour for the classical SA and Taguchi-SA for a non-linear multi-modal function respectively. The classical SA algorithm assesses new candidate solutions, x^p , randomly. The Taguchi-SA assesses new candidate solutions according to structured orthogonal tables which determines systematic and uniform sampling from the neighbourhood of the current candidate solution. This methodically improves the combination of design variables to the best candidate solution, x^{best} , within a defined search radius. Comparing Fig. 21b and c, the increased efficiency of the search behaviour is evident. Liu [191] also show the algorithm capable of handling functions with up to 100 dimensions.

Liu [191] applied the Taguchi-SA to aerodynamic design optimisation of high lift aerofoils and wing planforms, using a finite volume solver, were the enhanced SA algorithm was able to further optimise an already optimised high lift aerofoil. The classical Metropolis Criterion, and a fixed cooling rate of 0.85 was used. The primary objective was reduction of the nose-down pitching moment with secondary objectives of increasing the lift-to-drag

ratio. A reduction of 9.11% in the pitching moment was achieved, and the lift-to-drag ratio increased by 32.74%. Furthermore, the Taguchi-SA was used to optimise a supersonic wing planform with variables of sweep, aspect ratio, and taper ratio to minimise the drag rise experiences when transitioning through transonic flows. The algorithm was found to increase the sweep angle and span (increasing the aspect ratio), whereas the taper ratio of the wing was reduced; this reduced the average drag force between Mach 0.2 to 2 by 2.74%.

Mukesh and Lingadurai [183] apply the simplex simulated annealing (SIMPASA), established by Press and Teukolsky [192], to maximise the aerofoil lift coefficient at a fixed angle of attack, with respect to 12 design variables and subject to 5 constraints. Modifications to the SA algorithm enable it to make more appropriate moves to new candidate solutions by combining the Metropolis Criterion algorithm with the non-linear simplex algorithm. The classical SA and SIMPASA algorithms increased the lift coefficient by 22.8% and 27.4% respectively, showing robustness for global optimisation of non-convex, highly constrained aerodynamic optimisation. Mukesh and Lingadurai [183] provide no detail of the algorithm annealing schedule or other optimiser variables but state that the algorithm was sensitive to the number of iterations for each cooling cycle and initial temperature.

Wang and Damodaran [193] offer a computational brute force methodology to improve the SA algorithm through parallel-processing. They argue that stochastic methods such as SA are very appropriate for multi-disciplinary aerodynamic optimisations however incorporating modifications to the classical SA algorithm can only result in moderate computational savings. Thus, the performance of a serial and parallel SA algorithm are compared in a generic wing optimisation problem adopting the Breguet-Range equation. This combines aerodynamic and structural considerations in the maximisation of range with respect to 4 design variables which are subject to 6 constraints. All SA settings are constant between the serial and parallel optimisations. A key point, which the authors do not address, is that applying parallel computing to the SA algorithm is analogous to a genetic algorithm with zero crossover operation probability. Results showed that the parallel implementation of the SA optimiser over eight processors achieved a 4.4 speed up time and achieved the exact same solution quality as the serial SA. The serial SA algorithm required over 600 objective function evaluations, whereas the parallel SA only required 148 function evaluations. Hence, the number of reduced function calls in this aerodynamic design application is either chance or by design. In parallel, eight random initialisation (compared to one in serial) of candidate solutions in the design space increases the likelihood of one of those point being sufficiently closer to an optimal design. Alternatively, the SA algorithms initialised in parallel may have been modified to use some method for co-operation to navigate the search space.

Subsequently, Wang and Damodaran [194] explore broader application of parallel SA optimisation with CFD for different aerodynamic optimisations, including: 1) diffuser shape design; 2) convergent nozzle; 3) supersonic axi-symmetric nozzle. Based on this work, among others such as Aly et al. [195], for aerodynamic optimisation, there are loose guidelines for the selection of annealing schemes and selection of termination criteria.

5. Application of hybridised algorithms to aerodynamic optimisations

Numerous hybrid algorithms which incorporate elements from different optimisation algorithms exist and have shown successful application. Here we will only consider hybrid algorithm schemes applied to aerodynamic design problems; in an optimi-




Fig. 20. Different bat-wing topologies for wing morphology optimisation using simulated annealing [189].

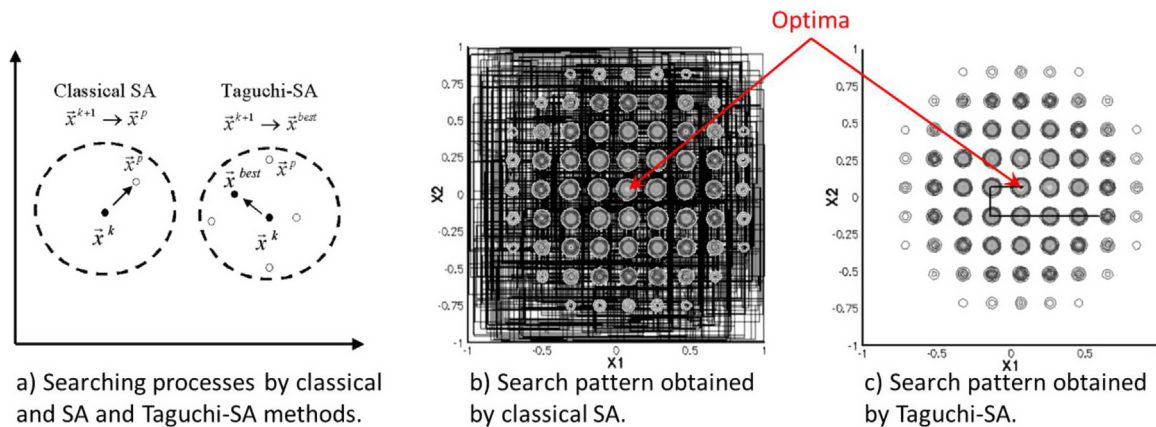


Fig. 21. Comparison of search processes and resulting search patterns obtained using classical SA and Tauchi-SA methods for a non-linear multi-modal function [191].

sation context, hybridisation can be defined as mixing two or more algorithms, with possibly, further complementary features.

Deterministic approaches take advantage of the analytical properties of the search space to generate a sequence of candidate solutions with systematic improvements, usually resulting in the number of design iterations required to be small; a major shortcoming is the dependency of a compatible design space and sufficient baseline geometry. Heuristic approaches do not take into account properties of the search space, but instead seek improvements to candidate solutions on the basis of experience and judgement with a probabilistic approach which can lead to large numbers of design iterations and long computing times. Therefore, hybrid approaches try to combine methods from these approaches in an attempt to mitigate the weaknesses each hold. Hybridisation does not exclusively require deterministic-heuristic combinations as heuristic-heuristic hybrid algorithms also exist.

The various methods of hybridisation between different types of algorithms can be classified into three main groups: 1) pre-hybridisation, where, for example, the population of the GA is pre-optimised using the GBM; 2) organic-hybridisation, in which the GBM is used as an operator within the GAs for improving each population member in each generation; and 3) post-hybridisation, in which the GAs final population is used to provide an initial design for the GBM. It should be noted that these classifications are not limited to the hybridisation of GAs and GBMs, however, the most frequent hybrid algorithm found in aerospace application have been designed in attempt to combine the best characteristics of GAs and GBM [196–203,23,122,98].

Gage et al. [196] present the post-hybridisation of a classical GA with sequential quadratic programming for the topological design of wings and trusses. This work is notable for hybridisation as it is one of the first hybrid methods (HM) employed in aerodynamic design optimisation. They demonstrated that post-hybridisation is effective for final refinement of the GA's candidate solutions. By switching to a GBM once the GAs population is sufficiently mature, computational demands can be reduced and superior solutions can be found relative to allowing the GA to continue. In more recent

work, Kim et al. [122] also use post-optimisation, to improve the aerodynamic and acoustic performance of a axial-flow fan, by combining the multi-objective real-encoded NSGA-II from which the Pareto-optimal solutions are further optimised using sequential quadratic programming. The specific difficulty in this method of hybridisation is the transition from a multi-objective problem to a single-objective problem. There are typically two ways to transition: 1) combine all of the objectives into a composite objective using a weighted-sum approach for example; and 2) sequential optimisation, optimising one objective at a time while treating all other objectives as equality constraints. Kim et al. [122] adopted the latter method, pointing out that it did not preserve Pareto optimality and created a set of optimal solutions for each objective with many duplications forming.

A pre-optimisation strategy is proposed by Chernukhin and Zingg [23] which takes advantage of the GA's stochastic search capability and the gradient-based optimiser SNOPT's ability to efficiently identify local optima and enforce constraints directly. They implement the developed HM into a series of different aerodynamic optimisation problems including: aerofoil optimisation, transonic wing-section optimisation, subsonic wing-section optimisation, and blended-wing-body optimisation. The HM's initial population is optimised by SNOPT for a limited number of design iterations on each candidate solution in order to balance biased and insufficient solution development. The improved candidate solutions are then stochastically perturbed by the GA's crossover operations (perturbs random solution variables) had to be avoided as it was found to conflicted with the refinement capability of the GBM. Pure-mutation operations (random replacement of solution variables) were included to improve the stochastic search. Similar pre-optimisation strategies have been employed by Xing and Damodaran [204,205] which combine GA stochastic searching and GBMs analytical optimisation procedures in the optimisation of nozzle shapes.

Compiled optimisation results from Chernukhin and Zingg [23] for a series of optimisation problems is shown in Fig. 22. From

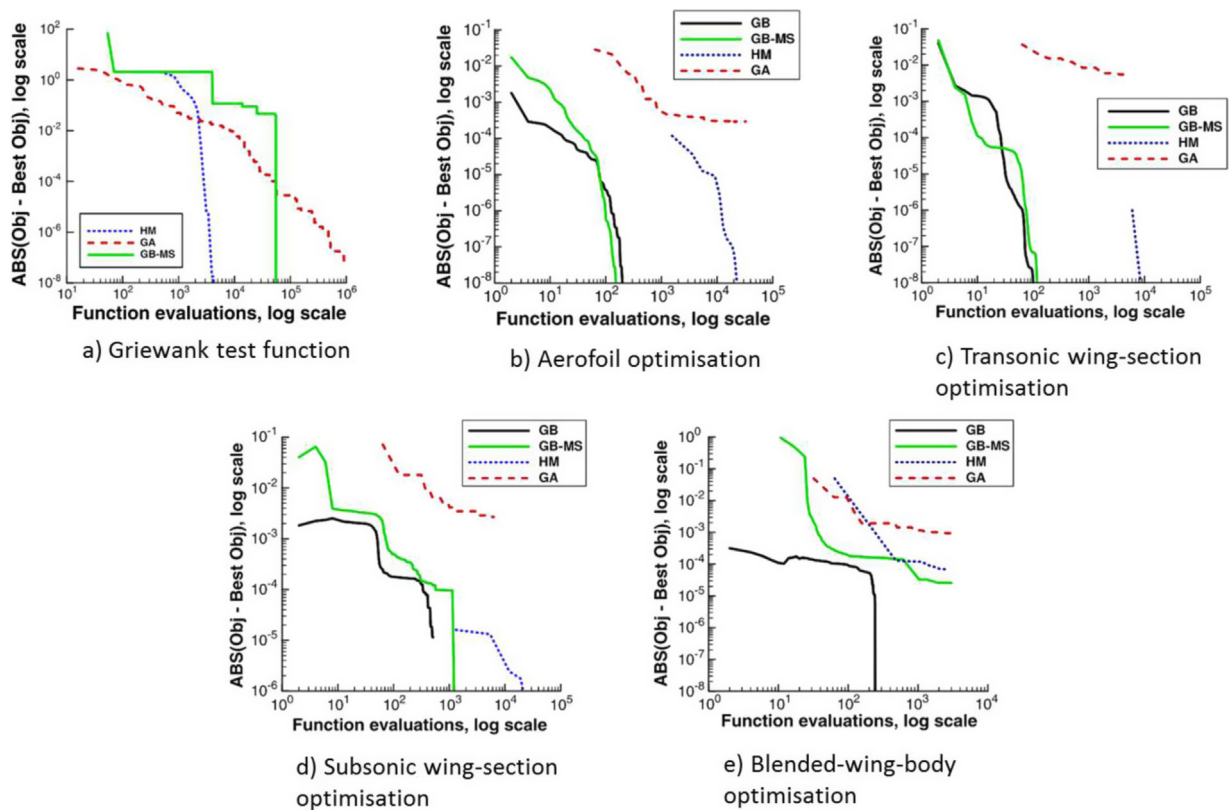


Fig. 22. Compiled optimisation results, from Chernukhin and Zingg [23], comparing optimiser convergence plots for a gradient-based (GB), a multi-start gradient-based (GB-MS), a genetic algorithm (GA), and a developed genetic algorithm/gradient-based hybrid method (HM) for different optimisation problems.

Fig. 22(a) the HM was found to significantly outperform other algorithms for highly-multi-modal problems; namely the Griewank function which offers hundreds of locally optimal solutions. Both the aerofoil and transonic wing optimisation are expected to be uni-modal base on the author's results. It is seen from Fig. 22(b) & (c) that the HM, achieving equivalent solution quality, added unnecessary computational expense relative to the GBMs which were several orders-of-magnitude more efficient; the HM and GA methods are useful to prove that no additional local optima exist. Both the subsonic wing and blended-wing body optimisation present mildly multi-modal problems in which 7 and 8 local optima were found respectively. The 8 local optima for the blended-wing-body optimisation can be seen in Fig. 1. For the subsonic wing design, Fig. 22(d), the GBM initialising with a single starting position within the design space could not find any globally optimal solution; however employment of the GBM with multi-start initial design points (GB-MS) found optimal designs with fewer function evaluations than the HM. Finally, in the convergence plot for the blended-wing-body, Fig. 22(e), the gradient method prove most effective. The GB-MS and the HM are seen to have similar convergence rates and are both expected to be able to find globally optimal solutions but were terminated due to time constraints. These results suggest that, for aerodynamic optimisation problems which are generally mildly multi-modal, the HM implemented here adds unnecessary complication and computational demand.

The classical GA has been extended for the purpose of more efficient aerodynamic shape design optimisation by Catalano et al. [98] to include two new operations based on gradient search in a hybrid organic-optimisation algorithm. Fig. 23 shows how the classical genetic algorithm has been modified to include such operators. The activation of each gradient operator is controlled probabilistically, inspired by classical crossover and mutation operation use. The first gradient-based operator, which acts on the whole population,

has two further probabilistic controls determining behaviour with each candidate solution. The first determines the maximum number of iterations and the second determines the sensitivity analysis method the gradient optimiser uses. The second gradient operator is relatively simpler, applying only one optimisation iteration to the current best solution according to the steepest descent rule with a random step size. Finally there is an elitist mechanism which preserves the fittest solutions from each generation and reintroduces the best known solution from all generations into the current population.

For the single-objective optimisation of an aerofoil with 24 variables, Catalano et al. [98] find that the complex HM developed is comparable to the standard GBM in terms of the number of function evaluations only achieving a 0.4% better objective. Under different settings (population, number of iterations, sensitivity analysis, etc.), they find that the hybrid algorithm slightly outperforms the gradient method, reducing the objective function further by 2.2%, but requires many more function evaluations. The classical GA is outperformed by all other algorithms. Therefore, the hybridised GA-GBM showed accelerated performance in aerodynamic optimisation relative to a standard GA and capable of matching, and in one case able to outperform, the GBM. The GBM (with a suitable sensitivity analysis) is shown overall more effective, the hybrid algorithm performed well however added unnecessary complexity to the optimisation framework.

PSO and GBM have also been combined for aerodynamic design. Jansen et al. [24] used a PSO algorithm post-optimised with the gradient optimiser, SNOPT, in the aerostuctural optimisation on non-planar wings using a panel method and potential flow theory. Little information is given on the hybridisation however the reasoning behind the method employed was to compensate for instances when the PSO converged prematurely. This ensured the

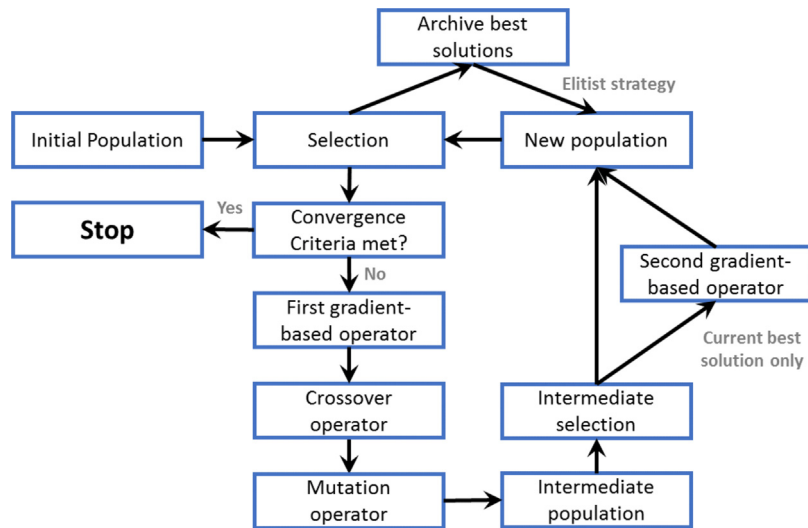


Fig. 23. Hybrid organic-optimisation algorithm [98].

PSO solutions were at a local minimum and offered further solution refinement where possible.

Azab and Ollivier-Gooch [168] propose an alternative technique which implements pre-optimisation of a swarm using SQP with adjoint gradients to find an initially high quality global swarm position before using the PSO algorithm. The objective is to define a high quality focal point in the design space that influences the swarm with the aim of reducing the total computational demand to find the global optima. In transonic drag optimisation with thickness and lift constraints, the pre-optimisation using the SQP method was not able to find a shock free optima but did reduce the drag. The SQP adjoint gradients tend to minimise the drag through poor mechanisms, thus becoming stuck in a local optima. The PSO algorithm essentially found the initial point provided by the SQP irrelevant as the SQP performance was limited in the given design space. They identify that the influence of numerical noise (partly due to the shock wave appearing and vanishing with geometric changes) coupled with the limitations of the Euler flow physics model are what prevented the GBM finding the appropriate gradient information. Additionally, when the objective space was very nearly flat, representing incremental improvements, the shallow gradients available in the objective space were overwhelmed by numerical noise meaning that pre-optimisation procedures did nothing. The computational expense of the hybrid scheme is reported at 4 to 9 times more than the SQP algorithm alone, but for all optimisation cases the HM found far superior optimal solutions. The authors did not compare the computational expense of the HM to the un-hybridised PSO algorithm.

Alternative hybrid methods based on hybridising two heuristic approaches also exist in aerodynamic optimisation literature. A modified binary genetic algorithm (GA) and SA have been used together for organic-hybridisation by Herbert-Acero et al. [206] in the optimisation of wind turbine rotors. A schematic diagram of the hybrid algorithm procedure is shown in Fig. 24. Fig. 24 shows that the hybrid procedure optimised the radial geometric distributions of the blade, using the SA algorithm, and the selection of aerofoil shape along the blade discretised span used the GA genetic operators. The GA did not optimise the combination of sectional aerofoil shape over the span, it acted to perturb the combination of sectional aerofoil (specifically the camber distribution) along the blade for which there were 100^{17} possible combinations. To demonstrate the hybrid framework developed, a SA algorithm was used to optimise radial geometric of different blades with fixed camber distributions from which the results were compared to hybrid procedure.

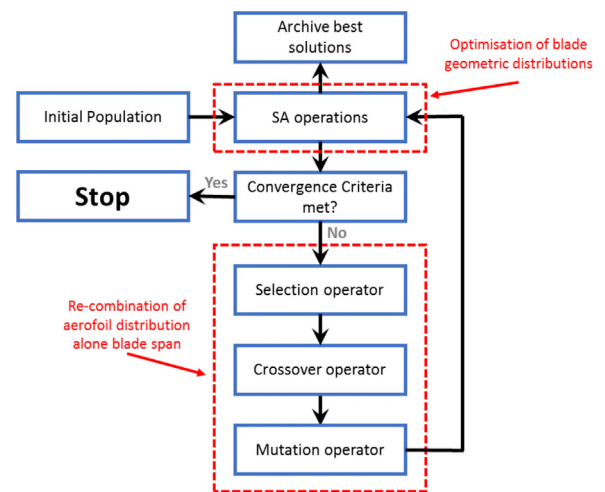


Fig. 24. Hybrid GA-SA organic-optimisation algorithm [206].

Fig. 25 compares the different geometries and aerodynamic performance curves for different optimised NACA 4-digit wind turbine blades. The genetic operations in the hybrid algorithm gave the optimisation process the ability to leverage the sectional aerofoil distribution to allow a nearly constant chord and thickness distribution. The obtained chord length attributed to aerodynamic benefits of operating at higher Reynolds numbers and the low thickness mitigated performance deterioration near stall. Furthermore, the results showed that the hybrid optimiser held distinct improvements in the blade design including: a reduction of the cut-in wind speed, increase aerodynamic efficiency, and an overall reduction of material used to manufacture the blades.

Other heuristic hybrids that exist include hybridisations of operations between GAs and PSO [207–209], or SA and PSO [210]. These studies typically attempt to handle highly multi-modal design spaces requiring large degrees of modification while addressing strongly stipulated aerodynamic constraints. For example, Khurana et al. [207] show that for aerofoil shape optimisation, the classical PSO algorithm was unable to optimise to a feasible optimal which met all constraints once the dimensionality of the problem exceeded 10 variables. Additionally, excessive computational resource was required for the evaluation of aerodynamically infeasible designs. Part of the problem was associated with the lack

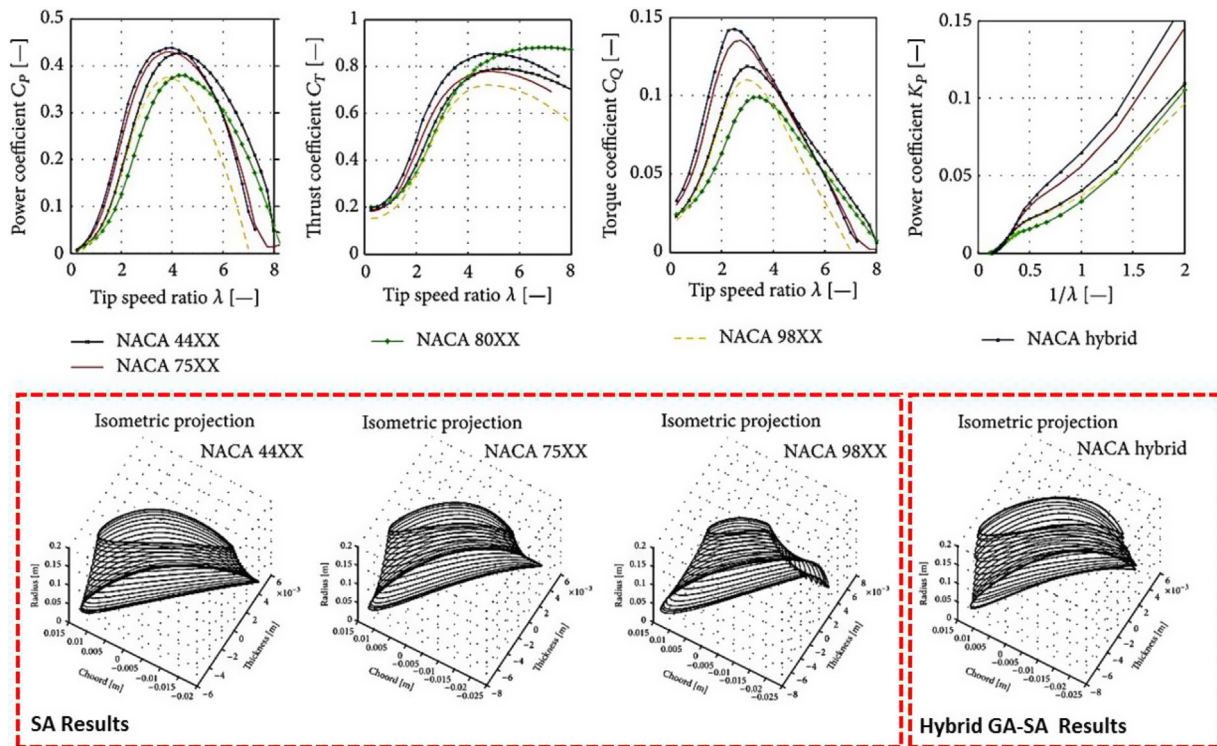


Fig. 25. Geometry and aerodynamic performance of optimised wind turbine blades, optimised by the SA and hybrid GA-SA algorithms [206].

of search diversity attributing to sub-optimal solutions. To overcome this they introduced systematic mutation operators from GAs and observed significant convergence improvements in fewer iterations.

6. Application of surrogate modelling to aerodynamic optimisation

Surrogate modelling can be viewed as a non-linear inverse problem with the aim of determining a continuous function that relates design variables to output responses from finite data. Surrogate assisted optimisation aims to alleviate the computational burden of the aerodynamic optimisation process by defining a simplified mathematical relationship allowing for fewer numerical simulations to be required. To interrogate the surrogate an optimisation algorithm is needed to perform a global search of the design space relating to the response surface. It is seen from the literature that surrogates are almost exclusively coupled with gradient-free population based optimisation methods, such as genetic or particle swarm algorithms [211,61,20]. With surrogate assisted stochastic optimisation Madsen et al. [212] optimised a diffuser shape, Li et al. [61] optimised the nacelle/pylon position for a wing, and Lundberg et al. [213] optimised ground vehicle aerodynamics.

Key steps are outlined by Hashimoto et al. [28] for a surrogate assisted aerodynamic optimisation which is summarised in Fig. 26. The construction of the surrogate generally consists of three steps: 1) design of experiment sample plan to generate initial sample points in the design space - point selection is discussed by Boopathy and Rumpfkeil [214]; 2) numerical simulations are performed to compute the output/performance of each sample point; 3) sample point data (input & output) are used by an approximation model to construct the surrogate. Replacing a particular problem analysis with a surrogate analysis does not affect the problem formulation, but it will strongly influence the solutions identified. Therefore, once the surrogate model is constructed it must be validated (sometimes considered a 4th step). This has the purpose of establishing

the predictive capabilities of the surrogate model in design regions away from known sample data. Keane [215] used both empirical and CFD data to create a surrogate for drag variations with gross changes in wing topology. This study suggested updating the initial data set iteratively by adding new points where good designs were found; with this update the initial design from the original response surface was greatly improved.

There are both parametric and non-parametric alternatives in constructing a surrogate model. Parametric approaches (such as kriging or polynomial regression) are model dependant forming a functional relationship between the response variables and the design variable samples that are known. Non-parametric approaches (such as radial-basis functions or neural networks) use local models in different regions of the sample data to build-up an overall framework of the model. Furthermore, surrogates can also be classified into regression type (polynomial regression, radial basis functions) which tend to be better suited to noisy functions, and interpolation type (kriging) creating best-fit response models. Usage of any of these models is not straight forward as the quantity and quality of information the user has to provide in the construction of the surrogate is not known a priori. Furthermore, the efficient exploitation of training data can be restricted by inherent problem complexity, constraints, design variable dimensionality, and accuracy vs. computational cost. Hence selection of the most appropriate surrogate is considered problem dependant, as it will directly influence the optimisation algorithms' decision making capability. There are no general rules leading to the choice of type of surrogate, generation of sample data for training and validation, and indeed the combination of surrogate model and optimisation algorithm. Different surrogate models will be better suited to different data sets and care must be taken to not over-generalise the problem or false optimisation may occur. A review and comprehensive discussion of surrogate modelling is presented by: Queipo et al. [216], Simpson et al. [217], Jin et al. [218], and Forrester and Keane [219].

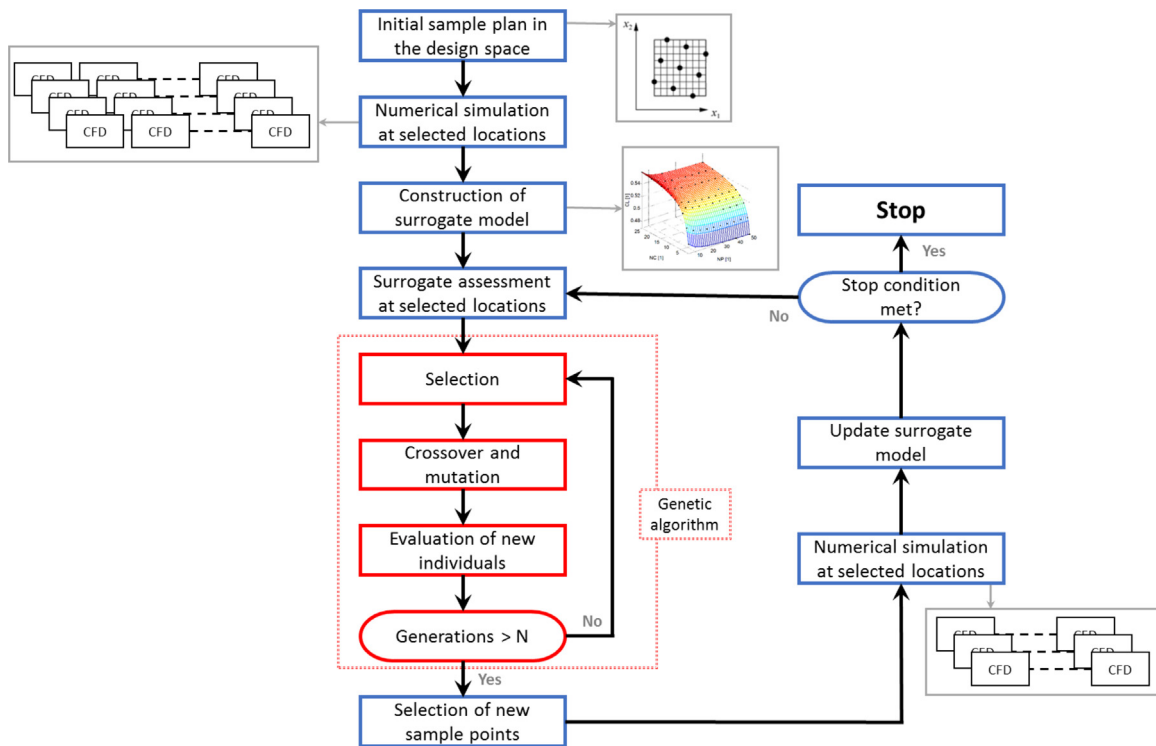


Fig. 26. Example structure for surrogate based optimisation with a standard genetic algorithm.

Parametric surrogates have been widely applied to aerodynamic optimisations due to their flexibility and ease of use. The resulting model will form a response surface that fits exactly to the sample data points. Therefore parametric models are very well suited to conditions where a design space is poorly/sparsely sampled, clustered, or noisy. These approaches do not rely on any specific model structure and are a successful statistical tool for modelling globally dispersed spatial observations. Perhaps the most straight-forward parametric surrogate is one formulated through polynomial regression. Lian and Liou [220] use 2nd-order polynomial response surface equations to construct the functional relationship between design variables and objectives enabling GAs to re-design centrifugal compressors and transonic compressor blades. As the surrogate greatly reduces the computational cost of function evaluations, the GA population and generation sizes were increased due to the freed computational resources. This facilitated increased exploration and identification of superior solutions.

Other studies have used polynomial regression response surfaces for the optimisation of helicopter rotor designs to increase efficiency at a reduced vibration over different operating conditions. Collins et al. [221] report development of these methods and show that 4th-order polynomials, requiring over 300 simulations to construct, are capable of accurately approximating performance metrics achieving a regression coefficient (R^2) of 0.99. This surrogate maintained robustness of the high fidelity numerical simulations at a greatly reduced computational time, however, scaling functions were required to map low-fidelity results to the higher-fidelity domain. Similar approximation models have been used by Leisink et al. [222] for the multi-point aerodynamic design of helicopter blades. The authors consider three design points to be optimised by a GA. In the initial stage of the optimisation, the GA operates at a low-fidelity level in a series of reduced design space. Following this, a high-fidelity surrogate model is created and the optimisation is continued. Their results compare the GA surrogate assisted optimisation with the un-assisted GA, the resulting Pareto optimal fronts are shown in Fig. 27. The multi-point

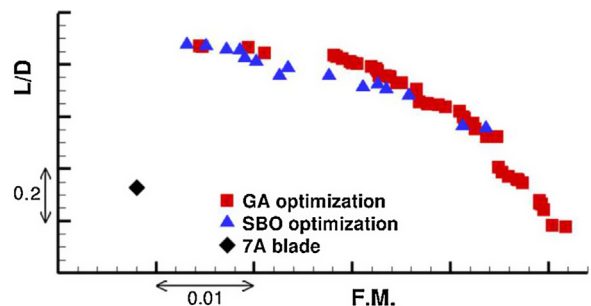


Fig. 27. Pareto optimal fronts for the GA and the surrogate assisted GA compared to the 7A helicopter blade baseline design [222].

optimisation for the surrogate assisted GA required a total of 180 simulations (80 simulation set up the initial surrogate with 5 subsequent updates of additional 20 simulations) while the lone GA required 1600 simulations. Pareto Results in Fig. 27 show reasonable agreement, however the surrogate assisted optimisation does not capture the entire extent of favourable solutions across the front. Hence better trade-off solutions and other favourable configurations that are found by the GA are completely missed by the surrogate optimisation. A common flaw with scaling low-fidelity results to a higher-fidelity domain is that improvement of a design predicted by a low-fidelity model does not guarantee an improvement in the high-fidelity problem [223]. Zadeh et al. [224] present methods for tuning design variables to cope with discrepancies between high and low-fidelity models.

Kriging surrogates (discussed in detail for aerodynamic surrogate modelling by Rosenbaum and Schulz [225]), are perhaps one of the most effective meta-models due to their ability to model complicated responses through either interpolation or regression. Kriging additionally provides an indication of prediction uncertainty, can be used to overcome some of the contrasting performance of low and high fidelity samples. The specific approach

for this is referred to as co-kriging, in which a low fidelity kriging surrogate is used to tune the hyperparameters of a higher fidelity surrogate model. Additionally, for kriging based algorithms, the quantity of initial sample information is independent of the number of design variables typically resulting in efficient algorithms for highly dimensional domains. Toal et al. [211] discuss effective hyper parameter tuning for co-kriging, and consider several methodologies. The co-kriging surrogate developed is later used by Toal et al. [226] for aerodynamic optimisations with a PSO algorithm, with which they also consider the difference in performance of kriging and co-kriging. The co-kriging demonstrates the capability of reducing the number of high-fidelity CFD analysis required by 40% for a multi-point optimisation in a highly dimensional multi-modal design space. This computational saving enabled the consideration of additional design conditions or fewer overall analyses to be carried out. The authors also highlight the necessity for care if the design space contains high non-linearity or discontinuity as more complex response curves will need to be developed. Elsewhere, co-kriging regression surrogates have been shown to utilise wind tunnel to calibrate the meta-model constructed from computational data [227,228].

Laurenceau et al. [229] present a comprehensive study comparing kriging-based and co-kriging-based optimisations for increasingly dimensional and multi-modal aerodynamic problems. They highlight that with increasing dimensionality, the number of computational analyses used to construct the surrogate is likely to become restricted, dependant upon the computational resources available. In limiting the sample size to 200 CFD analyses in a wing-drag optimisation it was found that the kriging based optimisation lacked sufficient accuracy leading to the optimiser wasting computational cost from excessive exploration of the design space. In comparison, the co-kriging-based optimisation performed $\approx 70\%$ less exploration of the design space while enabling a more optimal solution to be identified. Chung and Alonso [230], in the optimisation of a supersonic aircraft, also find that co-kriging-based approximation models are of higher accuracy than kriging-based ones. This again greatly improves the efficiency of the optimisation by reducing computational cost in a large multi-modal design space. Furthermore, Koziel et al. [26] adopt co-kriging for multi-objective aerofoil optimisations to improve the surrogate's flexibility resulting in stronger convergence properties of the optimisation procedure.

There are several examples within the literature which do not adopt multi-fidelity co-kriging models. Takenaka et al. [231] use a kriging assisted genetic algorithm for the design exploration of winglet designs. The selection of the initial 32 sample points for the surrogate model were made by the standard Latin Hypercube Sampling (LHS)[232] model; the initial samples are shown in Fig. 28. Based on this study the kriging surrogate model is very efficient for design problems with multiple conflicting objectives and a small number of design variables; six variables are used here. Their optimised winglet resulted in a total drag reduction of approximately 22 counts for an increase in root bending moment of 5.3%. The total drag reduction provided by the winglet was validated using wind tunnel data.

Hashimoto et al. [28] demonstrated that kriging is capable of fitting complex design spaces and reducing computational demands in the optimisation of high-wing aircraft. A particular trait which makes kriging attractive in complex design optimisations is that statistical improvement criteria, and related uncertainties, are readily available at any point in the design space without additional expense. This enables a more robust exploration of the surrogate by accounting for both the predicted value and its uncertainty simultaneously, transforming the objective function into the corresponding Expected Improvement (EI) function [233]. The EI function indicates the probability of a point being optimum

in the design space and also gives an indication to regions with high uncertainty. Design points with high EIs represent a balance between finding promising regions in the design space based on the surrogate prediction (local search) and finding regions of high uncertainty in the surrogate prediction (global search). Through selecting the best EI points to perform additional numerical simulations, the kriging model can be iteratively re-calibrated, as in Fig. 26, improving the model and searching for superior solution simultaneously. For the optimisation of helicopter blade design to reduce vibrations due to both blade-vortex interaction and dynamic stall, Glaz et al. [234] also use an iteratively updated kriging assisted GA to identify Pareto-optimal designs. The kriging model was capable of capturing fundamental trade-offs among different flight conditions and associated vibrational loading. In this work it is suggested that the application of kriging to aerodynamic optimisations is only appropriate when the time needed to generate the interpolation points is much greater than the time needed to interpolate the data [217].

Non-parametric surrogates are becoming more popular in aerodynamic optimisations, but can be more complex to implement. The increasing popularity is based on their capability to approximate any continuous behaviour with arbitrary precision of the host computing environment. For instance radial basis functions (RBF), although inferior to kriging models regarding interpolating accuracy [235], are easier to characterise and modify, and its superior smoothness can make it more suitable for many design spaces. Morris et al. [56] is seen to use RBF for CFD-based optimisations to provide both a framework to deform the mesh discretising the computational domain and also a way to interpolate forces and moments between them, as the two meshes are likely not coincident. Additionally, they can easily extend from two- to three-dimensional problems. Fincham and Friswell [33] use the RBF in a similar manner to optimise morphing aerofoils decomposing the RBF into three coupled data clusters representing the aerodynamic surface, changes to the aerodynamic surface, and the morphing actuation system. One of the advantages the authors found was that once the initial computation cost of evaluating the design points was established, the computational cost throughout the optimisation was constant (therefore predictable) and significantly reduced.

Bevan et al. [236,237] construct a radial basis function based surrogate model to optimise vortex-generators to alleviate separation and buffet onset for thick and highly loaded aerofoils. The authors state that the advantage of radial basis functions is their ability to fit scattered multivariate data exactly, thus preserving the CFD data at sample points. Lower-fidelity approximations of the higher-fidelity design space were necessary to preserve important flow physics required for the design to increase the accuracy of interpolating between known design points. Ong et al. [238] hybridised a GA and a sequential quadratic programming optimiser. Coupling this method with radial basis functions allowed them to decompose the full-scale problem into a sequence of sub-problems confined to regions defined by the surrogate. The series of local surrogate feed into the global system and is shown capable of coping with fundamental difficulties associated with dimensionality for a single-objective optimisation. The main purpose of constructing the surrogate in this way was to increase the predictability of new design points. In this work it is seen that neighbouring solutions had more influence on the design than remote ones.

Neural networks (multi-layered radial basis functions) form a more sophisticated behavioural modelling technique which may be unnecessary for most aerodynamic optimisations. The persistent difficulties associated with these methods is the need to choose the structure of the surrogate network and a suitable training sequence. Khurana et al. [207] suggest that neural networks work well when the relationship between the design space and the objective function is complex, less intuitive, and highly-dimensional.




Fig. 28. Takenaka et al. [231] 32 initial Latin Hypercube sample sites for the development of a kriging surrogate in the optimisation of a commercial aircraft winglet.

They develop a neural network, on a trial-and-error basis, to duplicate the performance of RANS simulations which a PSO algorithm searches for optimal aerofoil shapes. Massaro et al. [239] coupled a GA with a Neural Network for the multi-point optimisation of helicopter rotor performance. The Neural Network modelled the hover condition for the rotor in addition to the two different flight conditions enabling a crossover of active constraints.

In a situation where one may not know which surrogate model would perform best, the use of multiple or hybrid surrogates may be suitable and can offer several advantages to the optimisation decision making capability [216]. Viana et al. [240] study the application of several open source surrogate modelling tool boxes demonstrating that hybridising up to ten surrogates (one of which kriging) out performed the stand alone kriging surrogate by substantially reducing the number cycles required for convergence. In certain instances however while the number of cycles for the optimisation was reduced the number of functions calls required increased, i.e. computational demand increased as wall-clock time decreased, indicating parallel computation could be beneficial. Additionally results show that the rate of convergence did not scale with the number of surrogates in the hybrid model - 5 and 10 surrogate hybrid models achieved comparable results. The particular aspect of this study attractive to aerodynamic optimisation is that the use of multiple surrogates are shown to have a distinct advantage in highly dimensional design spaces with multiple design targets. Hybrid surrogate modelling under different techniques aid in design exploration and generate diversity as well as improve the surrogates prediction thus reducing the model uncertainty.

Tianyuan and Xiongqing [241] recognised in the aerostructural optimisation of an unmanned combat air vehicle that different output responses did not all coincide with the same point in the search domain. Thus the application of different surrogate models were used to map out different aspects of the response output. This can help minimise each surrogates generalisation error which is used to assess the quality of the approximation model for prediction and establish its capability (validation) for use in analysis and optimisation. A kriging model was found most suitable for predicting the drag coefficient, low radar cross section constraints and the structural weight while second-order polynomial regression models the internal volume of the structure. This implies that the internal volume varied in a much simpler manner with possibly fewer optima and complexity - this makes sense as the internal volume is likely

not to vary too much. The surrogates were constructed from the same 200 samples and are not iteratively updated as the relative error of the surrogates and experiments is less than 5% for the considered design space.

Namura et al. [14] follow similar reasoning and combine the multi-objective GA with two surrogate models in the optimisation of vortex generators for super-critical aerofoils. The first surrogate used a kriging model which was identified as more accurate for modelling the lift coefficient and the chord-wise separation location; the second was a hybrid radial basis function/kriging surrogate showed higher accuracy in estimating the lift-to-drag-ratio. The lift-to-drag ratio is a result of many factors and so can be expected to be complex with regions of discontinuity due to presence of separation. Both kriging and radial basis functions can adapt well to complex spaces, but by hybridising them the resulting approximation surface will be better for capturing both macro and micro design trends [242,243].

7. Conclusion

A survey of several important aerodynamic optimisation strategies are reported here. There is however, significant scope for further work to be carried out on the optimisation algorithms discussed and the strategies used to integrate them with aerodynamic based problems.

It is common knowledge that appropriateness of geometric parametrisation, problem definition and optimisation algorithm definitively depends on the nature of the problem at hand; i.e. number of variables, their scope, required fidelity, and ultimately how the designer embeds and tunes the algorithm. Although relative cost is only one of several important considerations in choosing an optimisation algorithm, for an effective aerodynamic optimisation process deep understanding and consideration must be given to: 1) the level of information fidelity required from the flow solver, dependant on the type of problem; 2) scope of parametrised design space; 3) types of design variables, e.g. discrete and/or continuous; 4) single or multi-objective optimisation; 5) constraints handling; 6) properties of the design space, e.g. number of local optima, discontinuities.

Gradient-based approaches using the discrete adjoint sensitivity analysis scales approximately linearly with the number of

design variables and is very capable of handling thousands of design variables and constraints. They inherently require a geometrically compact set of design variables and an initial user defined baseline geometry and enable systematic design capability in which high-fidelity simulations may be more useful if the analytical properties of the design space are continuous and exhibit low-modality. The gradient-based approach may be more appropriate for detailed aerodynamic design as they are only capable of offering a narrow range of solutions. The cost of gradient-free relative to gradient-based methods increases dramatically and they often require tighter convergence tolerances to be prescribed as they can take advantage of loose constraint definitions. In aerodynamic optimisations, overcoming the computational cost related to these algorithms often use methods that facilitate the progressive containment of shape parametrisation to refine the search space; this tries to reduce the number of wasteful numerical simulations. The benefit of these algorithms is in their simplicity and their parallel nature which is capable of trade-off multi-disciplinary performance metrics over a range of solutions.

Hybrid methods have become more popular in aerodynamic optimisation. Such hybridisations can be used to take advantage of the explicit strengths of certain algorithms in order to account for fundamental difficulties of other algorithm; this improving the overall algorithm performance for more effective and efficient problem solving. For aerodynamic optimisation hybrid algorithms in some instances have been found capable, but for most applications are found to over complicate the optimisation process. From a practical perspective in aerodynamic optimisation, much work remains to be done in benchmarking existing stand-alone architectures to conclude decisively if hybrid methods enhance optimisation capability or just create more complicated optimisation infrastructures.

In many circumstances the relationship between the design parameters and the objective function is highly non-linear, however the objective function relationship with the design variables if often found to follow predictable trends. Consequently, surrogates can help in understanding this relationship which moves the optimisation towards understanding the behaviour of different search components. This ultimately can lead to a design structure that incorporates more problem-specific knowledge and helps to overcome dependency on the user.

For all aerodynamic optimisations regardless of the overall architecture there are several general issues which must be addressed in-order to exploit the optimisation algorithm to its full potential. Despite the widespread success of many algorithms in different contexts, there will always be the persistent question related to the usefulness of a particular algorithm solving for a wide range of problems. The *No Free Lunch Theorem* for optimisation by Wolpert and Macready [244] states that the ultimate optimisation algorithm does not exist, and that all optimisation algorithms have the same average performance over a set of optimisation problems; this basically suggesting that different algorithms are better than others for particular classes of problems.

Finally, it is important to highlight that other optimisation schemes are available than those presented in this paper: they may not yet be popularised, to the authors knowledge, into geometry dependant aerodynamic optimisation design problems (e.g. Mixed Integer Optimisation or gravity flow optimisation). The architectures in this literature survey represent significant progress in aerodynamic optimisation, which has been an ongoing endeavour for over 60 years. Major advances have been found from developments in computational resources and high-fidelity modelling. Newest algorithms and implementation strategies are often decades old, where new algorithms are simply permutations of older concepts.

Acknowledgements

The authors would like to express our gratitude and appreciation to the low speed National Wind Tunnel Facility (EPSRC grant number: EP/L024888/1) at the University of Glasgow. Great appreciation must also be registered for Dr. Euan McGookin, whose support and technical expertise have been invaluable.

Additionally, the authors would like to thank Zephyr Business Services and the Royal Aeronautical Society Centennial Scholarship for funding this work.

References

- [1] European Commission. Flightpath 2050 Europe's Vision for Aviation.
- [2] J.E. Green, Civil aviation and the environment – the next frontier for the aerodynamicist, *Aeronaut. J.* 110 (1110) (2006) 469–486.
- [3] J. Nocedal, S.J. Wright, S.M. Robinson, *Numerical Optimization*, Springer, New York, 1999.
- [4] G. Anderson, M. Aftosmis, Adaptive Shape Parameterization for Aerodynamic Design, in: NASA Technical Report: NAS-2015-02, (May), 2015.
- [5] S. Wakayama, I. Kroo, Subsonic wing planform design using multidisciplinary optimization, *J. Aircraft* 32 (1995) 746–753.
- [6] N.P.M. Werter, R.D. Breuker, Aeroelastic tailoring and structural optimization using an advanced dynamic aeroelastic framework, in: International Forum on Aeroelasticity and Structural Dynamics, Saint Petersburg, Russia, 2015, pp. 1–20.
- [7] J. Sobieszczanski-Sobieski, R.T. Haftka, Multidisciplinary aerospace design optimization: survey of recent developments, in: 34th Aerospace Sciences Meeting and Exhibit, AIAA 96-0711, volume 70, Reno, Nevada, 1996, pp. 1–33.
- [8] T.F. Wunderlich, Multidisciplinary wing optimization of commercial aircraft with consideration of static aeroelasticity, *CEAS Aeronaut. J.* 6 (3) (2015) 407–427.
- [9] D.W. Levy, K.R. Laffin, E.N. Tinoco, J.C. Vassberg, M. Mani, B. Rider, C. Rumsey, R.A. Wahls, J.H. Morrison, O.P. Broderson, S. Crippa, D.J. Mavriplis, M. Murayama, Summary of data from the first AIAA CFD drag prediction workshop, in: 51th AIAA Aerospace Sciences Meeting, 2013, pp. 1–31.
- [10] C.J. Roy, E.N. Tinoco, Summary of data from the sixth AIAA CFD drag prediction workshop: case 1 code verification, in: 55th AIAA Aerospace Sciences Meeting, (January), 2017, pp. 1–19.
- [11] W. Yamazaki, K. Matsushima, K. Nakahashi, Aerodynamic design optimization using the drag-decomposition method, *AIAA J.* 46 (5) (2008) 1096–1106.
- [12] Z. Lyu, G. Kenway, J.R.R.A. Martins, Aerodynamic shape optimization investigations of the common research model wing benchmark, *AIAA J.* 53 (4) (2014) 968–985.
- [13] G. Kenway, J.R.R.A. Martins, Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration, *J. Aircraft* 51 (1) (2014) 144–160.
- [14] N. Namura, S. Obayashi, S. Jeong, Efficient global optimization of vortex generators on a super critical infinite-wing using kriging-based surrogate models., in: 52nd Aerospace Sciences Meeting, AIAA 2014-0904, 2014, pp. 1–12.
- [15] G.R. Anderson, M. Nemec, M.J. Aftosmis, Aerodynamic shape optimization benchmarks with error control and automatic parameterization., in: 53rd AIAA Aerospace Sciences Meeting, AIAA 2015-1719, 2015, pp. 1–18.
- [16] L.M. Rios, N.V. Sahinidis, Derivative-free optimization: a review of algorithms and comparison of software implementations, *J. Glob. Optim.* 56 (3) (2013) 1247–1293.
- [17] S. Amaran, N.V. Sahinidis, B. Sharda, S.J. Bury, Simulation optimization: a review of algorithms and applications, *Ann. Oper. Res.* 240 (1) (2016) 351–380.
- [18] D.W. Zingg, M. Nemec, T.H. Pulliam, A comparative evaluation of genetic and gradient-based algorithms applied to aerodynamic optimization, *Eur. J. Comput. Mech.* 17 (1–2) (2008) 103–126.
- [19] G.N. Vanderplaats, *Numerical Optimization Techniques for Engineering Design: With Applications*, McGraw-Hill Book Company, 1999.
- [20] Y. Zhang, Z.H. Han, L. Shi, W.P. Song, Multi-round surrogate-based optimization for benchmark aerodynamic design problems, in: 54th AIAA Aerospace Sciences Meeting, AIAA 2016-1545, number 127, San Diego, California, 2016, pp. 1–17.
- [21] J. Sobieszczanski-Sobieski, A linear decomposition method for large optimisation problems – blueprint for development, in: NASA TM-83248, 1982.
- [22] J. Sobieszczanski-Sobieski, Optimisation by decomposition: a step from hierarchic to non-hierarchic systems, in: NASA Technical Report: N89-25149, 1988.
- [23] O. Chernukhin, D.W. Zingg, Multimodality and global optimization in aerodynamic design, *AIAA J.* 51 (6) (2013) 1342–1354.
- [24] P.W. Jansen, R.E. Perez, J.R.R.A. Martins, Aerostructural optimization of nonplanar lifting surfaces, *J. Aircraft* 47 (5) (sep 2010) 1490–1503.

- [25] S. Peigin, B. Epstein, Multiconstrained aerodynamic design of business jet by CFD driven optimization tool, *Aerosp. Sci. Technol.* 12 (March (2)) (2008) 125–134.
- [26] S. Koziel, Y. Tesfahunegn, A. Amrit, L.T. Leifsson, Rapid multi-objective aerodynamic design using co-kriging and space mapping, in: 57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA 2016-0418, number January, San Antonio, TX, USA, 2016, pp. 1–10.
- [27] P.D. Vecchia, F. Nicolosi, Aerodynamic guidelines in the design and optimization of new regional turboprop aircraft, *Aerosp. Sci. Technol.* 38 (October) (2014) 88–104.
- [28] A.H. Hashimoto, S.J. Jeong, S.O. Obayashi, Aerodynamic optimization of near-future high-wing aircraft, *Trans. Jpn. Soc. Aeronaut. Space Sci.* 58 (2) (2015) 73–82.
- [29] M. Athanasopoulos, H. Ugail, G.G. Castro, Parametric design of aircraft geometry using partial differential equations, *Adv. Eng. Softw.* 40 (7) (2009) 479–486.
- [30] N.E. Sevant, M.J. Bloor, M.J. Wilson, Aerodynamic design of a flying wing using response surface methodology, *J. Aircraft* 37 (4) (2000) 562–569.
- [31] A. Jameson, L. Martinelli, N.A. Pierce, Optimum aerodynamic design using the Navier-Stokes equations, *Theor. Comput. Fluid Dynam.* 10 (1–4) (1998) 213–237.
- [32] G. Kenway, G. Kennedy, J.R.R.A. Martins, A CAD-free approach to high-fidelity aerostructural optimization, in: Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, AIAA 2010-9231, Fort Worth, TX, 2010, pp. 1–18.
- [33] J.H.S. Fincham, M.I. Friswell, Aerodynamic optimisation of a camber morphing aerofoil, *Aerosp. Sci. Technol.* 43 (2015) 245–255.
- [34] J. Hall, D.J. Poole, T.C.S. Rendall, C.B. Allen, Volumetric shape parameterisation for combined aerodynamic geometry and topology optimisation, in: 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA 2015-3354, number June, Dallas, TX, 2015, pp. 1–29.
- [35] S. Koziel, Z. Michalewicz, Evolutionary algorithms, homomorphous mappings, and constrained parameter optimization, *Evol. Comput.* 7 (1) (1999) 19–44.
- [36] J. Siens, M.S. Innocente, Particle swarm optimisation: fundamental study and its application to optimisation and to jetty scheduling problems, *Trends Eng. Comput. Technol.* 10 (2008) 3–126.
- [37] K. Deb, L. Thiele, M. Laumanns, E. Zitzler, Scalable multi-objective optimization test problems, *CEC-2002 2* (2002) 825–830.
- [38] T. Kato, K. Shimoyama, S. Obayashi, Evolutionary algorithm with parallel evaluation strategy of feasible and infeasible solutions considering total constraint violation, *IEEE 1* (978) (2015) 986–993.
- [39] T.D. Robinson, K.E. Willcox, M.S. Eldred, R. Haimes, Multifidelity optimization for variable-complexity design, in: Proceedings of the 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Portsmouth, VA, 2006, pp. 1–18, AIAA 2006-7114.
- [40] S. Gemma, F. Mastroddi, Multi-disciplinary and multi-objective optimization of an unconventional aircraft concept, in: 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA 2015-2327, Dallas, TX, 2015, pp. 1–20.
- [41] M.J.D. Powell, A direct search optimization method that models the objective and constraint functions by linear interpolation, *Adv. Optim. Numer. Anal.* 7 (1994) 51–67.
- [42] M.D.J. Powell, Direct search algorithms for optimization, *Acta Numer.* 7 (1998) 287–336.
- [43] A. March, K. Willcox, Constrained multifidelity optimization using model calibration, *Struct. Multidiscip. Optim.* 46 (2012) 93–109.
- [44] D.R. Jones, C.D. Perttunen, B.E. Stuckman, Lipschitzian optimization without the lipschitz constant, *J. Optim. Theory Appl.* 79 (1) (1993) 157–181.
- [45] D.R. Jones, *Encyclopedia of Optimization*, 2001.
- [46] K.W. Lee, W.H. Moase, S.Z. Khong, Aerodynamic shape optimization via global extremum seeking, *IEEE transactions on control system technology* (2015) 1–8.
- [47] F. Deng, Z. Jiao, J. Chen, Overall performance analysis – oriented aerodynamic configuration optimization design for hypersonic vehicles, *J. Spacecraft Rockets* (2017) 1–12.
- [48] S. Ebastien, L.E. Digabel, Algorithm 909: NOMAD: nonlinear optimization with the MADS algorithm, *ACM Trans. Math. Softw.* 37 (4) (2011) 1–15.
- [49] M.D. Ardema, J.V. Bowles, T. Whittaker, Near-optimal propulsion-system operation for an air-breathing launch vehicle, *J. Spacecraft Rockets* 32 (6) (1995) 951–956.
- [50] Y. Liu, L. Zhou, Y. Lu, Research on multidisciplinary analysis tool based on parametric model of hypersonic morphing waverider, *Appl. Math. Inform. Sci.* 9 (4) (2015) 2047–2056.
- [51] R. Hassan, B. Cohanin, O. Weck, G. Venter, A comparison of particle swarm optimisation and the genetic algorithm, in: 46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, AIAA 2005-1897, 2005, pp. 1–13.
- [52] P. Dellla Vecchia, Development of Methodologies for the Aerodynamic Design and Optimization of New Regional Turboprop Aircraft, University of Naples, 2013, PhD thesis.
- [53] J. Ren, A.S. Thelen, A. Amrit, X. Du, L.T. Leifsson, Y. Tesfahunegn, S. Koziel, Application of multifidelity optimization techniques to benchmark aerodynamic design problems, in: 54th AIAA Aerospace Sciences Meeting, (January), 2016.
- [54] Z. Zhang, P.C. Chen, Z. Zhou, S. Yang, Z. Wang, Q. Wang, Adjoint based structure and shape optimization with flutter constraints, in: 57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA 2016-1176, San Diego, California, 2016, pp. 1–23.
- [55] M. Nemeč, M.J. Aftosmis, Output error estimates and mesh refinement in aerodynamic shape optimization, in: 51st AIAA Aerospace Sciences Meeting, Grapevine, TX, 2013, pp. 1–18.
- [56] A.M. Morris, C.B. Allen, T.C.S. Rendall, CFD-based optimization of aerofoils using radial basis functions for domain element parameterization and mesh deformation, in: *Int. J. Numer. Methods Fluids*, (October 2007), 2008, pp. 601–629.
- [57] J.E. Hicken, D.W. Zingg, Aerodynamic optimization algorithm with integrated geometry parameterization and mesh movement, *AIAA J.* 48 (2) (2010) 400–413.
- [58] D. Li, R. Hartmann, Adjoint-based airfoil optimisation with discretization error control, *Int. J. Numer. Methods Fluids* 77 (2015) 1–17.
- [59] K. Chiba, A. Oyama, S. Obayashi, K. Nakahashi, H. Morino, Multidisciplinary design optimization and data mining for transonic regional-jet wing, *J. Aircraft* 44 (4) (2007) 1100–1112.
- [60] M. Murayama, K. Nakahashi, K. Matsushima, Unstructured dynamic mesh for large movement and deformation, in: AIAA Paper 2002-0122, 2002.
- [61] J. Li, Z. Gao, J. Huang, K. Zhao, Aerodynamic design optimization of nacelle/pylon position on an aircraft, *Chin. J. Aeronaut.* 26 (4) (2013) 850–857.
- [62] Z. Lyu, J.R.R.A. Martins, Aerodynamic design optimization studies of a blended-wing-body aircraft, *J. Aircraft* 51 (5) (2014) 1604–1617.
- [63] T.R. Brooks, G. Kennedy, J.R.R.A. Martins, High-fidelity aerostructural optimization of a high aspect ratio tow-steered wing, in: 57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA 2016-1179, (January), 2016, pp. 1–16.
- [64] D.A. Burdette, G. Kenway, J.R.R.A. Martins, Performance evaluation of a morphing trailing edge using multipoint aerostructural design optimization, in: 57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA 2016-0159, San Diego, California, 2016, pp. 1–14.
- [65] M. Biava, M. Carrion, G.N. Barakos, D. Stewart, Adjoint-based optimisation of ducted propellers for hybrid air vehicles, in: 54th AIAA Aerospace Sciences Meeting, AIAA 2016-1784, San Diego, California, 2016, pp. 1–13.
- [66] S.A. Ning, I. Kroo, Multidisciplinary considerations in the design of wings and wing tip devices, *J. Aircraft* 47 (2) (2010) 534–543.
- [67] S.N. Skinner, H. Zare-Behtash, Aerodynamic optimisation of non-planar lifting surfaces, in: 57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA 2016-0164, San Diego, California, 2016, pp. 1–21.
- [68] D. Goldberg, *Genetic Algorithms in Searching Optimisation and Machine Learning*, Addison-Wesley, Reading, MA, 1989.
- [69] N.F. Foster, G.S. Dulikravich, Three-dimensional aerodynamic shape optimisation using genetic and gradient search algorithms, *J. Spacecraft Rockets* 34 (1) (1997).
- [70] M. Secanell, A. Suleman, Numerical evaluation of optimization algorithms for low-reynolds-number aerodynamic shape optimization, *AIAA J.* 43 (10) (2005) 2262–2267.
- [71] Z. Lyu, Z. Xu, J.R.R.A. Martins, Benchmarking optimization algorithms for wing aerodynamic design optimization, in: The Eighth International Conference on Computational Fluid Dynamics, ICCFD8-2014-0203, Chengdu, Sichuan, China, 2014, pp. 1–18.
- [72] J.R.R.A. Martins, J.J. Alonso, J.J. Reuther, Complete Configuration aero-structural optimization using a coupled sensitivity analysis method complete configuration aero-structural optimization using a coupled sensitivity analysis method, in: 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimisation Conference, AIAA Paper 2002-5402, Stanford, CA, 2002, pp. 1–14.
- [73] J.E.V. Peter, R.P. Dwight, Numerical sensitivity analysis for aerodynamic optimization: a survey of approaches, *Comput. Fluids* 39 (3) (2010) 373–391.
- [74] J.R.R.A. Martins, J.T. Hwang, Review and unification of discrete methods for computing derivatives of single- and multi-disciplinary computational models, *AIAA J.* 51 (11) (2013) 2582–2599.
- [75] J.R.R.A. Martins, P. Sturdza, J.J. Alonso, The complex-step derivative approximation, *ACM Trans. Math. Softw.* 29 (3) (2003) 245–262.
- [76] G. Kenway, J.R.R.A. Martins, Aerostructural shape optimization of wind turbine blades considering site-specific winds, in: 12th AIAA/ISSMO Multidisciplinary Analysis and Optimisation Conference, AIAA 2008-6025, Victoria, British Columbia, Canada, 2008, pp. 1–12.
- [77] P.E. Gill, W. Murray, M.A. Saunders, SNOPT: an SQP algorithm for large-scale constrained optimization, *SIAM J. Optim.* 47 (1) (2005) 99–131.
- [78] O. Pironneau, On optimum design in fluid mechanics, *J. Fluid Mech.* 64 (1) (1974) 97–110.
- [79] D. Kraft, A software package for sequential quadratic programming, in: Technical report, Tech. Re DFVLR-FB 88-28, German Aerospace Center, 1988.
- [80] D. Kraft, Algorithm 733: TOMP-fortran modules for optimal control calculations, *ACM Trans. Math. Softw. (TOMS)* 20 (3) (1994) 262–281.
- [81] P.D. Frank, G.R. Shubin, A comparison of optimization-based approaches for a model computational aerodynamics design problem, *J. Comput. Phys.* 98 (1) (1992) 74–89.

- [82] A. Jameson, N.A. Pierce, L. Martinelli, Optimum aerodynamic design using the Navier-Stokes equations, in: 35th Aerospace Sciences Meeting and Exhibit, AIAA Paper 97-0101, number January, 1997.
- [83] C.A. Mader, J.R.R.A. Martins, J.J. Alonso, E.V. Der Weide, ADjoint: an approach for the rapid development of discrete adjoint solvers, *AIAA J.* 46 (4) (2008) 863–873.
- [84] Z. Lyu, G. Kenway, C. Paige, J.R.R.A. Martins, Automatic differentiation adjoint of the Reynolds-averaged Navier-Stokes equations with a turbulence model, in: 21st AIAA Computational Fluid Dynamics Conference, AIAA 2013-2581, San Diego, California, 2013, pp. 1–24.
- [85] J.E. Hicken, D.W. Zingg, Induced-Drage Minimization of nonplanar geometries based on the Euler equations, *AIAA J.* 48 (11) (2010) 2564–2575.
- [86] H. Gagnon, D.W. Zingg, High-fidelity aerodynamic shape optimization of unconventional aircraft through axial deformation, in: 52nd Aerospace Sciences Meeting, AIAA 2014-0908, National Harbor, Maryland, 2014, pp. 1–18.
- [87] J.R.R.A. Martins, J.J. Alonso, J.J. Reuther, A coupled-adjoint sensitivity analysis method for high-fidelity aero-structural design, *Optim. Eng.* 6 (1) (2005) 33–62.
- [88] G. Kenway, G.J. Kennedy, J.R.R.A. Martins, Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations, *AIAA J.* 52 (5) (2014) 935–951.
- [89] G. Kennedy, J.R.R.A. Martins, High aspect ratio wing design: optimal aerostructural tradeoffs for the next generation of materials, in: 52nd Aerospace Sciences Meeting, 2014, pp. 1–24.
- [90] D.A. Burdette, G. Kenway, Z. Lyu, J.R.R.A. Martins, Aerostructural design optimization of an adaptive morphing trailing edge wing, in: 15th AIAA/ISSMO Multidisciplinary Analysis and Optimisation Conference, AIAA paper 2014-3275, 2014, pp. 1–13.
- [91] A. Elham, M.J.L. van Tooren, Coupled adjoint aerostructural wing optimization using quasi-three-dimensional aerodynamic analysis, in: 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA 2015-2487, Dallas, TX, 2015, pp. 1–20.
- [92] C.A. Mader, J.R.R.A. Martins, Stability-constrained aerodynamic shape optimization of flying wings, *J. Aircraft* 50 (5) (2013) 1431–1449.
- [93] B. Grossman, Z. Gurdal, G.J. Strauch, W.M. Eppard, R.T. Haftka, Integrated aerodynamic/structural design of a sailplane wing, *J. Aircraft* 25 (9) (1988) 855–860.
- [94] B. Grossman, R.T. Haftka, P.J. Kao, D.M. Polen, M. Rais-Rohani, Integrated aerodynamic-structural design of a transport wing, *J. Aircraft* 27 (12) (1990) 1050–1056.
- [95] I.R. Chittick, J.R.R.A. Martins, An asymmetric suboptimization approach to aerostructural optimization, *Optim. Eng.* 10 (1) (2009) 133–152.
- [96] I. Kroo, Non-planar wing concepts for increased aircraft efficiency, in: VKI Lecture Series on Innovative Configurations and Advanced Concepts for Future Civil Aircraft, 2005, pp. 1–29.
- [97] S.J. Miller, Adaptive Wing Structures For Aeroelastic Drag Reduction and Loads Alleviation, University of Manchester, 2010, PhD thesis.
- [98] L.A. Catalano, D. Quagliarella, P.L. Vitagliano, Aerodynamic shape design using hybrid evolutionary computing and multigrid-aided finite-difference evaluation of flow sensitivities, *Int. J. Comput. Aided Eng. Softw.* 32 (2) (2014) 178–210.
- [99] G. Kenway, J.R.R.A. Martins, Aerodynamic shape optimization of the CRM configuration including buffet-onset conditions, in: 54th AIAA Aerospace Sciences Meeting, AIAA 2016-1294, San Diego, California, 2016, pp. 1–23.
- [100] M. Nemeč, D.W. Zingg, T.H. Pulliam, Multi-point and multi-objective aerodynamic shape optimisation, in: 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimisation Conference, AIAA Paper 2002-5548, Atlanta, Georgia, 2002.
- [101] S. Peigin, B. Epstein, Multipoint aerodynamic design of wing-body configurations for minimum drag, *J. Aircraft* 44 (3) (2007) 971–980.
- [102] G. Venter, Particle swarm optimization, in: 43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA 2002-1235, Denver, Colorado, 2002, pp. 1–9.
- [103] S.N. Sivanandam, S.N. Deepa, Introduction to Genetic Algorithms, Springer, 2008.
- [104] O. Castillo, P. Melin, J. Kacprzyk, Soft Computing for Hybrid Intelligent Systems, Springer, Berlin, Germany, 2008.
- [105] K. Chiba, S. Obayashi, K. Nakahashi, High-fidelity multidisciplinary design optimization of aerostructural wing shape for regional jet, in: 23rd AIAA Applied Aerodynamics Conference, AIAA 2005-5080, Toronto, Canada, 2005, pp. 621–635.
- [106] Z. Jinhuan, Z. Zhenggui, An improved genetic algorithm and its applications to the optimisation design of an aspirated compressor profile, *Int. J. Numer. Methods Fluids* (2015).
- [107] R.L. Haupt, S.E. Haupt, Practical Genetic Algorithms, 2nd ed., Wiley-Interscience, 2004.
- [108] M. Mitchell, An Introduction to Genetic Algorithms, first edition, MIT Press, 1998.
- [109] K.J. Rafferty, A Comparison Study of Search Heuristics for an Autonomous Multi-Vehicle Air-Sea Rescue System, University of Glasgow, 2014, PhD thesis.
- [110] K.G. Khoo, P.N. Suganthan, Evaluation of genetic operators and solution representations for shape recognition by genetic algorithms, *Pattern Recogn. Lett.* 23 (13) (2002) 1589–1597.
- [111] E.W. Mcgookin, Optimisation of Sliding Mode Controllers for Marine Applications: A Study of Methods and Implementation Issues, University of Glasgow, 1997, PhD thesis.
- [112] C.R. Reeves, Using genetic algorithms with small populations, in: Proceedings of the 5th International Conference on Genetic Algorithms, 1993, pp. 92–99.
- [113] O. Roeva, Improvement of genetic algorithm performance for identification of cultivation process models, in: 9th WSEAS International Conference on Evolutionary Computing, Sofia, Bulgaria, 2008, pp. 34–39.
- [114] O. Roeva, S. Fidanova, M. Paprzycki, Influence of the population size on the genetic algorithm performance in case of cultivation process modelling, in: Proceedings of the 2013 Federated Conference on Computer Science and Information Systems, 2013, pp. 371–376.
- [115] H.M. Pandey, A. Chaudhary, D. Mehrotra, A comparative review of approaches to prevent premature convergence in GA, *Appl. Soft Comput.* J. 24 (2014) 1047–1077.
- [116] M.E. Alfaro Cid, Optimisation of Time Domain Controllers For Supply Ships Using Genetic Algorithms and Genetic Programming, University of Glasgow, 2003, PhD thesis, Ph.D Thesis.
- [117] B.L. Miller, D.E. Goldberg, Genetic algorithms, tournament selection, and the effect of noise, *Complex Syst.* 9 (1995).
- [118] K.A. De Jong, W.M. Spears, A formal analysis of the role of multi-point crossover in genetic algorithms, *Ann. Math. Artif. Intell.* 5 (1) (1992) 1–26.
- [119] L.M. Schmitt, Theory of Genetic Algorithms II: models for genetic operators over the string-tensor representation of populations and convergence to global optima for arbitrary fitness function under scaling, *Theor. Comput. Sci.* 310 (1–3) (2004) 181–231.
- [120] K. Deb, Multi-objective Optimization using Evolutionary Algorithms, 2001.
- [121] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Trans. Evol. Comput.* 6 (2) (2002) 182–197.
- [122] J. Kim, B. Ovgor, K. Cha, J. Kim, S. Lee, K. Kim, Optimization of the aerodynamic and aeroacoustic performance of an axial-flow fan, *AIAA J.* 52 (9) (2014) 2032–2044.
- [123] G. Droandi, G. Gibertini, Aerodynamic blade design with multi-objective optimization for a tiltrotor aircraft, *Aircraft Eng. Aerosp. Technol.* 87 (2015) 19–29.
- [124] J.H. Holland, Adaption in Natural and Artificial Systems, The University of Michigan Press, 1975.
- [125] A. Kumar, Encoding schemes in genetic algorithm, *Int. J. Adv. Res. IT Eng.* 2 (3) (2013) 1–7.
- [126] G. Guglieri, Using of particle swarm for performance optimization of helicopter rotor blades, *Appl. Math.* 3 (2012) 1403–1408.
- [127] Z. Michalewicz, Genetic Algorithms + Data Structures = Evolution Programs, 3rd edition, Springer, 1995.
- [128] R. Poli, W.B. Langdon, N.F. Mcphee, A Field Guide to Genetic Programming, 2008, Number March.
- [129] M. Gen, R. Cheng, Genetic Algorithms and Engineering Optimisation, John Wiley & Sons, Inc., 2000.
- [130] T.Y. Lim, Structured population genetic algorithms: a literature survey, *Artif. Intell. Rev.* 41 (3) (feb 2012) 385–399.
- [131] D. Popvic, Retaining diversity of search point distribution through abreeder genetic algorithm for neural network learning, *IEEE Int. Conf. Neural Netw.* 49 (1997) 5–498.
- [132] C. Fernandes, A. Rosa, A study on non-random mating and varying population size in genetic algorithms using a royal road function, in: Proceedings of the 2001 Congress on Evolutionary Computation (IEEE Cat. No.01TH8546), 1, 2001, pp. 60–66.
- [133] A. Massaro, E. Benini, A surrogate-assisted evolutionary algorithm based on the genetic diversity objective, *Appl. Soft Comput.* 36 (2015) 87–100.
- [134] E. Benini, Three-dimensional multi-objective design optimization of a transonic compressor rotor, *J. Propul. Power* 20 (3) (2004) 559–565.
- [135] D.J. Zhao, Y.K. Wang, W.W. Cao, P. Zhou, Optimization of suction control on an airfoil using multi-island genetic algorithm, *Procedia Eng.* 99 (2015) 696–702.
- [136] A. Toffolo, E. Benini, Genetic diversity as an objective in multi-objective evolutionary algorithms, *Evol. Comput.* 11 (2) (2003) 151–167.
- [137] D. Dasgupta, D.R. Mcgregor, sGA: A Structured Genetic Algorithm. Technical report, University of Strathclyde, Glasgow, U.K.
- [138] B. Artushenko, Analysis of global exploration of island model genetic algorithm, in: CADSM 2009. 10th International Conference, 2, 2009, pp. 280–281.
- [139] J. Zhang, L. Xu, R. Gao, Multi-island genetic algorithm optimization of suspension system, *Telkomnika* 10 (7) (2012) 1685–1691.
- [140] L. Eshelman, J. Schaffer, Real-coded genetic algorithms and interval-schemata, in: Foundations of Genetic Algorithms, vol. 2, 1993, pp. 187–202.
- [141] W.T. McCormick, P.J. Schweitzer, T.W. White, W.T. McCormick, Problem decomposition and data reorganization by a clustering technique, *Oper. Res.* 20 (5) (1972) 993–1009.
- [142] G.A. Walters, D.K. Smith, Evolutionary design algorithm for optimal layout of tree networks, *Eng. Optim.* 24 (1995) 261–281.
- [143] J.M. Badyrka, R.M. Jenkins, R.J. Hartfield, Aerospace design: a comparative study of optimizers, in: 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, AIAA 2010-1311, Orlando, Florida, 2010.

- [144] A. Oyama, S. Obayashi, K. Nakahashi, Real-coded adaptive range genetic algorithm and its application to aerodynamic design, *JSME Int. J. Ser. A* 43 (2) (2000) 124–129.
- [145] A. Oyama, M. Liou, Transonic axial-flow blade optimization: evolutionary algorithms / three-dimensional Navier-Stokes solver, *J. Propul. Power* 20 (4) (2004) 612–619.
- [146] D.R. Jones, Large-scales multi-disciplinary mass optimization in the auto industry, in: *MOPTA08 Modelling and Optimization: Theory and Applications*, number August, Ontario, Canada, 2008.
- [147] M. Olhofer, Y. Jin, B. Sendhoff, Adaptive encoding for aerodynamic shape optimization using evolution strategies, in: *Proceedings of the 2001 Congress on Evolutionary Computation*, 2001, pp. 576–583.
- [148] J. Periaux, D.S. Lee, L.F. Gonzalez, K. Srinivas, Fast reconstruction of aerodynamic shapes using evolutionary algorithms and virtual nash strategies in a CFD design environment, *J. Comput. Appl. Math.* 232 (2009) 61–71.
- [149] D. Sasuki, S. Obayashi, Efficient search for trade-offs by adaptive range multi-objective genetic algorithms, *J. Aerosp. Comput. Inf. Commun.* 2 (April) (2005) 44–64.
- [150] K. Krishnakumar, Micro-genetic algorithms for stationary and non-stationary function optimisation, in: *Proceedings of Intelligent Control and Adaptive Systems*, 1989, pp. 289–296.
- [151] D. Sasaki, M. Morikawa, S. Obayashi, K. Nakahashi, Aerodynamic shape optimization of supersonic wings by adaptive range multiobjective genetic algorithms, in: *First International Conference on Evolutionary Multi-Criterion Optimization*, 2001, pp. 639–652.
- [152] D. Sasaki, S. Obayashi, H. Kim, Evolutionary algorithm vs. adjoint method applied to SST shape optimization, in: *Proceedings of the Ninth Annual Conference of the CFD Society of Canada*, Waterloo, 2001.
- [153] J.F. Wang, J. Periaux, M. Sefrioui, Parallel evolutionary algorithms for optimization problems in aerospace engineering, *J. Comput. Appl. Math.* 149 (1) (2002) 155–169.
- [154] D. Tse, L. Chan, Transonic airfoil design optimization using soft computing methods, *Canad. Aeronaut. Space J.* 46 (2000) 65–73.
- [155] T.L. Holst, T.H. Pilliam, Aerodynamic shape optimisation using real-number-encoded genetic algorithm, in: *19th Applied Aerodynamics Conference*, AIAA Paper 2001-2473, Anaheim, CA, 2001.
- [156] R. Eberhart, J. Kennedy, A new optimizer using particle swarm theory, in: *Proceedings of the Sixth International Symposium on Micro Machine and Human Science*, IEEE, 1995, pp. 39–43.
- [157] J. Kennedy, R.C. Eberhart, Particle swarm optimization, in: *Proceedings, IEEE International Conference on Neural Networks*, Perth, Australia, 1995, pp. 1942–1948.
- [158] Y. Wang, B. Zhang, Y. Chen, Robust airfoil optimization based on improved particle swarm optimization method, *Appl. Math. Mech.* 32 (10) (2011) 1245–1254.
- [159] Y. Zhang, S. Wang, G. Ji, A comprehensive survey on particle swarm optimization algorithm and its applications, *Math. Probl. Eng.* 2015 (2015) 1–38.
- [160] J. Kennedy, R.C. Eberhart, *Swarm Intelligence*, 1st edition, Morgan Kaufmann Publishers, 2001.
- [161] A. Campos, A. Pozo, E.P. Duarte, Evaluation of asynchronous multi-swarm particle optimisation on several topologies, *Concurr. Comput. Pract. Exp.* 25 (6) (2013) 1057–1071.
- [162] Y. Shi, R. Eberhart, A modified particle swarm optimizer, in: *International Conference on Evolutionary Computing*, Piscataway, NJ, 1998, pp. 69–73.
- [163] C. Grosan, A. Abraham, M. Nicora, Search optimization using hybrid particle sub-swarms and evolutionary algorithms, *Int. J. Simul. Syst. Sci. Technol.* 6 (10) (2005) 60–79.
- [164] G. Venter, J. Sobieszcanski-Sobieski, Multidisciplinary optimization of a transport aircraft wing using particle swarm optimization, *Struct. Multidiscip. Optim.* 26 (1–2) (2004) 121–131.
- [165] C. Praveen, R. Duvigneau, Low cost PSO using metamodels and inexact pre-evaluation: application to aerodynamic shape design, *Comput. Methods Appl. Mech. Eng.* 198 (2009) 1087–1096.
- [166] S. Haghghat, J.R.R.A. Martins, H.H.T. Lui, Aeroservoelastic design optimization of a flexible wing, *J. Aircraft* 49 (2) (mar 2012) 432–443.
- [167] B. Ouissam, D. Moulay, R. Mohamed, Title control of pitch angles to optimize the aerodynamic using particle swarm optimization, *WSEAS Trans. Syst. Control* 10 (2015) 22–28.
- [168] M. Azab, C. Ollivier-Gooch, High order aerodynamic optimization using new hybrid sequential quadratic programming-particle swarm intelligence technique, in: *50th AIAA Aerospace Sciences Meeting*, AIAA Paper 2012-0730, 2012, pp. 1–19.
- [169] R.E. Perez, K. Behdinan, Particle swarm optimization in structural design, in: *Swarm Intelligence: Focus on Ant and Particle Swarm Optimisation*, I-Tech Education and Publishing, Vienna, 2007, pp. 373–394.
- [170] M. Pant, R. Thangaraj, A. Abraham, Particle swarm optimization: performance tuning and empirical analysis, *Found. Comput. Intell.* 3 (2009) 101–128.
- [171] P.W. Jansen, *Aerostructural optimization of non-planar lifting surfaces*, University of Toronto, 2009, PhD thesis.
- [172] P.C. Fourie, A.A. Groenwold, The particle swarm optimization algorithm in size and shape optimization, *Struct. Multidiscip. Optim.* 23 (2002) 259–267.
- [173] A. Nejat, P. Mirzabeygi, M.S. Panahi, Airfoil shape optimization using improved multiobjective territorial particle swarm algorithm with the objective of improving stall characteristics, *Struct. Multidiscip. Optim.* 49 (6) (2014) 953–967.
- [174] R. Mendes, *Population Topologies and Their Influence in Particle Swarm Performance*, Universidade doMinho, 2004, PhD thesis.
- [175] M.B. Ghalia, Particle swarm optimization with an improved exploration-exploitation balance, *IEEE 978* (2008) 759–762.
- [176] R. Duvigneau, A multi-level particle swarm optimization strategy for aerodynamic shape optimization, in: *Evolutionary Methods for Design, Optimisation, and Control*, 2007.
- [177] X. Hu, X. You, Determination of the optimal control parameter range of air supply in an aircraft cabin, *Build. Simul.* 8 (4) (2015) 465–476.
- [178] J. Johnson, P. Picton, *Mechatronics: Concepts in Artificial Intelligence: Designing Intelligent Macheines Volume 2*, Butterworth Heinemann and The Open University, Oxford, 1996.
- [179] S. Kirkpatrick, C.D. Gelatt, M.P. Vecchi, Optimization by simulated annealing, *Science* 220 (4598) (1983) 671–680.
- [180] S. Kirkpatrick, Optimization by simulated annealing: quantitative studies, *J. Stat. Phys.* 34 (1984) 975–986.
- [181] S. Russell, P. Norvig, *Artificial intelligence: a modern approach*, 2nd edition, Prentice Hall, New Jersey, 1995.
- [182] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* 21 (6) (1953) 1087–1092.
- [183] R. Mukesh, K. Lingadurai, Aerodynamic optimization using simulated annealing and its variants, *Int. J. Eng. Trends Technol.* 2 (3) (2011) 73–77.
- [184] R.W. Eglese, Simulated annealing: a tool for operational research, *Eur. J. Oper. Res.* 46 (1990) 271–281.
- [185] W.T. Tiow, K.F.C. Yiu, M. Zangeneh, Application of simulated annealing to inverse design of transonic turbomachinery cascades, *J. Power Energy* 216 (2002) 59–73.
- [186] H.E. Romeijn, R.L. Smith, Simulated annealing for constrained global optimization, *J. Global Optim.* 5 (1994) 101–126.
- [187] W.L. Goffe, G.D. Ferrier, J. Rogers, Global optimization of statistical functions with simulated annealing, *J. Econometr.* 60 (1994) 65–99.
- [188] B. Behzadi, C. Ghotbi, A. Galindo, Application of the simplex simulated annealing technique to nonlinear parameter optimization for the SAFT-VR equation of state, *Chem. Eng. Sci.* 60 (2005) 6607–6621.
- [189] E. Leylek, J. Manzo, E. Garcia, Engineering notes: bat-inspired wing aerodynamics and optimisation, *J. Aircraft* 47 (1) (2010) 323–328.
- [190] J.E. Manzo, E.A. Leylek, E. Garcia, Drawing insight from nature: a bat wing for morphing aircraft, in: *Proceedings of ASMA Conference on Smart Materials, Adaptive Structures, and Intelligent Systems*, SMASIS08, Ellicott City, MD, 2008, pp. 671–678.
- [191] J. Liu, Novel Taguchi-simulated annealing method applied to airfoil and wing planform optimization, *J. Aircraft* 43 (1) (2006) 102–109.
- [192] W.H. Press, S.A. Teukolsky, Simulated annealing optimization over continuous spaces, *Comput. Phys.* 5 (426) (1991) 426–429.
- [193] X. Wang, M. Damodaran, Comparison of deterministic and stochastic optimisation algorithms for generic wing design problems, *J. Aircraft* 37 (5) (2000) 929–932.
- [194] X. Wang, M. Damodaran, Aerodynamic shape optimization using computational fluid dynamics and parallel simulated annealing algorithms, *AIAA J.* 39 (8) (2001) 1500–1508.
- [195] S. Aly, M. Ogot, R. Pelz, Stochastic approach to optimal aerodynamic shape design, *J. Aircraft* 33 (5) (1996) 956–961.
- [196] P. Gage, I. Kroo, P. Sobieski, Variable-complexity genetic algorithm for topological design, *AIAA J.* 33 (11) (1995) 2212–2217.
- [197] A.H.W. Bos, Aircraft conceptual design by genetic/gradients-guided optimization, *Eng. Appl. Artif. Intell.* 11 (3) (1998) 377–382.
- [198] A. Vicini, D. Quagliarella, Airfoil and wing design through hybrid optimization strategies, in: *AIAA Paper 1998-2729*, 1998.
- [199] F. Muyl, L. Dumas, V. Herbert, Hybrid method for aerodynamic shape optimization in automotive industry, *Comput. Fluids* 33 (5–6) (2004) 849–858, ISSN 00457930.
- [200] P. Gudla, R. Ganguli, An automated hybrid genetic-conjugate gradient algorithm for multimodal optimization problems, *Appl. Math. Comput.* 167 (2005) 1457–1474.
- [201] T. Goel, R. Vaidyanathan, R.T. Haftka, W. Shyy, N.V. Queipo, K. Tucker, Response surface approximation of pareto optimal front in multi-objective optimization, *Comput. Methods Appl. Mech. Eng.* 196 (2007) 879–893.
- [202] Y.S. Ong, K.Y. Lum, P.B. Nair, Hybrid evolutionary algorithm with Hermite radial basis function interpolants for computationally expensive adjoint solvers, *Comput. Optim. Appl.* 39 (1) (2008) 97–119.
- [203] S. Lehner, W.A. Crossley, Hybrid optimization for a combinatorial aircraft design problem, in: *9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*, AIAA 2009-7116, South Carolina, 2009, pp. 1–15.
- [204] X. Xing, M. Damodaran, Design of three dimensional nozzle shapes using hybrid optimisation techniques, in: *42nd AIAA Aerospace Sciences Meeting and Exhibit*, AIAA 2004-26, 2004.
- [205] X. Xing, M. Damodaran, Design of three-dimensional nozzle shape using NURBS, CFD and hybrid optimisation strategies, in: *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA-4368, 2004.
- [206] J. Herbert-Acero, J. Martínez-Lauranchet, O. Probst, S. Méndez-Díaz, K. Castillo-Villar, M. Valenzuela-Rendón, P. Réthoré, A hybrid

- metaheuristic-based approach for the aerodynamic optimization of small hybrid wind turbine rotors, *Math. Probl. Eng.* (2014) 1–18.
- [207] M.S. Khurana, H. Winarto, A.K. Sinha, Airfoil optimisation by swarm algorithm with mutation and artificial neural networks, in: 47th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, (January), 2009, pp. 1–19.
- [208] L. Cao, D. Zhang, Aerodynamic configuration optimization for hypersonic gliding vehicle based on improved hybrid multi-objective PSO algorithm, in: *IEEE: ICSPCC2015*, 2015, pp. 1–5.
- [209] M.A. Sahnehsaraei, M.J. Mahmoodabadi, M. Taherkhorsandi, S.M.M. Yazdi, A hybrid global optimisation algorithm: particle swarm optimisation in association with a genetic algorithm Complex System Modelling and Control Through Intelligent Soft Computations, 9, 2015, pp. 45–86.
- [210] Y. Wang, W. Li, Aircraft aerodynamic configuration design based on hybrid multi-objective particle swarm optimisation algorithm, *Chin. J. Aeronaut.* 29 (5) (2008) 1202–1206.
- [211] D.J.J. Toal, N.W. Bressloff, A.J. Keane, C.M.E. Holden, The development of a hybridized particle swarm for kriging hyperparameter tuning, *Eng. Optim.* 43 (6) (2011) 675–699.
- [212] J.I. Madsen, W. Shyy, R.T. Haftka, Response surface techniques for diffuser shape optimization, *AIAA J.* 38 (9) (2000) 1512–1518.
- [213] A. Lundberg, P. Hamlin, D. Shankar, A. Broniewicz, T. Walker, C. Landström, Automated Aerodynamic Vehicle Shape Optimization Using Neural Networks and Evolutionary Optimization, *SAE Int. J. Passenger Cars – Mech. Syst.* 8 (1) (2015) 242–251.
- [214] K. Boopathy, M. Rumpfkeil, A multivariate interpolation and regression enhanced kriging surrogate model, in: 21st AIAA Computational Fluid Dynamics Conference, AIAA 2013-2964, 2013, pp. 1–18.
- [215] A.J. Keane, Wing optimization using design of experiment, response surface, and data fusion methods introduction, *J. Aircraft* 4 (4) (2003).
- [216] N.V. Queipo, R. Haftka, W. Shyy, T. Goel, R. Vaidyanathan, P.K. Tucker, Surrogate-based analysis and optimization, *Progr. Aerosp. Sci.* 41 (1) (2005) 1–28.
- [217] T.W. Simpson, J. Peplinski, P.N. Kock, J.K. Allen, Metamodels for computer-based engineering design: survey and recommendations, *Eng. Comput.* 17 (2) (2001) 129–150.
- [218] R. Jin, W. Chen, T.W. Simpson, Comparative studies of metamodeling techniques under multiple modeling criteria, *J. Struct. Optim.* 23 (1) (2001) 1–13.
- [219] A.I.J. Forrester, A.J. Keane, Recent advances in surrogate-based optimization, *Progr. Aerosp. Sci.* 45 (2009) 50–79.
- [220] Y. Lian, M. Liou, Multiobjective optimisation using coupled response surface model and evolutionary algorithm, *AIAA J.* 43 (6) (2005) 1316–1325.
- [221] K.B. Collins, D.N. Mavris, Application of low- and high-fidelity simulation tools to helicopter rotor blade optimization, *J. Am. Helicopter Soc.* 042003 (2013) 1–11.
- [222] D. Leusink, D. Alfano, P. Cinnella, Multi-fidelity optimization strategy for the industrial aerodynamic design of helicopter rotor blades, *Aerosp. Sci. Technol.* 42 (2015) 136–147.
- [223] T.D. Robinson, M.S. Eldred, K.E. Willcox, R. Haimes, Surrogate-based optimization using multifidelity models with variable parameterization and corrected space mapping, *AIAA J.* 46 (11) (2008) 2814–2822.
- [224] P.M. Zadeh, V.V. Toropov, A.S. Wood, Metamodel-based collaborative optimization framework, *Struct. Multidiscip. Optim.* 38 (2) (2009) 103–115.
- [225] B. Rosenbaum, V. Schulz, Response surface methods for efficient aerodynamic surrogate models, *Comput. Flight Test.* 123 (2013) 113–129.
- [226] D.J.J. Toal, A.J. Keane, Efficient multipoint aerodynamic design optimization, *J. Aircraft* 48 (5) (2011) 1685–1695.
- [227] A.I.J. Forrester, Black-Box calibration for complex-system, *Philos. Trans. R. Soc.* 356 (2010) 7–3579.
- [228] Y. Kuya, K. Takeda, X. Zhang, A.I.J. Forrester, Multi fidelity surrogate modeling of experimental and computational aerodynamic data sets, *AIAA J.* 49 (2) (2011) 289–298.
- [229] J. Laurenceau, M. Meaux, M. Montagnac, P. Sagaut, Comparison of gradient-based and gradient-enhanced response-surface-based optimizers, *AIAA J.* 48 (5) (2010) 981–994.
- [230] H. Chung, J.J. Alonso, Using gradients to construct cokriging approximation models for high-dimensional design optimization problems, in: 401th AIAA Aerospace Sciences Meeting and Exhibit, AIAA 2002-0317, 2002, pp. 1–15.
- [231] K. Takenaka, K. Hatanaka, W. Yamazaki, K. Nakahashi, Multidisciplinary design exploration for a winglet, *J. Aircraft* 45 (5) (2008) 1601–1611.
- [232] M.D. McKay, R.J. Beckman, W.J. Conover, A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Am. Stat. Assoc. Am. Soc. Q.* 42 (1) (1979) 55–61.
- [233] D.R. Jones, M. Schonlau, J. William, Efficient global optimization of expensive black-box functions, *J. Glob. Optim.* 13 (4) (1998) 455–492.
- [234] B. Glaz, P. Friedmann, L. Liu, Helicopter vibration reduction throughout the entire flight envelope using surrogate-based optimization, *J. Am. Helicopter Soc.* 54 (012007) (2009) 1–15.
- [235] W. Jacquet, B. Truyen, P. De Groen, I. Lemahieu, J. Cornelis, Global Optimization In Inverse Problems: A Comparison of Kriging and Radial Basis Functions. Technical report., 2005.
- [236] R.L.T. Bevan, D.J. Poole, C.B. Allen, T.C.S. Rendall, Simulation and surrogate-based design of rectangular vortex generators for tiltrotor aircraft wings, in: 33rd AIAA Applied Aerodynamics Conference, AIAA 2015-2725, Dallas, TX, 2015, pp. 1–18.
- [237] R.L.T. Bevan, D.J. Poole, C.B. Allen, T.C.S. Rendall, Optimization of vane-type vortex generators for tiltrotor wings using computational fluid dynamics, in: 54th AIAA Aerospace Sciences Meeting, AIAA 2016-1304, San Diego, California, 2016, pp. 1–19.
- [238] Y.S. Ong, P.B. Nair, A.J. Keane, Evolutionary optimization of computationally expensive problems via surrogate modeling, *AIAA J.* 41 (4) (2003) 687–696.
- [239] A. Massaro, A. D'Andrea, B. Ernesto, Multiobjective-multipoint rotor blade optimization in forward flight conditions using surrogate-assisted memetic algorithms, in: Conference: 37th European Rotorcraft Forum, 2011.
- [240] F.A.C. Viana, R.T. Haftka, L.T. Watson, Efficient global optimization algorithm assisted by multiple surrogate techniques, *J. Glob. Optim.* 56 (2013) 669–689.
- [241] H. Tianyuan, Y. Xiongqing, Aerodynamic/stealthy/structural multidisciplinary design optimization of unmanned combat air vehicle, *Chin. J. Aeronaut.* 22 (4) (2009) 380–386.
- [242] V.R. Joseph, Y. Hung, A. Sudijianto, Blind kriging: a new method for developing metamodels, *ASMA J. Mech. Des.* 130 (2008) 1–8.
- [243] N. Namura, K. Shimoyama, S. Jeong, S. Obayashi, Kriging/RBF-hybrid response surface method for highly nonlinear functions, *J. Comput. Sci. Technol.* 6 (2012) 81–96.
- [244] D.H. Wolpert, W.G. Macready, No free lunch theorems for search, in: Technical Report SFI-TR-95-02-101, 1996.

Shaun N. Skinner – is a PhD candidate in the Aerospace Sciences research division in the School of Engineering at the University of Glasgow. His Ph.D. focuses on the numerical and experimental aerodynamic optimisation of non-planar wing configurations for drag, vibration, and load alleviation. Other research experience and interests in fluid dynamics includes: (i) supersonic film cooling; (ii) passive and active flow control for low and high speed flows; (iii) aero-elastics; (iv) experience in the design and usage of wing tunnel models and flow diagnostic equipment in experimental facilities ranging from low subsonic, incompressible flows, to supersonic flows, upto Mach 3.

Hossein Zare-Behtash – is a lecturer in the Aerospace Sciences research division in the School of Engineering – University of Glasgow. He is a Chartered Engineer, Member of the Royal Aeronautical Society, Member of the Institution of Mechanical Engineers, and Senior Member of the American Institute of Aeronautics and Astronautics. His research interests cover a spectrum of fluid dynamics challenges and applications related to fluid-structure interactions, such as: (i) active/passive flow control technologies for low-speed and transonic, applications such as plasma actuators, 3D bumps, and vortex generators, (ii) development/application of advanced flow diagnostic and instrumentation methods such as thermographic phosphors, and unsteady PSPs, (iii) extensive experience in the design and usage of a variety of experimental facilities ranging from low subsonic, incompressible, Mach numbers to hypersonics, Mach numbers 5 and above, (iv) unsteady aerodynamics of non-planar wing planforms – optimised using structured genetic algorithms.