Kedward, L. J., Allen, C. B., & Rendall, T. C. S. (2020). Gradient-Limiting Shape Control for Efficient Aerodynamic Optimisation. *AIAA Journal*. <https://doi.org/10.2514/1.J058977>

Peer reviewed version

Link to published version (if available):
[10.2514/1.J058977](https://doi.org/10.2514/1.J058977)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via American Institute of Aeronautics and Astronautics at <https://arc.aiaa.org/doi/full/10.2514/1.J058977> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Gradient-Limiting Shape Control for Efficient Aerodynamic Optimisation

L. J. Kedward [†] ; C. B. Allen [‡] ; T. C. S. Rendall [§]

Department of Aerospace Engineering, University of Bristol, Bristol, UK

Local shape control methods, such as B-Spline surfaces, are well-conditioned such that they allow high-fidelity design optimisation, however this comes at the cost of degraded optimisation convergence rate as control fidelity is refined due to the resulting exponential increase in the size of the design space. Moreover, optimisations in higher-fidelity design spaces become ill-posed due to high frequency shape components being insufficiently bounded; this can lead to non-smooth and oscillatory geometries that are invalid both in physicality (shape) and discretisation (mesh). This issue is addressed here by developing a geometrically-meaningful constraint to reduce the effective degrees of freedom and improve the design space thereby improving optimisation convergence rate and final result. A new approach to shape control is presented using coordinate control (x, z) to recover shape-relevant displacements and surface gradient constraints to ensure smooth and valid iterates. The new formulation transforms constraints directly onto design variables, and these bound the out-of-plane variations to ensure smooth shapes as well as the in-plane variations for mesh validity. Shape gradient constraints approximating a C^2 continuity condition are derived and demonstrated on a challenging test case: inviscid transonic drag minimisation of a symmetric NACA0012 aerofoil. Significantly the regularised shape problem is shown to have an optimisation convergence rate independent of both shape control fidelity and numerical mesh resolution, while still making use of increased control fidelity to achieve improved results. Consequently, a value of 1.6 drag counts is achieved on the test case, the lowest value achieved by any method.

Copyright © 2020 by Laurence Kedward

Aspects of this work have been presented at AIAA SCITECH and AVIATION forums: AIAA-2018-3951, AIAA-2019-1209, AIAA-2019-3175.

[†] PhD Student, AIAA Student Member, laurence.kedward@bristol.ac.uk, Bristol, BS8 1TR, UK

[‡] Professor of Computational Aerodynamics, AIAA Senior Member, c.b.allen@bristol.ac.uk, Bristol, BS8 1TR, UK

[§] Senior Lecturer, AIAA Member, thomas.rendall@bristol.ac.uk, Bristol, BS8 1TR, UK

Nomenclature

C_D	=	Drag coefficient
ϵ	=	Absolute constraint bound
h	=	Mean surface mesh spacing
N_j^k	=	k^{th} order B-Spline basis
s	=	Arc-length
t	=	Parametric ordinate
x, z	=	Cartesian coordinates
κ	=	Surface curvature
σ	=	Length scale constraint parameter
D	=	Derivative matrix
n	=	Unit normal vector
p	=	B-Spline control point
q	=	B-Spline derivative control point
t	=	Unit tangent vector
$x(t)$	=	Uniparametric curve coordinates
Δx	=	Discrete curve facet
α	=	Design variable vector
$\dot{\square}$	=	parametric derivative
$\square'$	=	arc-length derivative

I. Introduction

Computational fluid dynamics (CFD) has matured rapidly over the last half-century and, in combination with increasing computing power, has become a commonplace tool for modelling fluid flow in industry. In addition to reducing the reliance of industry on expensive and lengthy ground and in-flight tests, this accessibility to CFD is also being increasingly applied in numerical optimal design: the coupling of optimisation methods to numerical analysis tools to improve on a design objective. Shape optimisation in particular has many applications across aerospace, with optimal design objectives such as aerodynamic drag, structural weight, aeroacoustic signature, mission endurance, electromagnetic stealth as well as various multidisciplinary objectives.

Aerodynamic shape optimisation has reached a good level of maturity, as well as variety, among the tools and approaches adopted, particularly in the areas of adjoint methods, surface and volume mesh manipulation and quasi-Newton search methods. Much collaborative effort towards this has been in the AIAA Aerodynamic Design Optimization Discussion Group (ADODG*) through which a set of benchmark cases [1–3] have been defined and tackled; here the current state-of-the-art sees much recent work surrounding the Common Research Model wing [4, 5] and wing-body-tail [6, 7] configurations. The cutting-edge in the field now sees the explicit inclusion of aerodynamic phenomena, such as buffet onset [8] and aeroelastic flutter [9], into the optimisation thereby further refining the design problem being posed. Importantly, maturity in the field has been demonstrated by certain studies [2, 10] which compare different optimisation frameworks and achieve consistent optimised results therewith.

A shape optimisation framework often comprises three main components: a numerical analysis code (*e.g.* finite element analysis, CFD); a geometry representation and manipulation method (*e.g.* computer-aided-design model, shape basis functions); and a numerical optimisation method (*e.g.* gradient-descent, gravitational search). The optimisation process involves successive calls to the analysis tool by the optimiser to iteratively improve upon the objective. Since the analysis method itself often involves an expensive iterative process, optimal design still remains a computationally expensive process. Current state-of-the-art gradient-based search algorithms, such as those using quasi-Newton methods, have an optimisation cost proportional to the effective degrees-of-freedom in the problem; these degrees-of-freedom can be thought of as the subspace formed by the combination of design variables and constraints.

For shape optimisation, the design variables are defined by the shape control method acting as the interface between the search algorithm, which operates on the vector of design variables, and the analysis code, which is only concerned with the geometry. Hence the choice and configuration of the shape control methodology can have particular influence on the number of optimisation iterations, and therefore computational cost. This is demonstrated in large scale comparative studies by Masters *et al.* [11, 12] whereby the efficiency, robustness and final result are compared across seven different shape parameterisation methods.

Among such parameterisations, the ability to refine the design space for increased control fidelity varies significantly, an observation that can be attributed to the conditioning of such geometric parameterisations. Whereas some shape control methods rapidly become ill-conditioned when refined, which consequently adds little to the design space, local shape control methods, such as B-Splines, are well-conditioned where the basis retains good linear independence when refined; this allows high-fidelity shape control but comes at the cost of degraded optimisation performance as control fidelity is refined due to the resulting exponential increase in the size of the design space.

While increasing control fidelity, these additional degrees of-freedom also introduce extraneous geometries into the design space which, for local control methods, consist of non-smooth and oscillatory shapes. Such undesirable shapes have been observed in optimisation literature [12–14] for medium to high-fidelity B-Splines. The propensity for

---

*<https://sites.google.com/view/mcgill-computational-aerogroup/adodg> [retrieved February 2020]

gradient-based search algorithms to amplify high-frequency grid components and produce such oscillatory shapes is clear when one considers the limiting case of using mesh-point control whereby the design variables are the individual vertices of the discretised geometry, or mesh. This natural approach to shape control, which for a fixed mesh resolution provides the most comprehensive design space, is well-conditioned since there is no parametric coupling between vertices. However, even with the maturation of adjoint flow solvers[15] which provide surface sensitivities at a cost effectively independent of the number of design variables, mesh point control alone leads to a degeneration of surface smoothness since the gradients of the discrete shape functional have a lower regularity than the original shape [16, 17]. The same problem is encountered for shape and topology optimisation in solid mechanics where the nature of finite element analysis induces non-physical optima with jagged edges and so-called ‘chequerboard’ patterns [18, 19]. This problem is most commonly avoided by restricting the number of surface spline control points or using global volume splines which, while generally guaranteeing smoothness, do not remain well-conditioned with refinement. Moreover, it is rare within literature for authors to perform a control-fidelity study demonstrating the effect of refining their chosen shape control method; this may be due to computational constraints or the fact that shape control is peripheral to the research goal of the work.

In this work, the effective degrees of freedom of high-fidelity shape control methods are considered from a geometric standpoint, motivated by the observation that such problems are ill-posed by being insufficiently bounded; that is, they naturally include geometries that are invalid both in physicality (shape) and discretisation (mesh). Several methods for shape control exist, which address the ill-posed nature of the problem either directly or indirectly; however such methods often *a priori* constrain the design space at low fidelities and perform poorly when shape control is refined. The goal of the work here is to present a new approach to shape control which explicitly removes non-smooth and oscillatory shapes from the design space by limiting contributions from high-frequency components. This filtering approach, implemented using constraints on surface gradients, reduces the effective degrees of freedom in the design problem in a geometrically meaningful way while still allowing the high-fidelity design space to be exploited.

II. Background

A variety of shape control methods are encountered in aerodynamic shape optimisation all of which address shape smoothness in some way either directly or indirectly; commonly-used methods are presented in Table 1 and discussed briefly here. A simple but effective approach to addressing smoothness during optimisation is to penalise bad shapes by augmenting the objective function with a weighted penalty term given by a function of some surface derivative such as curvature [16] or alternatively on a calculated quantity such as pressure distribution [40]. This method is a form of regularisation which arises in a variety of fields for solving ill-posed problems. One way of viewing regularisation is the addition of prior knowledge in the form of assumptions that characterise the solution; the solution of the regularised problem is that which is most probable under the assumptions made. Within shape optimisation, penalisation of shape

Table 1 Common shape control methods for optimisation

Shape parameterisation <i>Continuous shape functions are used to describe geometry such that the number of design variables is much less than the number of mesh nodes</i>	B-Spline surfaces [20] Bezier surfaces [21] Volume splines: Bezier[22], B-Spline[23], NURBS[3], RBF[24] Hicks-Henne [25] CSTs [26] Fictitious loads [27] SVD [28] Partial differential equation [29] Volume of solid active contour [30, 31]
Multilevel optimisation <i>Successive optimisations are performed where shape control is refined after each stage</i>	Subdivision [32, 33] B-Spline/Bezier knot insertion [34] RBF [35]
Sensitivity filtering <i>Mesh vertices are used for shape control and surface sensitivities are filtered to remove high frequency components</i>	Implicit Sobolev [36] Explicit Gaussian [37] Multilevel preconditioner [38]
Shape filtering <i>Mesh vertices are used for shape control and both surface sensitivities and shape updates are filtered to remove high frequency components</i>	Gaussian [18] Vertex Morphing [39]

derivatives is equivalent therefore to an *a priori* characterisation of the differentiability of the solution shape. These methods are simple to implement since they only augment the objective with an explicit term, however the choice of weighting parameter, which represents an arbitrary trade-off between performance and smoothness, is not explicit but still directly influences the final result.

The most prevalent approach to addressing the challenges of shape optimisation is shape parameterisation, referring to methods that describe the geometry using a lower-dimension representation with intuitive design controls. Examples include: spline surfaces; volume splines; and various other shape functions derived analytically (Hicks-Henne and shape-class functions), from physical considerations (fictitious loads) or numerically (orthogonal modes). All such methods typically reduce to a linear combination of shape basis functions; different parameterisations vary only in the form and fidelity of the basis functions. Inherent in the reduction of dimensionality is the introduction of smoothing, specifically the amplitude of higher order derivatives are implicitly bound by that of the basis functions. This advantage was emphasised by Braibant and Fleury when advocating the use of B-Splines for shape optimisation [20]. However the reduction in dimensionality inevitably corresponds to a reduction in the available design space and therefore attainable optimum; as such, the choice of parameterisation fidelity becomes important and this is investigated extensively by Masters *et al.* for several different parameterisation methods [11, 12]. High-fidelity parameterisations (many design

variables) often exhibit poor optimisation performance due to degraded design space conditioning or the exponential increase in design space and generation of non-smooth shapes [12–14] as shape control locality approaches that of mesh-point control. In their studies, one parameterisation is shown to perform particularly well, orthogonal modes extracted by singular value decomposition (SVD). This method uses a database of aerofoils to extract shape modes, a subset of which is used as a basis for shape optimisation. The chosen SVD modes correspond to high predominance in the aerofoil library which also therefore represents filtering and removal of low-importance high-frequency shape information. Similar treatment of high-frequency information occurs in multilevel methods which progressively or adaptively enrich the parameterisation during successive optimisations [32–35, 41, 42].

Splines and general basis function approaches describe the shape explicitly as a function of the design variables; implicit methods also exist however, whereby the design variables constrain a mathematical problem, the solution to which defines the shape. For example, the partial differential equation (PDE) approach by Bloor and Wilson [29] solves a fourth-order PDE where the boundary conditions are defined by the design variables. More recently, the restricted snake volume of solid (RSVS) method of Payot *et al.* [30, 31] defines shapes as the minimal surface enclosing specified volume fractions on a background grid. In this approach the volume fractions of the background grid cells are the design variables which, in combination with the restricted snakes active contour method, are able to produce shapes of arbitrary topology.

An alternative to shape parameterisation is direct mesh point control which uses the individual coordinates of nodes in the discretised analysis mesh as design variables. The advantage of this method is that the design space is only limited by the analysis discretisation - the most expansive design space possible. Furthermore, it is trivial to implement compared to setting up shape parameterisation functions, and generalisable to arbitrary geometries, assuming a discrete surface mesh exists. However, despite the theoretical simplicity, practical implementation requires that either the gradients or surface perturbations be smoothed to maintain a smooth surface. Sensitivity filtering involves passing the raw shape sensitivity information through some form of low-pass filter before it is presented to the search algorithm, the intended effect of this being that the removal of high-frequency components in the shape sensitivity leads to a smooth search direction and therefore maintains surface smoothness. Common within aerodynamic shape optimisation is a weighted Sobolev gradient as suggested by Jameson[36] and used in combination with design variables defined as normal perturbations to the initial geometry. This second-order implicit smoothing method is equivalent to applying the inverse of the mesh Laplacian [17] to the surface sensitivities and is most often combined with a simple steepest descent search forming the so-called Smoothed Descent Method [43]. Studies by Jameson notably conclude that the optimisation convergence of the smoothed steepest descent method is independent of the number of design variables [36]. Stück *et al.* [37] instead use a Gaussian kernel to explicitly smooth shape sensitivities and the resulting filter is shown to be first-order equivalent to the Sobolev method; however the implementation has improved flexibility and efficiency compared to the implicit Sobolev method which requires the solution of a partial differential equation.

Explicit and implicit smoothing methods are studied by Jaworski and Müller [44] within a multilevel shape optimisation procedure. Vazquez *et al.* [38] draw analogy between sensitivity filtering and preconditioning of discrete partial differential equations, and derive a multi-level filter to compensate the loss of regularity in the gradient update.

The methods discussed so far primarily control the regularity of the shape-relevant variation (normal to the surface), however the mesh-relevant (in-plane) variation must also be addressed to maintain validity of the surface mesh. The objective of this work is to generalise the regularisation effects of the methods discussed above (penalty-regularisation, shape parameterisation, filtering) to allow the use of high-fidelity well-conditioned shape control methods and remove dependence of the optimisation convergence rate on shape control fidelity. Shape gradient constraints are derived and it is demonstrated that a sufficiently constrained problem space is all that is required for a well-posed and consistent shape optimisation problem without limitation on the numerical method used for solution.

The remainder of the paper is organised as follows: in the following section the key challenges of shape control for optimisation are presented leading to the motivation for the present work. In section IV two-dimensional shape constraints are derived for both mesh point control and cubic B-Splines. Finally sections V and VI present the methodology and results for applying the shape constraint formulation to drag minimisation using a gradient-based adjoint optimisation framework.

III. Shape Control for Optimisation

An advantage of using shape parameterisation is that the shape control fidelity (number of design variables), which defines the extents of the design space, can be selected as is required; for example B-Spline control polygons can be refined via knot-insertion to generate higher-fidelity control, and shape functions, such as Hicks-Henne, can be defined at an arbitrary number of positions on the geometry. High-fidelity shape control is of particular benefit for aerodynamic design objectives due to their non-linearity and the necessary coupling of design variables defining a geometric boundary. It is well known that high control fidelity is required for aerofoil optimisation [33] and more generally in the chord-wise direction [4], especially at the leading edge [25]. Moreover, high-fidelity localised shape control is found to be particularly beneficial when dealing with similarly localised geometric constraints such as thickness and when performing multipoint optimisations involving multiple flight conditions [10].

However high-fidelity shape parameterisations generally suffer from degraded performance during numerical optimisation as demonstrated by Masters *et al.* [11, 12]. In their studies it was found that as the number of design variables increased, the convergence of optimisation was poor and the final result was often sub-optimal. The choice of control fidelity therefore becomes a trade-off between optimisation performance and extent of the design space. The cause of degraded efficiency at high-fidelity can be categorised as either an ill-posed problem (presence of undesirable shapes) or an ill-conditioned problem corresponding to basis functions of local and global support respectively.

In the former case, for local shape parameterisations (*e.g.* mesh points and B-Splines), as the number of design

variables increases the shape functions become more localised, which introduces both increased design fidelity but also oscillatory shapes characterised by high amplitudes in the high frequency components. The good linear-independence of such local control methods means that the useable design space grows at an exponential rate when design variables are added, which consequently results in degraded optimisation convergence rate. Moreover, the shape optimisation process is increasingly prone to local minima formed by undesirable oscillatory shapes since there is no explicit mechanism to avoid non-smooth search directions. The poor performance of B-Splines at high-fidelity matches that observed by Reuther and Jameson [13] and again by Castonguay *et al.* [14] who subsequently showed that this effect is eliminated when used with sensitivity filtering. However gradient smoothing is not general since the discrete gradients are no longer consistent with the optimisation problem which is particularly important for constrained SQP methods. Masters *et al.* [12] found that increasing the order of B-Splines, which increases the support and smoothness of the basis, was also able to alleviate some of the convergence issues, though this removes the advantages of using low-order B-Splines in their compact support and consequent good conditioning.

In the latter case, for shape functions which exhibit a large zone of influence and remain smooth with refinement, the conditioning of the shape control generally increases significantly as the number of design variables increases. Whereas local basis functions retain linear independence as they are refined, shape basis functions with large support suffer from deteriorating linear independence such that increasing the number of design variables give diminishing additions to the design space. This means that similar shapes do not have similar control vectors, *i.e.* shapes that are geometrically similar do not lie close together within the design space. This is especially undesirable for gradient-based optimisers which generate a sequence of solutions by stepping through the design space. The effect of this being that convergence rate is unaffected by refinement, since the useable design space is not practically improved due to ill-conditioning, which therefore means that the final attainable result of the optimisation is also not improved. For example, in [45] high-order Bezier (96 knots) became ‘saturated’ as the design variables hit their respective upper bounds during optimisation and preventing further improvement. This is explained by the global support of simple Bezier curves leading to a highly-conditioned design space where small shape modifications require large changes to all design variables.

The two effects, non-smoothness and ill-conditioning, at high shape control fidelities are demonstrated in Figures 1 and 2. Figure 1 shows the basis functions for cubic B-Splines at two control fidelities; when control fidelity is increased the basis functions retain good linear independence but become highly localised thereby allowing non-smooth shapes. Conversely when the Hicks-Henne basis, shown in Figure 2, is refined the basis functions retain smoothness but become increasingly indistinguishable which represents deteriorating linear independence (ill-conditioning).

A. Design space conditioning

To demonstrate the effect of design space conditioning on shape optimisation, an inverse design problem is considered for a two-dimensional aerofoil. The objective is the RMS difference with a target pressure distribution which is evaluated

Fig. 1 Cubic B-Spline shape functions

Fig. 2 Hicks-Henne shape functions

with a panel method. Gradients for the inverse problem are evaluated discretely using automatic differentiation and optimisation is performed using the SNOPT [46] SQP package. Shape control is performed using both cubic B-Splines and Hicks-Henne functions at varying levels of fidelity.

Shown in Figure 3 are the optimisation convergence histories for both shape control methods applied to the inverse design problem. For the B-Spline parameterisation, the commonly-observed deterioration in initial convergence rate due to increased number of design variables is clearly seen prior to the quadratic convergence achieved in the last stages when the quasi-Newton Hessian approximation has been sufficiently constructed. However the same behaviour is not seen for the Hicks-Henne parameterisation. The B-Spline shape control exhibits deteriorating initial convergence rate with increases in control fidelity due to its good linear-independence which means that the design space grows exponentially when refined. As a result the B-Spline case is able to achieve better overall results albeit at a slower pace. By contrast, the Hicks-Henne parameterisation does not experience any significant change in initial convergence rate as the number of design variables increases. Moreover, there is little improvement in optimised result when increasing the control fidelity of the Hicks-Henne parameterisation. This illustrates the effect of an ill-conditioned shape control in restricting the ‘useable’ design space; whereas the design space has additional degrees of freedom in the mathematical sense, the poor conditioning means that numerical methods such as gradient-based optimisation cannot make good use of them.

This is further illustrated quantitatively in Figure 4(a) which shows the condition number for each shape control basis and Figure 4(b) which shows the magnitude of the optimised design variable vectors when performing inverse pressure

Fig. 3 Inverse pressure design convergence

Fig. 4 Shape control conditioning

design. For the B-Spline case, the conditioning of the basis is good and largely independent of the number of design variables; this is reflected in the magnitude of the optimised design variable vector which remains well-bounded and close the geometry it is representing. By contrast, the Hicks-Henne parameterisation rapidly becomes ill-conditioned as it is refined. This results in increasingly larger design variable values as the optimiser attempts to explore the ill-conditioned design space. The magnitude of the resulting optimised design variables do not match that of the underlying geometry and importantly, the optimisation has not achieved optimised objective values comparable to those from the B-Spline parameterisation.

B. Control orientation

Also of importance for shape optimisation, in addition to smoothness and conditioning, are the actual degrees of freedom available to the search algorithm. Among previous studies it is common for the parameterisation to be manually constrained such that only one ordinate is free to move [45], or that certain degrees of freedom (*e.g.* sweep, taper)

are specified by hand [24], or combinations of both [47]. The parameterisation of a single coordinate, thickness, for aerofoils is convenient and commonly adopted since aerofoils are predominantly planar such that most of the surface normals are closely aligned with the vertical coordinate. However this has three drawbacks: first, the single coordinate parameterisation is restricted in its ability to produce local rotation of the surface; second, there is poorly defined control in areas where the local normals do not initially align with the vertical (*e.g.* the leading edge); and finally the notion of a dominant coordinate direction is not general and does not extend to three-dimensional geometries. In a similar way, manually specifying degrees of freedom relies on user-knowledge which, in the general case, does not necessarily know which controls are useful beforehand. The implication is that to allow shape relevant perturbations (displacements normal to the current shape) at all stages of optimisation, shape control must be permitted in all coordinate directions. A side-effect of this is that consideration must be made for the grid-relevant, in-plane, variation of the discretised surface.

Based on this, the rationale for the work here is to perform localised shape control in all coordinate directions and constrain the larger resulting design space by utilising a geometric constraint to maintain validity and continuity of the discrete surface. The use of local shape control methods ensures the design space is well-conditioned which means that refinements are meaningful and which therefore allows high-fidelity design optimisation. Constraint-based continuity has been applied in [48] and [49] for structural topology optimisation where first spatial derivatives were bounded to enforce approximate C^0 on the solution. Xu *et al.* [50] use a non-linear constraint to enforce geometric continuity (G^2) between separate spline patches in their CAD-based shape optimisation. Similarly Cinquegrana *et al.* [51] include a penalty on derivatives of surface curvature to counter oscillations induced by a NURBS parameterisation. In the work here an approximate C^2 condition is applied across the entire discretised shape to control both shape smoothness and mesh quality.

IV. Shape Constraints

In this section a constraint-based approach is developed based on an approximate C^2 condition to allow general parameterisation of more than one ordinate of the discretised surface mesh when using two different local control methods: mesh-point control and B-Spline control.

A. Geometric smoothness

In this section the concept of smoothness is presented within the mathematical framework of a plane uniparametric curve (Figure 5(a)) in $\mathbb{R}^2$; there is no loss of generality here since the same framework can be applied to surfaces in $\mathbb{R}^3$ by application along isoparametric curves. The curve is described continuously by:

$$\mathbf{x} = \mathbf{x}(t) = \begin{bmatrix} x(t) \\ y(t) \end{bmatrix}, \quad t \in [a, b] \subset \mathbb{R} \quad (1)$$

Fig. 5 Two-dimensional geometry representations

where the component functions x, y are differentiable with respect to the parameter t , denoted here by dot, such that the non-vanishing tangent vector is given by:

$$\dot{\mathbf{x}}(t) = \begin{bmatrix} \dot{x}(t) \\ \dot{y}(t) \end{bmatrix} \neq \mathbf{0}, \quad \forall t \quad (2)$$

A local coordinate frame at the point $\mathbf{x}(t)$ is constructed exposing familiar geometric quantities; the unit tangent vector $\mathbf{t}$ and unit normal vector $\mathbf{n}$ are given by:

$$\mathbf{t} = \frac{\dot{\mathbf{x}}}{\|\dot{\mathbf{x}}\|}, \quad \mathbf{n} = \frac{\dot{\mathbf{t}}}{\|\dot{\mathbf{t}}\|} \quad (3)$$

Any parametric curve can be re-parameterised naturally by the arc length parameterisation, $s(t)$:

$$s(t) = \int_a^t \|\dot{\mathbf{x}}\| dt \quad (4)$$

which yields the following relationships:

$$\mathbf{x}' = \mathbf{t}, \quad \mathbf{x}'' = \kappa \mathbf{n} \quad (5)$$

where κ is the local curvature and prime ' denotes differentiation with respect to arc length..

One way to mathematically describe smoothness is differentiability. The definition of $\mathbf{x}(t)$ above states that the curve is at least once differentiable with respect to the parameter t and hence this curve has parametric continuity 1, denoted C^1 , which implies continuously varying tangent vectors. Parametric curves that are twice differentiable with respect to their *current* parameterisation are C^2 and will be curvature continuous. However this definition is dependent on the particular parameterisation and indeed a curve can have continuity of curvature while not being C^2 . Hence the more general class of curves is defined as *Geometrically continuous* $G^2 \supset C^2$ which contains those curves that are twice differentiable with respect to the invariant arc length parameterisation. Using subscript + and - to represent the left and right limits respectively, then the conditions for C^2 and the more general G^2 can be expressed as:

$$C^2 : \quad \ddot{\mathbf{x}}_+(t) - \ddot{\mathbf{x}}_-(t) = \mathbf{0} \quad (6)$$

Fig. 6 Geometric and parametric continuity

$$G^2 : \quad \mathbf{x}''_+(s) - \mathbf{x}''_-(s) = \mathbf{0} \quad (7)$$

The difference between the stricter C^2 condition and the more general G^2 can be explained by expanding the term $\ddot{\mathbf{x}}$ using the chain rule:

$$\ddot{\mathbf{x}} = \mathbf{x}''(\dot{s})^2 + \mathbf{x}'(\ddot{s}) \quad (8)$$

Substituting the relations of equation 5 this becomes:

$$\ddot{\mathbf{x}} = \kappa(\dot{s})^2 \mathbf{n} + (\ddot{s}) \mathbf{t} \quad (9)$$

The parametric derivative $\ddot{\mathbf{x}}$ contains a normal component corresponding to the local curvature, and an in-plane component representing the ‘acceleration’ of the parameterisation. Notice that the condition for G^2 (equation 7), which uses the arc length derivative $\mathbf{x}''$, is concerned only with the normal component (for curvature continuity) whereas the stricter C^2 condition (equation 6), defined using the parametric derivative $\ddot{\mathbf{x}}$, also requires continuity of the parameterisation; this is illustrated in Figure 6 where ‘velocity’ refers to the rate of change in the parameter and ‘acceleration’ refers to the rate of change in velocity. *i.e.* Parametric continuity is a special case of geometric continuity. This makes intuitive sense since only the normal component is visible (shape-relevant) and so the class of visually smooth curves should not exclude those with discontinuous parameterisations hence the classification of geometric continuity G^k . Therefore all geometrically continuous curves permit an arclength parameterisation $t \rightarrow s$ (equation 4), however the same cannot be said about the inverse process, formulating a C^k parameterisation $s \rightarrow t$ for a G^k curve.

In this work, the parametric derivative will be used to enforce the stricter C^2 condition since, by suitable choice of the parameter t , the in-plane component can be used to represent the smoothness of the discretisation. Hence by enforcing the C^2 condition both the shape relevant and mesh relevant components can be controlled simultaneously.

B. Gradient constraint for mesh point control

Consider now a mesh made up of vertices sampled at discrete parameter locations (Figure 5(b)) described by:

$$\mathbf{x}_i = \mathbf{x}(t_i), \quad i \in [0, n] \subset \mathbb{Z} \quad (10)$$

In this discrete setting differentiation with respect to the parameterisation t (denoted by dot) is approximated by discrete difference in the logical domain such that the facets of the mesh are given by:

$$\dot{\mathbf{x}} = \Delta \mathbf{x}_i = \mathbf{x}_{i+1} - \mathbf{x}_i \neq \mathbf{0}, \quad i \in [1, n] \subset \mathbb{Z} \quad (11)$$

Similarly the relations of equation 3 give approximations to the unit tangent and normal vectors defined at the surface elements and vertices respectively by:

$$\mathbf{t}_i = \mathbf{t}(\Delta \mathbf{x}_i) = \frac{\Delta \mathbf{x}_i}{\|\Delta \mathbf{x}_i\|}, \quad \mathbf{n}_i = \mathbf{n}(\mathbf{x}_i) = \frac{\mathbf{t}_{i+1} - \mathbf{t}_i}{\|\mathbf{t}_{i+1} - \mathbf{t}_i\|} \quad (12)$$

and the arc length parameterisation is replaced by an approximation using chord length:

$$s_i = \sum_1^i \|\Delta \mathbf{x}_i\| \quad (13)$$

To enforce some smoothness restriction on the mesh, an analogous condition to that of C^2 (equation 6), based on the parametric derivative $\dot{\mathbf{x}}$ is sought. Since the parameter t corresponds to the mesh point index i , then the second parametric derivative $\ddot{\mathbf{x}}$ is simply the uniform mesh Laplacian given by:

$$\ddot{\mathbf{x}}_i = \Delta \mathbf{x}_{i+1} - \Delta \mathbf{x}_i = \mathbf{x}_{i+1} - 2\mathbf{x}_i + \mathbf{x}_{i-1} \quad (14)$$

On the discretised surface, the left and right limits of equation 6 cannot shrink to smaller than a surface element and so instead the following condition is formulated:

$$\|\ddot{\mathbf{x}}_{i+1} - \ddot{\mathbf{x}}_i\|_\infty \leq \epsilon, \quad \forall i \quad (15)$$

for some small bound ϵ , *i.e.* the variation in the mesh Laplacian across surface facets is bounded.

This condition can be interpreted by expansion and manipulation of the uniform mesh Laplacian $\ddot{\mathbf{x}}$, which gives the

Fig. 7 A section of a piecewise linear discretised curve showing the uniform mesh Laplacian (*left*); the uniform mesh Laplacian expressed in the local frame (*right*).

following equivalent form in the approximate local frame:

$$\ddot{x} = \frac{\Delta S_{i+1} + \Delta S_i}{2} \|t_{i+1} - t_i\| n_i + \frac{t_{i+1} + t_i}{2} (\Delta S_{i+1} - \Delta S_i) \quad (16)$$

Since $\|t_{i+1} - t_i\| = \sqrt{2(1 - \cos\theta)} \approx \theta$ for small values of the exterior angle θ , and that $\theta = \kappa s$ then this becomes:

$$\ddot{x} = \kappa_i (\Delta S_{i+1/2})^2 n_i + (\Delta S_{i+1} - \Delta S_i) t_{i+1/2} \quad (17)$$

Note the resemblance to the exact form given in equation 9 where the nodal values of s and t have become approximated by simple averages from the neighbouring elements. Hence the condition in equation 15 bounds both normal and in-plane contributions representing variations in local curvature (or exterior angle) and element size respectively. These two components of the uniform mesh Laplacian are illustrated in Figure 7. An intuitive effect of this is that regions of high curvature will require smaller elements and vice-versa for flat regions.

Note that it is the coordinate components of the gradient that are bounded and not explicit bounds on the the normal and in-plane components; separation of the normal and in-plane components is possible however this is non-linear with respect to the shape control, which introduces unnecessary complexity to the optimisation process.

C. Gradient constraint for B-Spline control

Whereas mesh point control is trivial to setup and gives the most design flexibility, this level of control fidelity is often excessive for most applications. Moreover mesh point control introduces dependency on the discretisation method and the quality of the resulting surface mesh, as well as requiring the non-trivial transfer of the optimal result back to a continuous representation. In this section equivalent gradient constraints are derived for curves represented by B-Spline functions, though the same methodology can be applied to all spline-like representations and permits an efficient extension to surfaces in three-dimensions. The B-spline basis has several properties that are particularly desirable for

shape control; first, the basis functions have local support such that the effect of any control point (design variable) is limited to a local neighbourhood of the geometry; second, it has linear independence and is therefore guaranteed to form a well-conditioned polynomial basis regardless of the desired number of degrees of freedom (number of control points); finally, the resulting function has well defined smoothness properties, specifically C^{k-2} for a k^{th} -order spline in regions away from repeated knots. For a comprehensive theory of B-splines, see Farin [52].

Whereas the analytic spline has guaranteed smoothness, a numerical discretisation thereof does not; that is to say, without re-sampling, manipulations to the analytic form can still result in poor discretisations. Re-sampling is undesirable in gradient-based shape optimisation since the process cannot guarantee continuity for gradient evaluation and also implies a costly volume mesh regeneration which suffers from the same differentiability concern.

In this section, the same discrete surface gradient constraints are posed now using the spline formulation; in addition to affording all the aforementioned advantages of using B-Splines, and as will be demonstrated in the results, this also remedies degraded optimisation performance of B-Splines at increased fidelity by reducing the effective degrees of freedom. Considering now a geometry described continuously by a k^{th} -order B-Spline representation:

$$\mathbf{x}(t) = \sum_{j=0}^{n-1} N_j^k(t) \mathbf{p}_j \quad (18)$$

Where $\mathbf{p}_j$ are the B-spline control points and $N_j^k(t)$ are the B-spline basis functions (shown in Figure 1) defined in parameter space by scaling and shifting according to some knot vector $\mathbf{u}$.

A useful feature of B-spline functions is that their derivatives are conveniently given by another B-spline function of lower degree. The first derivative with respect to the parameterisation t is:

$$\dot{\mathbf{x}}(t) = \sum_{j=0}^{n-2} \tilde{N}_j^{k-1}(t) \mathbf{q}_j \quad (19)$$

where the basis functions $\tilde{N}_j^{k-1}(t)$ are defined on a modified knot vector $\tilde{\mathbf{u}}$ formed by truncating the first and last knots of the original $\mathbf{u}$ and the new ‘control points’ $\mathbf{q}_j$ are given by:

$$\mathbf{q}_j = \frac{k-1}{u_{j+k} - u_{j+1}} (\mathbf{p}_{j+1} - \mathbf{p}_j) \quad (20)$$

In a similar way the second derivative $\ddot{\mathbf{x}}$ can be evaluated and the following condition enforced:

$$\|\ddot{\mathbf{x}}_{j+1} - \ddot{\mathbf{x}}_j\|_{\infty} \leq \epsilon \quad (21)$$

where the derivatives $\ddot{\mathbf{x}}_j$ are evaluated at the knot parameters u_j . In this way the constraints are only applied at a sample

of locations along the curve and the number of constraints is proportional to the number of control points.

This spline-based constraint is formulated in parametric space which allows an easy extension of the method to surfaces in three dimensions. Specifically, the same constraints can be applied along isoparametric curves defined on the surface at each corresponding tensor-product knot location; this introduces two sets of constraints one for each parametric direction, however the constraints remains linear and proportional to the number of control points.

D. Constraint bound and filtering

As demonstrated by the formulation of the previous sections, the constraint bound is equivalent to bounding the third derivative with respect to a parameterisation that corresponds to the mesh point or B-Spline connectivity, the effect being that a bounded third derivative ensures that the second derivative is a continuous function which therefore means that the original function is twice continuously differentiable, or approximately C^2 . This can be further interpreted as a low-pass filter on the wavelength in parametric space whereby tightening the constraint bound (smaller ϵ) lowers the cut-off frequency and more high-frequency components are attenuated. Choosing the constraint bound therefore corresponds to choosing the smallest *meaningful* wavelength for the problem. This could be a physically motivated, such as feature sizes and manufacturing tolerances, or numerical, related to the spatial discretisation (*i.e.* mesh quality). In fact for aerodynamic optimisation problems, numerical concerns are likely to be the driving factor in choosing the length scale since manufacturing tolerances are small and small geometric features are neglected (*'de-featured'*) for numerical analysis. For example, the numerically-motivated length scale could be chosen to be the smallest geometric wavelength that can be represented by a discretisation of fixed resolution and which conforms to some measure of mesh quality. Measures of mesh quality are beyond the scope of the present study and instead a sweep of values for the constraint bound will be performed to demonstrate its effect on the shape optimisation problem.

E. Numerical implementation

A consequence of bounding the parametric derivative as defined previously is that the resulting smoothness constraints for mesh point control and B-Splines are linear with respect to the coordinate design variables and can hence be implemented in following form:

$$\begin{bmatrix} D \\ -D \end{bmatrix} \alpha < \epsilon \begin{bmatrix} 1 \\ \vdots \end{bmatrix} \quad (22)$$

Here α is the design variable vector consisting of stacked mesh point or control polygon coordinates for mesh point and B-Spline shape control respectively and ϵ is the constraint bound which is henceforth given by:

$$\epsilon = \sigma h^3 \quad (23)$$

where h is the mean mesh spacing of the initial grid and matrix $\mathbf{D}$ is either a third difference matrix for mesh point control or a spline derivative matrix for B-Spline control. σ is the grid-independent constraint parameter related to the filter cut-off imposed by the constraint. By decreasing the constraint bound σ , the cut-off frequency is lowered such that more high-frequency components are attenuated. Evaluation of the gradient constraint on a database[11] of 1296 aerofoils[†] shows that values of σ between 25 and 100 are more than sufficient for common aerofoil variations. For more general geometry, the same analysis of the baseline geometry can give an expected upper bound on σ for subsequent shape optimisation.

The space and time costs associated with applying the constraints are dependent on the specific search algorithm, however the constraint matrices are banded diagonal such that the number of non-zero elements scales linearly with respect to the number of degrees of freedom; hence algorithms that implement sparse methods are especially applicable. The optimiser used in this work (see section V) selects feasible search directions at a cost proportional to the square of the system size [53] and which is negligible compared to the cost of the objective function evaluations. The general continuity constraint assumes non-vanishing surface tangents and so it must be applied in combination with a non-linear constraint preventing surface elements from shrinking below a certain size, which is here taken to be the smallest element on the initial mesh.

V. Aerodynamic Optimisation Methodology

This section describes the numerical methods comprising the aerodynamic optimisation framework used herein; specifically the methods used for flow solution, mesh deformation and gradient-based optimisation. In the following section, this framework is used in combination with the gradient constraints of the previous section on a two-dimensional inviscid drag minimisation problem.

A. Flow solver: SU2

For this work the Stanford University Unstructured (SU²) [54] flow solver is adopted. This open-source software was constructed with aerodynamic shape optimisation in mind and hence has both continuous and discrete adjoint implementations [55]. The main flow solver implements both the compressible Euler and RANS equations using an unstructured finite volume method.

Multigrid acceleration is available as well as MPI parallel processing. The SU² suite also includes other modules for tasks such as shape parameterisation, mesh adaption and mesh deformation, however only the CFD module with continuous adjoint is used here for obtaining flow solutions and objective sensitivities. During optimisation, the primal and adjoint solutions are converged down to an absolute residual of 10^{-8} on density (and corresponding dual), except in cases where the primal has restarted from a previous solution when a tighter tolerance of 10^{-10} is used.

[†] Available as supplementary material at: <https://arc.aiaa.org/doi/suppl/10.2514/1.J054943> [Accessed February 2020]

B. Volume mesh handling

Fig. 8 513×257 O-Mesh generated around NACA 0012 aerofoil

High quality single-block O-meshes are generated here by a conformal mapping approach, with an orthogonality smoothing algorithm. Figure 8 shows three views of the complete 513×257 mesh, used later (L2 mesh in section VI.E); for clarity, only every other point in both directions is shown in the main picture, with every point shown in the two close-ups (the leading and trailing edge images have the same window size). During shape optimisation, mesh deformation is used to produce new meshes for the displaced surface geometry from the initial volume mesh. Not only is this computationally cheaper than regenerating a mesh for each geometry iteration but it also maintains consistency of the discretisation error which is highly desirable during iterative numerical optimisation.

In this work interpolation using multiscale radial basis functions (RBFs) [56] is used. Interpolation using radial basis functions (RBFs) has recently become a prominent mesh deformation method boasting excellent robustness and quality-preserving characteristics [57–59]. Moreover, the method is completely generic, operating on point-clouds alone,

and is perfectly parallel. The multiscale RBF method [56] is particularly effective, both increasing the computational efficiency and improving the system conditioning over conventional or reduced datapoint RBF methods, by using variable length scales depending on boundary point locations.

C. Search algorithm: SNOPT

The SNOPT [46] (Sparse Non-linear Optimiser) package is used here for gradient-based optimisation. This package implements a Sequential Quadratic Programming (SQP) algorithm for solving general non-linear constrained optimisation problems. The power of this package lies in its ability to efficiently and robustly handle large problems (≈ 1000 s of variables and constraints) while allowing precise constraint satisfaction.

The SQP algorithm operates iteratively whereby successive search directions are found from the solution of a quadratic programming (QP) sub-problem and a line-search is used to determine step length. In this work a non-derivative line-search is chosen. The sub-problems are formed from quadratic approximations to the augmented objective function (Lagrangian) and linearisations of the constraints. The quadratic approximation is initialised with an identity matrix and BFGS updates are used to approach the Hessian of the Lagrangian. The quadratic sub-problems are well-posed and the function values and derivatives thereof are easily evaluated; difficulty is not usually encountered when solving these sub-problems.

VI. Inviscid Drag Minimisation

A. Test case: Symmetric transonic drag reduction of NACA0012

This inviscid drag reduction case is one of several benchmark cases laid out by the AIAA Aerodynamic Design Optimisation Discussion Group (ADODG[‡]) to aid in the collaborative evaluation and comparison among shape-optimisation practitioners. The case is defined as minimising the inviscid drag at Mach 0.85 and 0 degree incidence subject to the constraint that the solution shape must lie outside the profile of the initial shape, a NACA0012 aerofoil:

$$\begin{aligned}
 \min_{\alpha} \quad & C_D \\
 \text{s.t.} \quad & y \geq y_0 \\
 & 0.85 \text{ Mach} \\
 & 0^\circ \text{ incidence}
 \end{aligned} \tag{24}$$

This case is chosen since it has been studied extensively in literature due to the unique aerodynamic and geometric challenges it poses. In fact, from a geometric standpoint it is the most challenging test case possible and so makes an ideal benchmark to demonstrate the current method. This does not restrict applicability to RANS-based objectives

---

[‡]<https://sites.google.com/view/mcgill-computational-aerogroup/adodg> [retrieved July 2019]

which generally present as better-conditioned optimisation problems than Euler-based ones due to the additional physical modelling and removal of non-unique solutions. This problem was formulated by Vassberg [60] specifically for studying different optimisation methods since it permits a large design space and requires only moderate computational expense. The case only exhibits drag due to shocks, with Vassberg stating that a shock-free design is anticipated to be impossible. Starting from a baseline of 469 drag counts, Vassberg achieved an optimised result of 103.8 drag counts while also identifying a pathological nature in the symmetric problem in the form of asymmetric instability. A detailed study by Destarac *et al.* [61] additionally demonstrated Mach hysteresis in various optimised shapes due to the existence of non-unique flow solutions. In addition to these numerical challenges, the associated geometric problem is demanding, requiring a high-fidelity and flexible design space as well as a robust and quality-preserving mesh deformation method.

Poole *et al.* [62] used orthogonal modes extracted by singular value decomposition, achieving 85.6 drag counts with only 15 design variables. Carrier *et al.* [45] apply several different optimisation frameworks, eventually obtaining a result of 36.7 drag counts, 1.1 drag counts of which they show corresponds to spurious drag. He *et al.* [63] use a volume B-Spline obtaining a result of 7.60 drag counts with 50 design variables, the lowest published result for comparable mesh resolution (131K cells).

Masters *et al.* compared six different parameterisations on the benchmark problem [12] as well as applying a multilevel subdivision parameterisation [33]. The authors achieved a result of 27.8 drag counts on a 257×257 mesh using single-level B-Splines (16 design variables) and a result of 15.7 drag counts on the same mesh using multilevel subdivision curves; when performed on a 1025×513 mesh the multilevel subdivision scheme was able achieve 4.2 drag counts. Two distinct solution shapes were observed by the authors differing in boat-tail angle and thickness, where the thicker shapes with larger-boat tail angles corresponded to different trailing-edge flow structures with lower overall inviscid drag. Two techniques found to be particularly effective by Masters *et al.* [33] for this test case are also used here: first, the explicit imposition of a symmetry boundary condition on the volume mesh to stabilise convergence of the flow; and second, the use of flow-restarts during optimisation to minimise the influence of flow hysteresis during line-searches.

B. Geometric constraint

Since shape control is permitted in all axes, the constraint on lying outside the initial profile requires some extra consideration. A linearisation is made at each point on the surface and constraints applied such that the displacement of each point must be above the initial local tangent plane, *i.e.* the dot product of displacement and initial normal must be positive. An advantage of this is that the constraint can be implemented linearly and is hence guaranteed to be satisfied (when using SQP). Whereas this constraint is only geometrically approximate, it is found to be sufficiently accurate for well-discretised surface meshes. Moreover this can be considered ‘conservative’ for convex shapes in the sense that when the approximate constraint is active, the real constraint remains strictly feasible.

C. Results: Mesh point control

Fig. 9 Convergence of symmetric transonic drag minimisation using gradient constrained mesh point control.

The transonic test case described above was run on a 257×257 half-mesh with a farfield distance of 50 chords (see Figure 8). Shape control is implemented using the two ordinates of the 257 vertices of the surface mesh such that there are 514 design variables. The surface gradient constraint is applied at each point (equation 22) as a linear constraint in SNOPT and multiple optimisations have been performed for varying values of the continuity parameter σ , summarised in Table 2. In this work, *minor iterations* refer to those of the line-searches performed at each major iteration. The number of minor iterations is equal to the sum of the number of flow solutions and adjoint solutions (where one adjoint solution is required at each major iteration) and hence gives an indication of computational work.

Table 2 Gradient-limited mesh-point optimisation results

σ	Drag (counts)	Major iterations	Minor iterations
25	18.99	43	145
50	11.04	52	179
100	7.68	60	229

Figure 9 shows the optimisation convergence histories for the sweep of parameter values and Figure 10 shows the resulting optimised profiles. All cases recover similar profiles, specifically those exhibiting large boat-tail angles and a maximum thickness at approximately fifty percent chord as found in [12]. The optimised pressure distributions, shown in Figure 11, clearly show the removal of the strong shock and the tendency towards a flat distribution with minimal variation at the leading edge acceleration and trailing edge compression. The same result is also clear in the Mach and pressure contours shown in Figure 12.

Recall that low values of σ correspond to a tighter gradient constraint and hence smaller design space. In the case

of the two tightest gradient constraints considered here, $\sigma = 25$ and $\sigma = 50$, the higher optimal drag values indicate that the respective design spaces may be over-constrained for the shape problem compared to larger values of σ which achieved better results. In the resulting profiles (Figure 10), the lesser performing cases have not achieved as tight a corner at the leading edge and have lower overall thicknesses indicating that the gradient constraint has restricted the short wavelength variations required in these areas. This is confirmed when plotting the constrained surface derivatives, shown in Figure 13. For $\sigma = 25$ and $\sigma = 50$, large portions of the surface constraints are active at the solution, particularly near the leading and trailing edges. In contrast, where $\sigma > 100$, all cases achieved below 8 drag counts with variations of less than a drag count between them. Correspondingly, the constrained surface derivatives are much less active at the solution. In general, as the gradient constraint is loosened (larger σ) the attainable result improves corresponding to a tighter leading edge corner, a maximum thickness closer to 50% and a smaller boat-tail angle. The improved results for larger σ also come at increased computational cost of optimisation (more function calls), however the increase is not substantial and is expected since the feasible design space is larger for larger values of σ .

All cases with $\sigma \geq 100$ achieve comparable results with less than a drag count in variation and a best result of 7.68 drag counts for $\sigma = 100$. The previous best result for this test case, at the same mesh resolution (257×257), is 15.7 drag counts by Masters *et al.* using multilevel subdivision curves [33]. The improvement achieved here is partly attributed to the use of coordinate control (control over both x and z coordinates) which has allowed unrestricted representation of the blunt leading and trailing edges observed in the optimised results. This is also the case for B-Spline control presented in the next section. Coordinate control (of mesh points or spline control points) is made practical by the gradient constraints which, as well as regularising the localised shape control, introduce consideration for surface connectivity by their additional in-plane components. As a result the deformed surface and volume meshes are of higher quality than would be produced by unconstrained coordinate control or thickness modifications alone.

(a)

(b) Maximum thickness location

Fig. 10 Initial and optimised shape of symmetric transonic drag minimisation using gradient-limited mesh point control

Fig. 11 Optimised pressure distributions from gradient-limited mesh-point control (*top*), showing detail at the trailing edge (*bottom*)

(a) NACA 0012

(b) Optimisation result ($\sigma = 100$)

Fig. 12 Mach and pressure fields for baseline and optimised profiles

(a) $\sigma = 25$

(b) $\sigma = 50$

(c) $\sigma = 100$

(d) $\sigma = 150$

(e) $\sigma = 200$

Fig. 13 Constrained third differences at optimal solutions normalised by mesh spacing

D. Results: B-Spline control

Based on the results of the mesh-point optimisations, a value of $\sigma = 100$ was chosen to perform optimisations using B-Spline shape control. The geometrically meaningful reduction in the design space imposed by the gradient constraints mean that the optimisation problem should remain well-posed (by excluding undesirable geometries) even with increases in B-Spline control fidelity. The intended effect of this is to remove dependence of the optimisation convergence rate on the B-Spline control fidelity. This will be demonstrated in this section by comparing optimisation with varying numbers of B-Spline control points for cases with and without the gradient constraints.

Least-squares fitting of the surface geometry is used to generate cubic B-Spline curves at a range of control fidelities where the number of control points varies between 8 and 64. The input surface mesh was parameterised using arc length and the knot sequence chosen to match the underlying point distribution such that more control points are located towards the leading and trailing edges where the mesh densities are highest. Figure 14 shows the resulting control polygon for $N = 12$ control points. The two ordinates of the B-Spline control points are used for shape control such that the number of design variables is twice the number of control points. The B-Spline gradient constraints, which are only applied between spline knots, are implemented as a linear constraint in SNOPT.

Fig. 14 B-Spline control polygons for $N = 12$ control points: initial parameterisation and optimised result ($\sigma = 100$)

Also shown in Figure 14 is the resulting control polygon and aerofoil profile after optimisation on the transonic NACA0012 case. The initial and optimised control polygons in this example demonstrate the importance of performing shape control with all coordinates (both x and z). Note the geometric changes at the leading and trailing edges where there are increases in radius of curvature and surface rotation respectively; in these cases, especially for the surface rotation at the trailing edge, there is significant movement of the control points in the x axis as well as the z . Without this extra freedom, excessive control fidelity would be required to reproduce the same deformation using, for example, vertical only displacements. Moreover the resulting surface discretisation would be of poor quality since it is necessarily

modified as a side-effect of seeking a normal shape-relevant displacement. In contrast, the coordinate shape control used here, in combination with shape constraints, enables efficient and flexible shape relevant perturbations to be made during optimisation while maintaining a high quality surface discretisation.

Figure 15 shows the convergence histories for both constrained and unconstrained cases with varying numbers of control points. Final optimisation results are summarised in Tables 4 and 3. Note that in the gradient-limiting case, all runs are performed with the same value of the constraint parameter, $\sigma = 100$. Immediately it is clear that the rate of convergence under the gradient constraint is independent of the control fidelity (number of control points) of the B-Spline shape control. Importantly however, the optimiser is still able to make use of the increased control fidelity to achieve improvements in the final result; hence the surface gradient constraint does not remove useful geometries from the higher fidelity design spaces.

By contrast, when used without the gradient constraint (Figure 15(a)) the convergence rate deteriorates with increased control fidelity. This behaviour, which is also seen in previous studies [12–14, 33, 64], is expected for increases in problem dimensionality and reflects the good linear-independence of the B-Spline shape basis even with refinement. However the exponential increase in the size of the design space, and consequent deterioration in convergence rate, can be avoided by filtering the allowable contributions from the additional low wavelength shape components; this reduces the effective design space in a geometrically meaningful way such that convergence rate is largely independent of chosen shape control fidelity. *i.e.* the convergence of the optimisation problem is dependent on the underlying physical problem and not the chosen solution method. This finding complements the results of Jameson on the independence of the smoothed steepest descent method to number of design variables [36]. For unconstrained steepest descent, sensitivity smoothing can also be viewed as projecting the gradient into some constraint subspace; hence smoothed steepest descent imposes a minimum length scale on the shape control (determined by the filter radius) and thereby also fixes the convergence rate for shape optimisation.

In addition to improving the convergence rate, the removal of spurious high frequency shape components also affects the path taken by the optimiser which therefore affects the final result. For this test case, the addition of the gradient constraints improves the final optimised objective at all control fidelities. Shown in Figure 16 are the resulting profiles for both unconstrained and constrained cases. For the unconstrained case (Figure 16(a)), the lower control fidelities tested (N8, N12, N16) do not have sufficient flexibility whereas the high control fidelities (N24, N32, N64) are dominated by a ‘step’ feature on the aft-surface and oscillations at the wavelength of the control points. The aft-surface feature appears spuriously in response to the shock in the initial flow; however this short wavelength phenomenon is undesirable from a geometric point of view and importantly, it leads to a strong local minima. The surface oscillations in the unconstrained results are most clearly evident when considering the optimised pressure distributions, shown in Figure 17, due to their analogy with surface curvature. In comparison, when optimised with surface gradient constraints, the B-Spline control polygons remain regular, without forming short wavelength features or oscillations on the surface.

This demonstrates the ability of the gradient constraint to sufficiently regularise the B-Spline basis while retaining the desirable property of a localised well-conditioned basis needed for high-fidelity shape control.

Fig. 15 Convergence of symmetric transonic drag minimisation using B-Spline control

Figure 18 shows the constrained spline third derivatives at the optimal solutions. For all cases, the gradient constraint is inactive for the majority of the spline knots across which it is applied, indicating that the limiting factor in the shape optimisation problem is the fidelity (number of control points) of the spline and not the gradient constraint. This further demonstrates the effect of the constraint in constraining the design space such that the shape problem is well-posed without limiting the attainable result.

Table 3 Unconstrained B-Spline optimisation results

No. control points	Drag (counts)	Major iterations	Minor iterations
8	97.83	10	107
12	51.88	67	315
16	36.01	54	345
24	129.0	36	227
32	127.6	27	174
64	222.2	55	268

Table 4 Gradient-limited B-Spline optimisation results ($\sigma = 100$)

No. control points	Drag (counts)	Major iterations	Minor iterations
8	52.53	10	31
12	18.98	203	542
16	11.77	46	173
24	8.14	41	187
32	8.25	38	171
64	7.88	62	271

(a) Without gradient constraints

(b) Gradient constrained ($\sigma = 100$)

Fig. 16 Initial and optimised profiles for symmetric transonic drag minimisation using B-Spline control

Fig. 17 Initial and optimised pressure distributions for symmetric transonic drag minimisation using B-Spline control

Fig. 18 Constrained spline third derivatives at optimal solutions

E. Effect of mesh resolution

The effect of mesh resolution was also investigated, with a fixed number of B-Spline control points ($N = 64$) and fixed constraint parameter ($\sigma = 100$) such that the shape control fidelity is held constant. Three mesh resolutions were investigated, summarised in Table 5; the grid dimensions indicated are $N_i \times N_j$ corresponding to number of grid points on the half surface and to the farfield respectively. For all resolutions considered here, the optimum drag values decrease monotonically as the mesh is refined. Figure 19 shows convergence of the objective during optimisation; despite the number of major iterations varying between mesh resolutions, the primary rate of convergence during the initial phase of optimisation, where the majority of objective reduction is achieved, is approximately constant between mesh resolutions. In the same way as for increased B-Spline control fidelity, improved results have been achieved while bounding the convergence of the optimisation. However in this case, the improvement in final objective is attributable to increased resolution of the physical problem and its gradients. At the highest mesh density an optimised value of 1.60 drag counts is achieved; to the authors' knowledge this is by far the lowest result to be published.

Table 5 Mesh resolution results

	Mesh resolution	Drag (counts)	Major iterations	Minor iterations
L1	129×129	21.20	39	171
L2	257×257	7.88	62	271
L3	513×513	1.60	73	308

Fig. 19 Optimisation convergence using gradient constrained B-Spline control at varying mesh resolutions.

VII. Conclusions

A new shape control method has been presented that permits local shape design optimisation and high-fidelity refinement without an exponential increase in the size of the design space and consequent degradation in optimisation convergence rate. Shape constraints have been derived to ensure surface smoothness and mesh validity when using local shape control for optimisation and which reduce the effective degrees of freedom in a geometrically meaningful manner; unlike existing shape control methods, the constraint-based approach is not restricted to low control fidelities or require the manual definition of shape functions and deformation modes. The resulting shape optimisation problem is well-posed, as the design space is sufficiently bounded, and well-conditioned since local control methods (mesh-points and B-Splines) are used. The local control methods with coordinate (x,z) control provide high-fidelity shape-relevant control, whereas the surface constraints reduce the effective degrees of freedom by excluding non-physical shapes from the design space. Moreover the constraint formulation also implicitly controls mesh validity by taking into account the in-plane surface variation. The constraint-based approach has been tested on a challenging two-dimensional inviscid drag minimisation problem and the effects of the constraint bound magnitude and varying control fidelities have been investigated. When used with B-Spline curves, the gradient-constraint excludes control polygons with excessive deformations or oscillatory arrangements and ensures the underlying discretisation remains smooth. Furthermore the regularised B-Spline shape control is shown to have an optimisation convergence rate independent of numerical grid resolution as well as shape control fidelity, while still making use of the increased design space to achieve improved results. As a result, high-fidelity shape optimisation is possible at a reasonable computational cost. At the highest mesh density the gradient limiting method applied with a 64 control point B-Spline curve achieved an optimised result of 1.60 drag counts for the standard test case, the lowest published result so far.

Acknowledgments

This work was supported by the EPSRC and the University of Bristol through Doctoral Training Grant EP/M507994/1. This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol - <http://www.bristol.ac.uk/acrc>.

References

- [1] Bisson, F., and Nadarajah, S., "Adjoint-Based Aerodynamic Optimization of Benchmark Problems," *53rd AIAA Aerospace Sciences Meeting, Kissimmee, Florida*, 2015. URL <https://doi.org/10.2514/6.2015-1948>.
- [2] Rizzi, A. W., Zhang, M., Bisson, F., Nadarajah, S., and Vos, J. B., "Comparative Study of Two Optimization Frameworks Applied to Case III: Induced-Drag Minimization," *54th AIAA Aerospace Sciences Meeting, San Diego, California*, 2016. URL <https://arc.aiaa.org/doi/abs/10.2514/6.2016-1296>.

- [3] Gagnon, H., and Zingg, D. W., “Two-Level Free-Form and Axial Deformation for Exploratory Aerodynamic Shape Optimization,” *AIAA Journal*, Vol. 53, No. 7, 2015, pp. 2015–2026. URL <https://doi.org/10.2514/1.J053575>.
- [4] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark,” *AIAA Journal*, Vol. 53, No. 4, 2015, pp. 968–985. URL <https://doi.org/10.2514/1.J053318>.
- [5] Lee, C., Koo, D., and Zingg, D. W., “Comparison of B-Spline Surface and Free-Form Deformation Geometry Control for Aerodynamic Optimization,” *AIAA Journal*, Vol. 55, No. 1, 2017, pp. 228–240. URL <https://doi.org/10.2514/1.J055102>.
- [6] Dumont, A., Méheut, M., Baumgärtner, D., and Bletzinger, K.-U., “Aerodynamic Shape Optimization Progress on ADODG Benchmark Problems Using the elsA Software,” *35th AIAA Applied Aerodynamics Conference, Denver Colorado*, 2017. URL <https://arc.aiaa.org/doi/abs/10.2514/6.2017-4081>.
- [7] Chen, S., Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “Aerodynamic Shape Optimization of Common Research Model Wing–Body–Tail Configuration,” *Journal of Aircraft*, Vol. 53, No. 1, 2016, pp. 276–293. URL <https://doi.org/10.2514/1.C033328>.
- [8] Kenway, G. K. W., and Martins, J. R. R. A., “Buffet-Onset Constraint Formulation for Aerodynamic Shape Optimization,” *AIAA Journal*, Vol. 55, No. 6, 2017, pp. 1930–1947. URL <https://doi.org/10.2514/1.J055172>.
- [9] Jonsson, E., Riso, C., Lupp, C. A., Cesnik, C. E., Martins, J. R., and Epureanu, B. I., “Flutter and post-flutter constraints in aircraft design optimization,” *Progress in Aerospace Sciences*, Vol. 109, 2019, p. 100537. URL <https://doi.org/10.1016/j.paerosci.2019.04.001>.
- [10] Reist, T. A., Koo, D., Zingg, D. W., Bochud, P., Castonguay, P., and Leblond, D., “Cross Validation of Aerodynamic Shape Optimization Methodologies for Aircraft Wing-Body Optimization,” *AIAA Journal*, Vol. Article in Advance, Published online 6 February, 2020. URL <https://doi.org/10.2514/1.J059091>.
- [11] Masters, D. A., Taylor, N. J., Rendall, T., Allen, C. B., and Poole, D. J., “A Geometric Comparison of Aerofoil Shape Parameterisation Methods,” *AIAA Journal*, Vol. 55, No. 5, 2017, pp. 1575–1589. URL <https://doi.org/10.2514/1.J054943>.
- [12] Masters, D. A., Poole, D. J., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., “Influence of Shape Parameterization on a Benchmark Aerodynamic Optimization Problem,” *Journal of Aircraft*, Vol. 54, No. 6, 2017, pp. 2242–2256. URL <https://doi.org/10.2514/1.C034006>.
- [13] Reuther, J., and Jameson, A., “A Comparison of Design Variables for Control Theory Based Airfoil Optimization,” *Technical Report, NASA Research Institute for Advanced Computer Science*, 1995.
- [14] Castonguay, P., and Nadarajah, S., “Effect of Shape Parameterization on Aerodynamic Shape Optimization,” *45th AIAA Aerospace Sciences Meeting and Exhibit*, 2007, pp. 1–20. URL <https://doi.org/10.2514/6.2007-59>.

- [15] Jameson, A., "Aerodynamic Design via Control Theory," *Journal of Scientific Computing*, Vol. 3, No. 3, 1988. URL <https://doi.org/10.1007/BF01061285>.
- [16] Mohammadi, B., and Pironneau, O., "Shape Optimization in Fluid Mechanics," *Annual Review of Fluid Mechanics*, Vol. 36, No. 1, 2004, pp. 255–279. URL <https://doi.org/10.1146/annurev.fluid.36.050802.121926>.
- [17] Schmidt, S., "Efficient Large Scale Aerodynamic Design Based on Shape Calculus," Ph.D. thesis, University of Trier, Germany, 2010. URL http://ubt.opus.hbz-nrw.de/frontdoor.php?source_opus=569&la=en.
- [18] Le, C., Bruns, T., and Tortorelli, D., "A gradient-based, parameter-free approach to shape optimization," *Computer Methods in Applied Mechanics and Engineering*, Vol. 200, No. 9-12, 2011, pp. 985–996. URL <https://doi.org/10.1016/j.cma.2010.10.004>.
- [19] Bletzinger, K. U., "A consistent frame for sensitivity filtering and the vertex assigned morphing of optimal shape," *Structural and Multidisciplinary Optimization*, Vol. 49, No. 6, 2014, pp. 873–895. URL <https://doi.org/10.1007/s00158-013-1031-5>.
- [20] Braibant, V., and Fleury, C., "Shape optimal design using B-splines," *Computer Methods in Applied Mechanics and Engineering*, Vol. 44, No. 3, 1984, pp. 247–267. URL [https://doi.org/10.1016/0045-7825\(84\)90132-4](https://doi.org/10.1016/0045-7825(84)90132-4).
- [21] Desideri, J.-A., "Hierarchical Optimum-Shape Algorithms using Embedded Bezier Parameterizations," *Numerical Methods for Scientific Computing Variational Problems and Applications*, edited by Y. Kuznetsov, P. Neittanmaki, and . O. Pironneau (CIMNE, Barcelona, 2003).
- [22] Sederberg, T. W., and Parry, S. R., "Free-form deformation of solid geometric models," *ACM SIGGRAPH Computer Graphics*, Vol. 20, No. 4, 1986, pp. 151–160. URL <https://doi.org/10.1145/15886.15903>.
- [23] Hicken, J. E., and Zingg, D. W., "Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement," *AIAA Journal*, Vol. 48, No. 2, 2010, pp. 400–413. URL <https://doi.org/10.2514/1.44033>.
- [24] Morris, A. M., Allen, C. B., and Rendall, T. C. S., "High-fidelity aerodynamic shape optimization of modern transport wing using efficient hierarchical parameterization," *International Journal for Numerical Methods in Fluids*, Vol. 63, 2010, pp. 297–312. URL <https://doi.org/10.1002/flid.2067>.
- [25] Hicks, R. M., and Henne, P. A., "Wing Design by Numerical Optimization," *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412. URL <https://doi.org/10.2514/3.58379>.
- [26] Kulfan, B. M., "A Universal Parametric Geometry Representation Method – "CST"," *45th AIAA Aerospace Sciences Meeting and Exhibit*, 2007. URL <https://doi.org/10.2514/1.29958>.
- [27] Belegundu, A., and Rajan, S. D., "A shape optimization approach based on natural design variables and shape functions," *Computer Methods in Applied Mechanics and Engineering*, Vol. 66, No. 1, 1988, pp. 87–106. URL [https://doi.org/10.1016/0045-7825\(88\)90061-8](https://doi.org/10.1016/0045-7825(88)90061-8).

- [28] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Metric-Based Mathematical Derivation of Efficient Airfoil Design Variables," *AIAA Journal*, Vol. 53, No. 5, 2015, pp. 1349–1361. URL <https://doi.org/10.2514/1.J053427>.
- [29] Bloor, M. I. G., and Wilson, M. J., "Efficient parametrization of generic aircraft geometry," *Journal of Aircraft*, Vol. 32, No. 6, 1995, pp. 1269–1275. URL <https://doi.org/10.2514/3.46874>.
- [30] Payot, A. D., Rendall, T. C., and Allen, C. B., "Restricted snakes volume of solid (RSVS): A parameterisation method for topology optimisation of external aerodynamics," *Computers and Fluids*, Vol. 182, 2019, pp. 60–84. URL <https://doi.org/10.1016/j.compfluid.2019.02.008>.
- [31] Payot, A. D., Rendall, T., and Allen, C. B., "Parametric Surfaces with Volume of Solid Control for Optimisation of Three Dimensional Aerodynamic Topologies," *AIAA Scitech 2019 Forum*, San Diego, California, 2019. URL <https://doi.org/10.2514/6.2019-1472>.
- [32] Bandara, K., Ruberg, T., and Cirak, F., "Shape optimisation with multiresolution subdivision surfaces and immersed finite elements," *Computer Methods in Applied Mechanics and Engineering*, Vol. 300, 2016, pp. 510–539. URL <https://doi.org/10.1016/j.cma.2015.11.015>.
- [33] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., "Multilevel Subdivision Parameterization Scheme for Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 55, No. 10, 2017, pp. 1–16. URL <https://doi.org/10.2514/1.J055785>.
- [34] Han, X., and Zingg, D. W., "An adaptive geometry parametrization for aerodynamic shape optimization," *Optimization and Engineering*, Vol. 15, No. 1, 2014, pp. 69–91. URL <https://doi.org/10.1007/s11081-013-9213-y>.
- [35] Anderson, G. R., Nemec, M., and Aftosmis, M. J., "Aerodynamic Shape Optimization Benchmarks with Error Control and Automatic Parameterization," *53rd AIAA Aerospace Sciences Meeting*, 2015, pp. 1–18. URL <https://doi.org/10.2514/6.2015-1719>.
- [36] Jameson, A., and Vassberg, J. C., "Studies of Alternative Numerical Optimization Methods Applied to the Brachistochrone Problem," *Computational Fluid Dynamics Journal*, Vol. 9, No. 3, 2000.
- [37] Stück, A., and Rung, T., "Adjoint RANS with filtered shape derivatives for hydrodynamic optimisation," *Computers and Fluids*, Vol. 47, No. 1, 2011, pp. 22–32. URL <https://doi.org/10.1016/j.compfluid.2011.01.041>.
- [38] Vazquez, M., Dervieux, A., and Koobus, B., "Multilevel optimization of a supersonic aircraft," *Finite Elements in Analysis and Design*, Vol. 40, 2004, pp. 2101–2124. URL <https://doi.org/10.1016/j.fva.2004.06.001>.
- [39] Hojjat, M., Stavropoulou, E., and Bletzinger, K. U., "The Vertex Morphing method for node-based shape optimization," *Computer Methods in Applied Mechanics and Engineering*, Vol. 268, 2014, pp. 494–513. URL <https://doi.org/10.1016/j.cma.2013.10.015>.

- [40] Poole, D. J., Allen, C. B., and Rendall, T. C., “High-fidelity aerodynamic shape optimization using efficient orthogonal modal design variables with a constrained global optimizer,” *Computers and Fluids*, Vol. 143, 2017, pp. 1–15. URL <https://doi.org/10.1016/j.compfluid.2016.11.002>.
- [41] Beux, F., and Dervieux, A., “A hierarchical approach for shape optimization,” *Engineering Computations*, Vol. 11, No. 1, 1994, pp. 25–48. URL <https://doi.org/10.1108/02644409410799191>.
- [42] Majd, B. A. E., Désidéri, J.-A., and Duvigneau, R., “Multilevel strategies for parametric shape optimization in aerodynamics,” *European Journal of Computational Mechanics*, Vol. 17, No. 1-2, 2008, pp. 149–168. URL <https://doi.org/10.3166/remn.17.149-168>.
- [43] Vassberg, J. C., and Jameson, A., “Industrial Applications of Aerodynamic Shape Optimization,” *VKI Lecutre-II, Brussels, Belgium*, Von Karman Institute, 2014. URL <http://aero-comlab.stanford.edu/Papers/Vassberg.VKI-Lecture02.2014.pdf>.
- [44] Jaworski, A., and Müller, J. D., “Toward modular multigrid design optimisation,” *Advances in Automatic Differentiation. Lecture Notes in Computational Science and Engineering*, Vol. 64, Springer, Berlin, Heidelberg, 2008, pp. 281–291. URL https://doi.org/10.1007/978-3-540-68942-3_25.
- [45] Carrier, G., Destarac, D., Dumont, A., Meheut, M., Salah El Din, I., Peter, J., Ben Khelil, S., Brezillon, J., and Pestana, M., “Gradient-Based Aerodynamic Optimization with the elsA Software,” *52nd AIAA Aerospace Sciences Meeting*, 2014, pp. 1–31. URL <https://doi.org/10.2514/6.2014-0568>.
- [46] Gill, P. E., Murray, W., and Saunders, M. A., “SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization,” *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131. URL <https://doi.org/10.1137/S0036144504446096>.
- [47] Morris, A. M., Allen, C. B., and Rendall, T. C. S., “CFD-based optimization of aerofoils using radial basis functions for domain element parameterization and mesh deformation,” *International Journal for Numerical Methods in Fluids*, Vol. 58, 2008, pp. 827–860. URL <https://doi.org/10.1002/flid.176>.
- [48] Niordson, F., “Optimal design of elastic plates with a constraint on the slope of the thickness function,” *International Journal of Solids and Structures*, Vol. 19, No. 2, 1983, pp. 141–151. URL [https://doi.org/10.1016/0020-7683\(83\)90005-7](https://doi.org/10.1016/0020-7683(83)90005-7).
- [49] Petersson, J., and Sigmund, O., “Slope constrained topology optimization,” *International Journal for Numerical Methods in Engineering*, Vol. 41, No. 8, 1998, pp. 1417–1434. URL [https://doi.org/10.1002/\(SICI\)1097-0207\(19980430\)41:8<1417::AID-NME344>3.0.CO;2-N](https://doi.org/10.1002/(SICI)1097-0207(19980430)41:8<1417::AID-NME344>3.0.CO;2-N).
- [50] Xu, S., Jahn, W., and Müller, J.-D., “CAD-based shape optimisation with CFD using a discrete adjoint,” *International Journal for Numerical Methods in Fluids*, Vol. 74, No. 3, 2014, pp. 153–168. URL <https://doi.org/10.1002/flid.3844>.
- [51] Cinquegrana, D., and Iuliano, E., “Investigation of adaptive design variables bounds in dimensionality reduction for aerodynamic shape optimization,” *Computers and Fluids*, Vol. 174, No. July, 2018, pp. 89–109. URL <https://doi.org/10.1016/j.compfluid.2018.07.012>.

- [52] *Curves and surface for computer aided geometric design : a practical guide*, Academic Press, Inc., 1993.
- [53] Gill, P. E., and Murray, W., “Numerically stable methods for quadratic programming,” *Mathematical Programming*, Vol. 14, No. 1, 1978, pp. 349–372. URL <https://doi.org/10.1007/BF01588976>.
- [54] Palacios, F., Colonno, M. R., Aranake, A. C., Campos, A., Copeland, S. R., Economon, T. D., Lonkar, A. K., Lukaczyk, T. W., Taylor, T. W. R., and Alonso, J. J., “Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design,” *51st AIAA Aerospace Sciences Meeting*, 2013. URL <https://doi.org/10.2514/6.2013-287>.
- [55] Economon, T. D., Alonso, J. J., Albring, T. A., and Gauger, N. R., “Adjoint Formulation Investigations of Benchmark Aerodynamic Design Cases in SU2,” *35th AIAA Applied Aerodynamics Conference*, 2017, pp. 1–13. URL <https://doi.org/10.2514/6.2017-4363>.
- [56] Kedward, L., Allen, C. B., and Rendall, T. C., “Efficient and exact mesh deformation using multiscale RBF interpolation,” *Journal of Computational Physics*, Vol. 345, 2017, pp. 732–751. URL <https://doi.org/10.1016/j.jcp.2017.05.042>.
- [57] de Boer, A., van der Schoot, M. S., and Bijl, H., “Mesh deformation based on radial basis function interpolation,” *Computers and Structures*, Vol. 85, No. 11-14, 2007, pp. 784–795. URL <https://doi.org/10.1016/j.compstruc.2007.01.013>.
- [58] Rendall, T. C. S., and Allen, C. B., “Unified fluid-structure interpolation and mesh motion using radial basis functions,” *International Journal for Numerical Methods in Engineering*, Vol. 74, No. 10, 2008. doi:<https://doi.org/10.1002/nme.2219>.
- [59] Rendall, T. C., and Allen, C. B., “Reduced surface point selection options for efficient mesh deformation using radial basis functions,” *Journal of Computational Physics*, Vol. 229, No. 8, 2010, pp. 2810–2820. URL <https://doi.org/10.1016/j.jcp.2009.12.006>.
- [60] Vassberg, J., Harrison, N., Roman, D., and Jameson, A., “A Systematic Study on the Impact of Dimensionality for a Two-Dimensional Aerodynamic Optimization Model Problem,” *29th AIAA Applied Aerodynamics Conference*, 2011, pp. 1–19. URL <https://doi.org/10.2514/6.2011-3176>.
- [61] Destarac, D., Carrier, G., Anderson, G. R., Nadarajah, S., Poole, D. J., Vassberg, J. C., and Zingg, D. W., “Example of a Pitfall in Aerodynamic Shape Optimization,” *AIAA Journal*, Vol. 56, No. 4, 2018, pp. 1–9. URL <https://doi.org/10.2514/1.J056128>.
- [62] Poole, D. J., Allen, C. B., and Rendall, T., “Control Point-Based Aerodynamic Shape Optimization Applied to AIAA ADODG Test Cases,” *53rd AIAA Aerospace Sciences Meeting*, 2015. doi:<https://doi.org/10.2514/6.2015-1947>.
- [63] He, X., Li, J., Mader, C. A., Yildirim, A., and Martins, J. R., “Robust aerodynamic shape optimization—From a circle to an airfoil,” *Aerospace Science and Technology*, Vol. 87, 2019, pp. 48 – 61. doi:<https://doi.org/10.1016/j.ast.2019.01.051>.

- [64] Nadarajah, S., Castonguay, P., and Mousavi, A., “Survey of Shape Parameterization Techniques and its Effect on Three-Dimensional Aerodynamic Shape Optimization,” *18th AIAA Computational Fluid Dynamics Conference*, 2007, pp. 1–23. URL <https://doi.org/10.2514/6.2007-3837>.