Impact of Shape Parameterisation on Aerodynamic Optimisation of Benchmark Problem

D. A. Masters*, D. J. Poole †,
Department of Aerospace Engineering, University of Bristol

N. J. Taylor‡,
MBDA UK Ltd, Filton

T. C. S. Rendall§ and C. B. Allen¶
Department of Aerospace Engineering, University of Bristol

This paper presents an investigation into the influence of shape parameterisation and dimensionality on the optimisation of a benchmark case described by the AIAA Aerodynamic Design Optimisation Discussion Group. This problem specifies the drag minimisation of a NACA0012 under inviscid flow conditions at $M = 0.85$ and $\alpha = 0$ subject to the constraint that local thickness must only increase. The work presented here applies six different shape parameterisation schemes to this optimisation problem with between 4 and 40 design variables. The parameterisation methods used are: Bèzier Surface FFD; B-Splines; CSTs; Hicks-Henne bump functions; a Radial Basis Function domain element method (RBF-DE) and a Singular Value Decomposition (SVD) method. The optimisation framework used consists of a gradient based SQP optimiser coupled with the SU^2 adjoint Euler solver which enables the efficient calculation of the design variable gradients. Results for the all the parameterisation methods are presented with the best results for each method ranging between 25 and 56 drag counts from an initial value of 469. The optimal result was achieved with the B-Spline method with 16 design variables. A further validation of the results is then presented and the presence of hysteresis is explored.

I. Introduction

With the rise in the use of optimisation techniques in aerodynamic design, a significant effort is being made to improve the effectiveness of all of its constituent parts. To aide this effort a series of benchmark cases have been set out by AIAA Aerodynamic Design Optimisation Discussion Group (ADODG)[a] to allow for comparisons to be made between different optimisation frameworks. This paper investigates the influence of dimensionality and shape parameterisation on the optimisation of ADODG benchmark case one, an inviscid, symmetric drag minimisation case.

This case has previously been investigated with a variety of different parameterisation methods: Bèzier Curves[1, 2]; B-Splines[3, 4, 5, 6, 7]; NURBS[8]; Class/Shape Transformations (CSTs)[9]; Hicks-Henne bump functions[10]; Bèzier Surface FFD[11]; PARSEC[12]; Radial basis function domain element (RBF-DE) deformation[13, 14] and Singular Value Decompositions (SVD)[14, 15]. Due to the large range of factors influencing the results it is difficult to isolate contribution of the parameterisation method from these previous studies. For example, the choice of flow solver[16] or optimisation algorithm[14] has been shown to have a significant impact on the optimisation results for this case. For this reason this work tests this case with a large range of parameterisation methods under the same optimisation framework across a range of design variables.

The aim of this paper is to apply a variety of shape parameterisation techniques to ADODG benchmark case one and asses the suitability of each method with between four and 40 design variables. This has been carried out using the gradient based SQP optimiser SNOPT[17] coupled with the SU^2[18] Euler solver, calculating the design variable gradients using the adjoint method. Further analysis and validation has then be applied to selected final optimised shapes to further understand the flow behaviour and hysteresis at the specified design conditions.

*Graduate Student, AIAA Student Member, dominic.masters@bristol.ac.uk, Bristol, BS8 1TR, UK
†Graduate Student, AIAA Student Member, d.j.poole@bristol.ac.uk, Bristol, BS8 1TR, UK
‡Capability Leader, Aerodynamic Tools & Methods, AIAA Senior Member, nigel.j.taylor@mbda-systems.com, WG3, PO Box 5, Filton, Bristol, BS34 7QW, UK
§Lecturer, AIAA Member, thomas.rendall@bristol.ac.uk, Bristol, BS8 1TR, UK
¶Professor of Computational Aerodynamics, AIAA Senior Member, c.b.allen@bristol.ac.uk, Bristol, BS8 1TR, UK
[a]https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG

II. NACA0012 Inviscid Drag Optimisation

A. Case Specifications

The benchmark case[19] considered in this paper is the inviscid, drag minimisation of a NACA0012 with freestream Mach number (M) of 0.85 subject to a thickness constraint at zero degrees incidence (α). Formally described as

$$\begin{aligned} \text{Minimise} \quad & C_D \\ \text{subject to:} \quad & M = 0.85, \\ & \alpha = 0, \\ & z \geq z^{baseline} \quad \text{on upper surface } \forall x \in [0,1], \\ & z \leq z^{baseline} \quad \text{on lower surface } \forall x \in [0,1]. \end{aligned}$$

This problem is based on work by Vassberg *et al.*[2] though with a slight modification to the baseline NACA0012 geometry to ensure a sharp trailing edge. It is defined as:

$$z^{baseline} = \pm 0.6(0.2969\sqrt{x} - 0.1260x - 0.3516x^2 + 0.2843x^3 - 0.136x^4). \tag{1}$$

B. Previous Work

This test case has been investigated a significant number of times previously and in all cases the drag has been reduced by increasing the aerofoil thickness aft of the max thickness point. Meheut[20] compared a number of optimised aerofoil shapes from six different institutes[4, 5, 7, 13, 14, 16] with the same flow solver. This produced a range of final drag results (all similar to values declared in the original publications) ranging from 32 to 86 drag counts on a mesh with 1024 points around the aerofoil. This paper also highlighted a number of unexpected flow characteristics present for these highly optimised aerofoils. Most notable was the appearance of hysteresis at the specified Mach number. They showed that for the aerofoil produced by Carrier *et al.*[4] two possible solutions existed, one with ~60 drag counts and one with ~110 drag counts. Similar behaviour was found by Lee *et al.*[7] who found solutions with ~40 and ~130 drag counts; an even larger difference in drag. They also found solutions with non-zero lift and an asymmetric flow field despite the fully symmetrical problem and mesh. They cited the occurrences of these unexpected lifting solutions (which produce high drag values) in the gradient evaluations as a contributor to some of the optimisation procedures stagnating prematurely. Nadarajah[5] found similar problems, with the flow converging to oscillatory results due to the existence of multiple solutions. This was however overcome by enforcing a symmetry boundary condition along the wake. This single change improved the final result of an optimisation procedure significantly, reducing the drag from 55.3 to 37.5 counts on a 768×128 mesh. This result was then further reduced to 25.2 counts when run on a super-fine 3072×512 mesh.

III. Parameterisation Methods

This work considers the impact of parameterisation on the case outlined above. The six methods considered are: Bèzier Surface FFD; B-Splines; CSTs; Hicks-Henne bump functions; a Radial Basis Function domain element method (RBF-DE) and a Singular Value Decomposition (SVD) method. This section outlines these methods and their definitions.

A. Bèzier Surface FFD

A Bèzier surface is a B-spline surface of Bèzier curves. Although these are usually used to create surfaces in three-dimensional space, they can also be used as a deformation tool in two-dimensional space by constraining the control points to a plane. To create a deformable domain from this surface a rectangular lattice of $(m+1) \times (n+1)$ uniformly spaced control points, $\mathbf{P}_{ij}$, is placed around an initial aerofoil. Then, given an undeformed domain $\mathcal{A}(x,z) \in [x_{min}, x_{max}] \times [z_{min}, z_{max}]$, the initial control point positions are defined as

$$\mathbf{P}_{ij}^{initial} = \left(x_{min} + \frac{i}{m}\left(x_{max} - x_{min}\right), z_{min} + \frac{j}{n}\left(z_{max} - z_{min}\right), 0 \right) \tag{2}$$

for $i = 0, \ldots, m$, $j = 0, \ldots, n$.

The two-dimensional Bèzier surface, $\Psi(u,v)$, spanning the deformed domain $\mathcal{D}(x,z)$, is then given by

$$\Psi(u,v) = \sum_{j=0}^{n} \sum_{i=0}^{m} B_{i,m}(u) B_{j,n}(v) \mathbf{P}_{ij} \tag{3}$$

where $u, v \in [0, 1]$ and $B_{i,m}$ are Bernstein polynomials.

To create the one-to-one deformation mapping required, the undeformed domain $\mathcal{A}$ is then normalised to the unit domain $\mathcal{N}(u, v) \in [0, 1] \times [0, 1]$ by the transformation

$$u(x) = \frac{x - z_{min}}{z_{max} - z_{min}}, \qquad v(z) = \frac{z - x_{min}}{x_{max} - x_{min}}, \tag{4}$$

which implies that the required deformation transformation, from $\mathcal{A}(x, z) \rightarrow \mathcal{D}(x, z)$, is given by

$$\Psi(u(x), v(z)) = \sum_{j=0}^{n} \sum_{i=0}^{m} B_{i,m}(u(x)) B_{j,n}(v(z)) \mathbf{P}_{ij}. \tag{5}$$

The deformation of the initial aerofoil with respect to the control point positions P_{ij} is therefore defined as

$$\mathbf{X} = \sum_{j=0}^{n} \sum_{i=0}^{m} B_{i,m}(u(\mathbf{x}^{initial})) B_{j,n}(\mathbf{z}^{initial}) \mathbf{P}_{ij}. \tag{6}$$

For this paper this method has been implemented with four rows in the control point lattice. This configuration has been used as it was shown to provide the most efficient coverage of the aerofoil design space[21]. Movement of each control point has been restricted to the z-direction and symmetry has been preserved by pairing equivalent upper and lower control points symmetrically to a single design variable. The number of design variables for each case is therefore equal to twice the number of columns in the control point lattice. Figure 1a shows an example deformation for an 8 design variable configuration with associated basis functions shown in 1b.

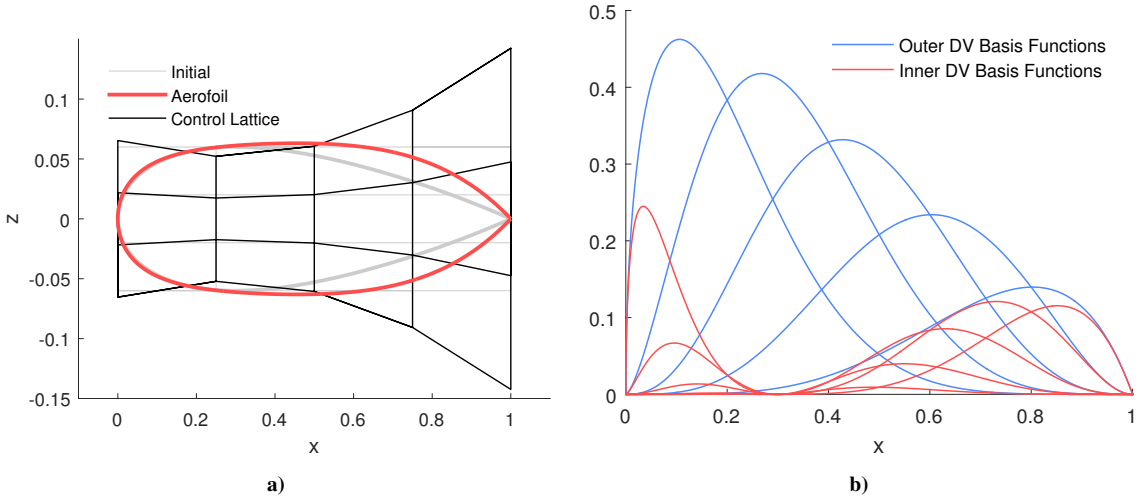

Figure 1. Example deformation of 4×6 Bèzier surface control lattice and NACA0012 where (a) shows the surface and control point positions and (b) shows the basis functions for the symmetric design variables that act on the upper surface of the aerofoil.

B. B-Splines

B-Splines are a widely used method for producing piecewise polynomial curves. Much like many of the other parametrisation methods, B-Splines rely on a set of basis functions, however in this case, the coefficients are defined spatially by a set of discrete control points $\mathbf{P}_i \in \mathbb{R}^3$. Given these assumptions a B-Spline curve parametrised by the scalar $u \in [r_0, r_l]$, is defined as

$$G(u) = \sum_{i=0}^{n-1} N_{i,p}(u) \mathbf{P}_i \tag{7}$$

where the $n = l - k$ basis functions of order k are given by

$$N_{i,0}(u) = \begin{cases} 1 & r_i \leq u < r_{i+1} \\ 0 & \text{otherwise} \end{cases}, \tag{8}$$

$$N_{i,k}(u) = \frac{u - r_i}{r_{i+k} - r_i} N_{i,k-1}(u) + \frac{r_{i+k+1} - u}{r_{i+k+1} - r_{i+1}} N_{i+1,k-1}(u), \tag{9}$$

with the increasing knot vector

$$\mathbf{r} = [r_0, \ldots, r_l], \qquad r_i \leq r_{i+1}. \tag{10}$$

The knot vector then takes the form

$$\mathbf{u} = [\underbrace{r_0, \ldots, r_0}_{k+1}, r_{k+1}, \ldots, r_{n-1}, \underbrace{r_l, \ldots, r_l}_{k+1}] \tag{11}$$

where the values $r_{k+1}, \ldots, r_{n-1}$ indicate the knot points where the basis functions meet. In this case, the B-Spline uniformity requires that the knots $r_k, \ldots, r_n$ are equally distributed. Additionally if $k = n$, and consequently $l = 0$, the B-Splines are called 'Bèzier Curves' and if $u \in [0, 1]$ the basis functions are the Bernstein polynomials of order $k - 1$.

A useful property of B-Splines is that the basis order k controls the locality of the influence of the control points, meaning that for a low order curve the influence of any change in control point position will be more localised compared to if a high order curve was used. Using low order splines can, however, impact the smoothness and overall fidelity of the curve.

B-splines can be used to represent aerofoils in a variety of different ways but for this study each aerofoil is represented by two distinct B-splines. For each B-spline $\mathbf{P}_0$ is fixed at the leading edge $(0, 0)$, $\mathbf{P}_{n+1}$ is at the trailing edge $(1, 0)$ and $\mathbf{P}_1$ is aligned vertically with the leading edge. The other points $\mathbf{P}_i$ are then distributed on a half cosine scale between (0,1) in the chord-wise direction and only allowed to vary in the vertical direction, i.e.

$$\mathbf{P}_0 = (0, 0), \quad \mathbf{P}_i = \left(\frac{1}{2}\left[1 - \cos\left(\frac{\pi(i-1)}{n+1}\right)\right], a_i\right), \quad \mathbf{P}_{n+1} = (1, 0), \tag{12}$$

where a_i denotes a design variable. Again symmetrical control points are paired to create one design variable that controls the local thickness of the aerofoil. Figure 2a shows an example of this configuration.

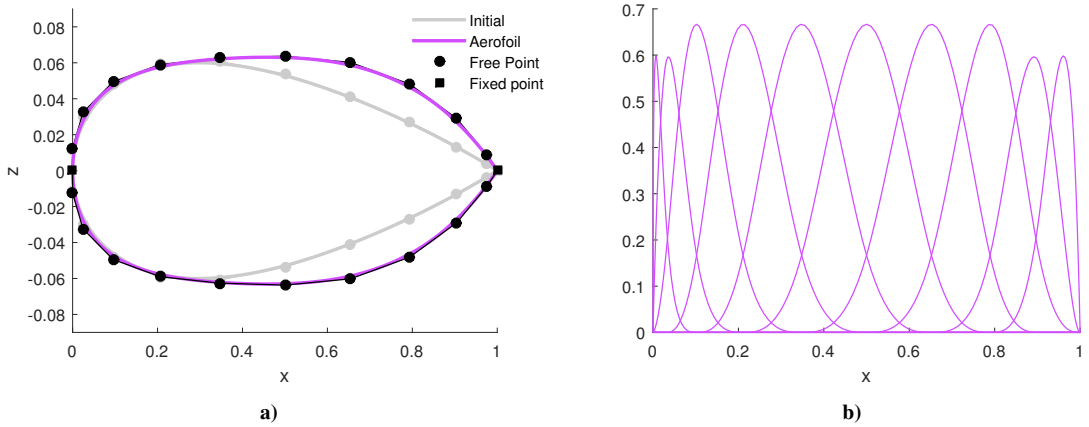

Figure 2. Example deformation of a cubic B-Spline configuration with 10 symmetric design variables where (a) shows the surface and control point positions and (b) shows the basis functions for the symmetric design variables that act on the upper surface of the aerofoil.

C. Class Function/ Shape Function Transformations (CST)

The CST method was developed by Kulfan[22, 23] primarily as a method of defining a wide range of aerofoils with relatively few design variables; however, the method can also be extended to other shapes such as square-like and circle-like objects. It is defined as

$$z_{upper} = C_{N2}^{N1}(x) \cdot S_{upper}(x) + x \cdot z_{te,upper}, \tag{13}$$

$$z_{lower} = C_{N2}^{N1}(x) \cdot S_{lower}(x) + x \cdot z_{te,lower}, \tag{14}$$

where the class function

$$C_{N2}^{N1}(x) = x^{N1} \cdot (1 - x)^{N2}, \tag{15}$$

and z_{te} defines the trailing edge thickness and $x \in [0, 1]$.

The values of $N1$ and $N2$ control the overall shape of the parameterisation, where a value of 1 creates a sharp edge, 0.5 a rounded edge and $0 < \epsilon \ll 1$ a flat box shape. The aerofoil class is therefore defined by $N1 = 0.5$ and $N2 = 1$, giving the round leading edge and sharp trailing edge required.

Kulfan[24] suggested defining $S(x)$ as the linear combination of Bernstein polynomials i.e.

$$S(x) = \sum_{i=0}^{n} a_i B_{i,n}(x) \tag{16}$$

where

$$B_{i,n}(x) = \binom{n}{i} x^i (1-x)^{n-i}, \tag{17}$$

a_i is the Bernstein coefficient and n is the degree of the polynomials. The class of Bernstein polynomials are a set of single sign C^n continuous functions defined on the interval $x \in [0, 1]$ and, in this region, are mathematically equivalent to the set of standard polynomials of the form $c_i x^i$.

Kulfan[25] later presented a leading edge modification (LEM) to the CST method, including an extra polynomial and coefficient, to improve the fidelity at the leading edge. This proposed adding an additional shape term such that

$$S(x) = \sum_{i=0}^{n} a_i B_{i,n}(x) + a_{n+1} x^{0.5} (1-x)^{n-0.5}. \tag{18}$$

For the tests performed in this paper the leading edge modification was always used. This is done because in other work from the authors[21] it was shown that the leading edge modification significantly helps with geometric accuracy of the parameterisation for aerofoils with high or low leading edge radius and previous work done on this case[20] suggests that a very large leading edge radius can be expected. The upper and lower surfaces will again be taken to be symmetric, reducing the number of design variables by half. The basis functions for a six design variable CST configuration are shown in 3a.

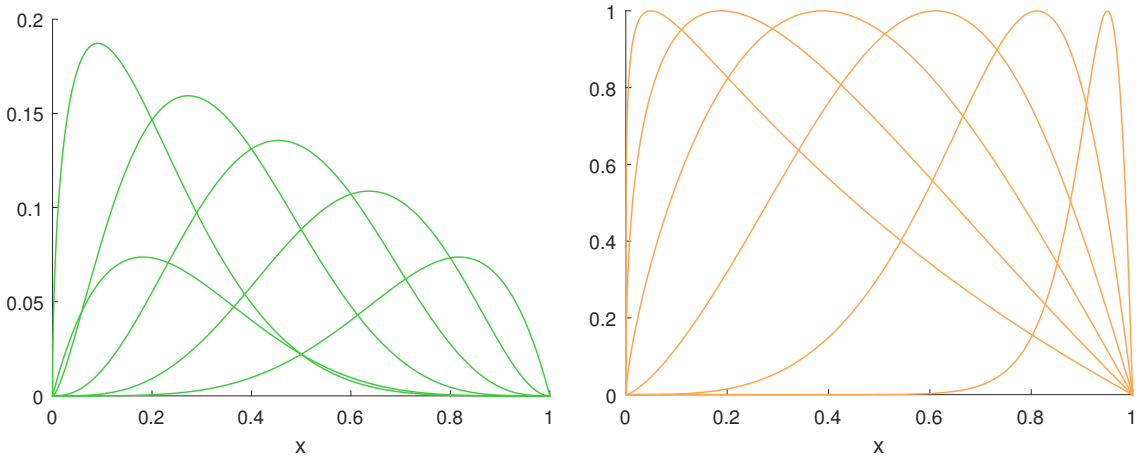

Figure 3. Basis functions for six design variable configurations of the CST method (a) and Hicks-Henne bump functions (b).

D. Hicks-Henne Bump Functions

Hicks-Henne bump functions use a base aerofoil definition plus a linear combination of a set of n basis functions defined between 0 and 1 to determine the final aerofoil shape. Each surface is defined by

$$\mathbf{z} = \mathbf{z}^{initial} + \sum_{i=0}^{n} a_i \phi_i(x) \tag{19}$$

for basis functions $\phi_i(x)$ and coefficients a_i for $i = 1, \ldots, n$.

The basis functions proposed by Hicks and Henne[26] were the sine functions

$$\phi_i(x) = \sin^{t_i}\left(\pi x^{\ln(0.5)/\ln(h_i)}\right), \tag{20}$$

where h_i is the location of the maxima of the basis function and t_i controls the width of the functions. Each bump function is therefore defined by three variables, each of which can be optimised or fixed. It is however preferable in many cases to fix both the position and width as this ensures the parameterisation is a linear function of the design variables; this is the approach used here. The bump positions will be defined as

$$h_i = \frac{1}{2}\left[1 - \cos\left(\frac{i\pi}{n+1}\right)\right], \quad i = 1, \ldots, n. \tag{21}$$

as used by Wu[27] and Masters[21, 28] and thickness parameters will be set to $t_i = 1$.

Due to the symmetry of the problem being investigated the upper and lower surface design variables were again taken to be equal.

E. Radial Basis Function Domain Element Method

The RBF domain element (RBF-DE) approach is a full domain deformation method like the Bèzier surface, so creates new aerofoil shapes based on the deformation of an initial aerofoil. The deformation method itself differs however, deforming by preserving the exact movement of a set of control points then creating a deformation field defined by radial basis function interpolation. The general theory of RBFs is outlined by Wendland[29] and Buhmann[30]; the formulation used here is presented extensively in Rendall and Allen[31] and its use as a parametrisation technique in Morris *et al.*[32].

The general solution for the deformation from the undeformed domain $\mathcal{A}(X)$ to the deformed domain $\mathcal{D}(X)$ is given by

$$\Gamma(X) = \sum_{i=1}^{n} \boldsymbol{\beta}_{DE_i} \phi(\|X - X_{DE_i}\|) + \mathbf{p}(X) \tag{22}$$

where DE_i indicates the ith domain element control point, X_{DE_i} its centre and $\boldsymbol{\beta}_{DE_i}$ its coefficient vector. $\mathbf{p}(\mathbf{x})$ is a linear polynomial used to ensure that translation and rotation are captured without added shape deformation.

The coefficients $\boldsymbol{\beta}_{DE_i}$ are found by requiring the exact recovery of the original function when the control points are in their original positions. The system is then completed by the additional requirement

$$\sum_{i=1}^{N} \boldsymbol{\beta}_{DE_i} \cdot \mathbf{q}(X) = 0 \tag{23}$$

where $\mathbf{q}(X)$ is a polynomial with order less than or equal to $\mathbf{p}(X)$.

When a discrete set of points in the original domain is to be transformed the problem can be formulated with matrix multiplication. Exact recovery of original points implies that

$$\mathbf{X}_{DE} = \mathbf{D}\boldsymbol{\Lambda} \tag{24}$$

where

$$\mathbf{X}_{DE} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ x_{DE_1} & z_{DE_1} \\ \vdots & \vdots \\ x_{DE_N} & z_{DE_N} \end{pmatrix}, \quad \boldsymbol{\Lambda} = \begin{pmatrix} \eta_0^x & \eta_0^z \\ \eta_x^x & \eta_x^z \\ \eta_z^x & \eta_z^z \\ \beta_{DE_1}^x & \beta_{DE_1}^z \\ \vdots & \vdots \\ \beta_{DE_N}^x & \beta_{DE_N}^z \end{pmatrix} \tag{25}$$

and

$$\mathbf{D} = \begin{pmatrix} 0 & 0 & 0 & 1 & 1 & \cdots & 1 \\ 0 & 0 & 0 & x_{DE_1} & x_{DE_2} & \cdots & x_{DE_N} \\ 0 & 0 & 0 & z_{DE_1} & z_{DE_2} & \cdots & z_{DE_N} \\ 1 & x_{DE_1} & z_{DE_1} & \phi_{DE_1DE_1} & \phi_{DE_1DE_2} & \cdots & \phi_{DE_1DE_N} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{DE_N} & z_{DE_N} & \phi_{DE_NDE_1} & \phi_{DE_NDE_2} & \cdots & \phi_{DE_NDE_N} \end{pmatrix} \tag{26}$$

with

$$\phi_{DE_iDE_j} = \begin{cases} \phi(\|X_{DE_i} - X_{DE_j}\|/S_R) & \text{for } \|X_{DE_i} - X_{DE_j}\| \leq S_R \\ 0 & \text{for } \|X_{DE_i} - X_{DE_j}\| > S_R \end{cases} \tag{27}$$

indicating the basis function on the distance between X_{DE_i} and X_{DE_j} and the subscript DE representing a domain element control point. S_R denotes the support radius which specifies the radius of the influence of each control point.

To locate the deformed points the following matrix must be formed where the subscript a indicates the original position of an aerofoil point:

$$\mathbf{A} = \begin{pmatrix} 1 & x_{a_1} & z_{a_1} & \phi_{a_1DE_1} & \phi_{a_1DE_2} & \cdots & \phi_{a_1DE_N} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{a_N} & z_{a_N} & \phi_{a_NDE_1} & \phi_{a_NDE_2} & \cdots & \phi_{a_NDE_N} \end{pmatrix}. \tag{28}$$

The matrix of the deformed aerofoil points, $\mathbf{X}^{deform}$, is then calculated as

$$\mathbf{X}^{deform} = \mathbf{A}\mathbf{\Lambda} \tag{29}$$

$$= \mathbf{A}\mathbf{D}^{-1}\mathbf{X}_{DE} \tag{30}$$

$$= \mathbf{H}\mathbf{X}_{DE}. \tag{31}$$

Note that as $\mathbf{H}$ is invariant of the current control point positions it only needs to be calculated once. It should be noted that this method can also be applied without the first three 'polynomial' rows or columns. The effect of this is that translations and rotations are not retained exactly but it ensures that deformations do not propagate past the support radius.

There are a few factors that affect the use of the RBF-DE method for reconstructing aerofoils; the support radius, the radial basis function, the initial aerofoil used, the number and initial position of the control points as well as the direction of their movement. For this study a support radius of 1 chord will be used throughout as well a radial basis function of Wendland's C2 function

$$\phi(x) = (1 - x)^4(4x + 1). \tag{32}$$

This leaves the positions of the initial control point positions to be defined. Contrary to the Bèzier surface method where the control points must be defined on a fixed uniform lattice, the initial RBF-DE control points can be placed anywhere. This flexibility gives the user great control over the influence and locality of the deformation though means that a comprehensive search for their best locations is challenging[33].

For this study, two different initial control point schemes were be considered. A set of 'off surface' control points defined on an ellipse around the aerofoil and a set of 'on surface' control points defined on the surface of the initial NACA0012, both shown in figure 4. For both configurations it should be noted that each configuration always contains the points included in the coarser levels with an additional set of bisecting points. Similarly to the other methods the design variables are chosen to be the symmetric pairs moving symmetrically in the z direction to create a local thickness change. Additional for the 'on surface' configuration the control points at the leading and trailing edges are held stationary.

F. SVD Method

The Singular Value Decomposition (SVD) method uses proper orthogonal decomposition to derive a set of ordered, orthogonal basis modes from a set of pre-determined training aerofoils. New aerofoil shapes can then be constructed as a linear combination of these modes where the fidelity of the construction is determined by the number of modes used. This technique was first employed by Toal *et al.*[34] then by Ghoman *et al.*[35] and Poole *et al.*[36]. Ghoman *et al.*[35] used a series of supercritical aerofoils to derive the modes and showed that other supercritical aerofoils could efficiently be reconstructed. Poole at al.[36] then extended this to show that a broad range of aerofoils could be represented given a wide choice of training aerofoils.

When constructing the aerofoil shape modes it is crucial that the training library is normalised such that the aerofoil shapes are defined equivalently and discretized by the same number of points. In this work they are transformed to have a sharp trailing edge and discretized such that all the aerofoils have an equal distribution of points along the x-axis. This second condition means that, in this instance, the modes only need to be constructed in the z direction.

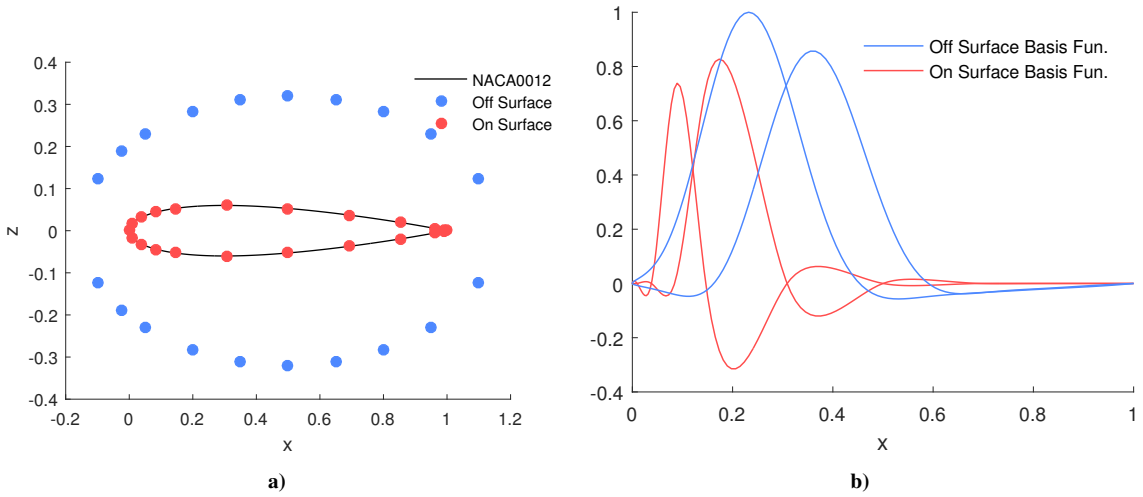

Figure 4. Comparison of 'on surface' and 'off surface' RBF-DE configurations for 10 symmetric design variables. (a) Shows the initial control point positions used and (b) shows a couple of the the basis function associated with this configuration (acting on the upper surface).

To formulate these modes the z coordinates of the training aerofoils are first used to form the rows of the matrix

$$\mathbf{T} = \begin{bmatrix} z_1^1 & z_2^1 & \cdots & z_N^1 \\ z_1^2 & z_2^2 & \cdots & z_N^2 \\ \vdots & \vdots & \ddots & \vdots \\ z_1^M & z_2^M & \cdots & z_N^M \end{bmatrix} \tag{33}$$

for M training aerofoils each of length N.

This matrix is then decomposed into the singular value decomposition

$$\mathbf{T} = \mathbf{U} \cdot \mathbf{\Sigma} \cdot \mathbf{V} \tag{34}$$

where columns of $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_{min(N,M)}]$ represent the ordered, orthogonal aerofoil modes and the diagonal values of Σ represent the energy of each mode. The energy is a measure of the importance of the modes within the training library and equivalently can be considered as a 'typical value' for its use.

New aerofoils are then constructed as a linear combination of these modes such that

$$\mathbf{z} = \sum_i a_i s_i \mathbf{v}_i \tag{35}$$

for some scaling $s_i$and where a_i represents the design variables. Two possible scalings are considered in this work, $s_i = 1$ and $s_i = \Sigma_{ii}$.

For this work the training library consisted of all the symmetric aerofoils from the UIUC aerofoil library[b] (smoothed and normalised as described in Masters[28]) and the symmetric NACA 4 series aerofoils with max thickness coefficients from 6 to 24. This made up a training library of 122 aerofoils. The modes produced are shown in figure 5 with the associated energies shown in figure 6.

[b] http://aerospace.illinois.edu/m-selig/ads/coord_database.html

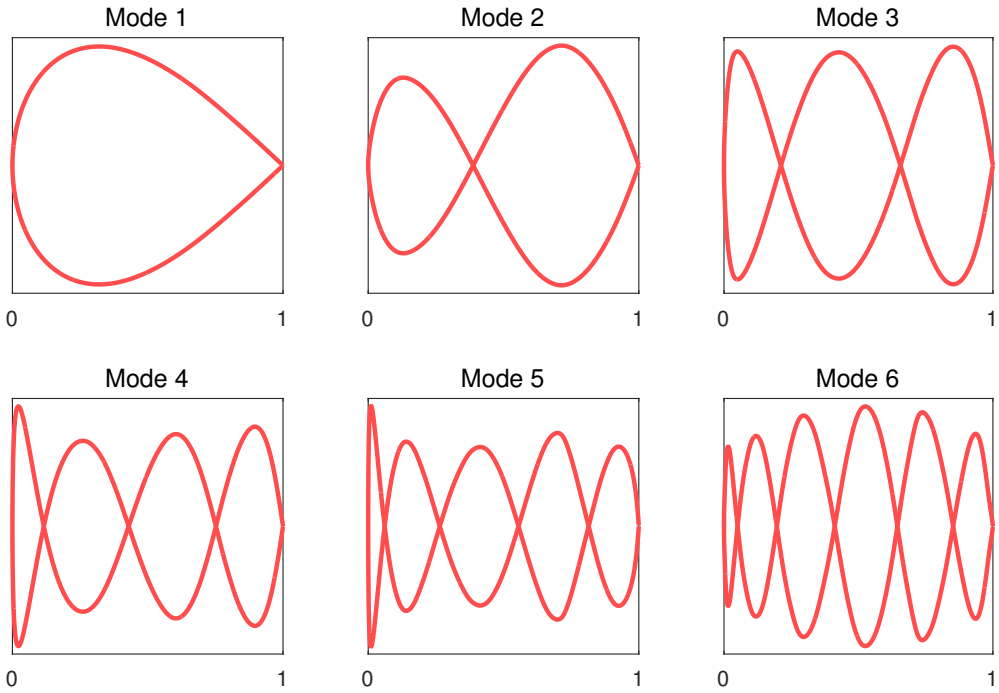

Figure 5. The first 6 SVD modes created with the training library of symmetric aerofoils.

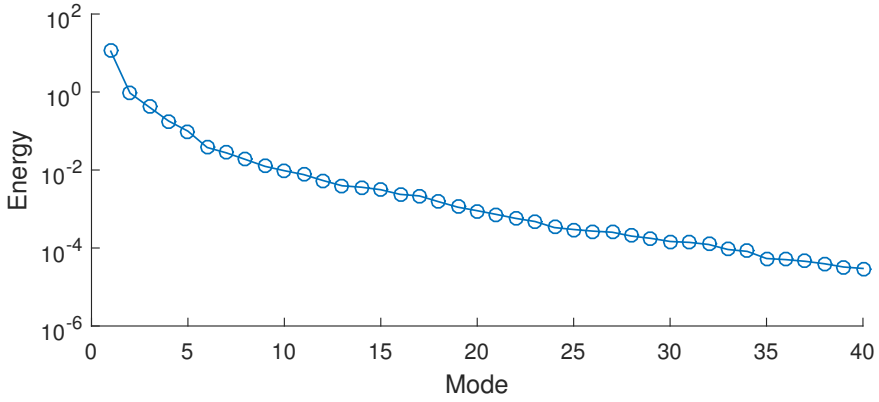

Figure 6. Energy associated with the first 40 SVD modes.

IV. Optimisation Framework

In this work the open-source, unstructured CFD code SU^2[18] was used for all the optimisations, and surface sensitivities were calculated through the adjoint equations. Each flow solution was converged down to a maximum change in the drag coefficient of less than 10^{-7} and convergence acceleration was achieved through the use of multigrid. Due to the symmetry of the problem only half of the aerofoil was solved for each simulation with a symmetry condition applied along the line $z = 0$.

The computational meshes were created as structured O-meshes with a distance to the far-field of 50 chord lengths and an equal ratio of cells around the aerofoil to cells to the far-field. The original problem specification[19] recommends that the initial mesh is grid independent to within 0.1 drag counts. A full mesh convergence study for this grid generator is presented in Poole *et al.*[14] (though with the alternative solver used in section VI), which identified that this condition was sufficiently met for the 257×257 mesh. These tests have been repeated with SU^2 which confirmed these results. For this reason all of the optimisation procedures were performed with the 257×257 mesh shown in figure 7a. In addition the final results were also run on the finer 513×513 mesh to confirm the results. These two mesh sizes will be referred to here on after as the 'optimisation mesh' and 'fine mesh' and are shown in figure 7.

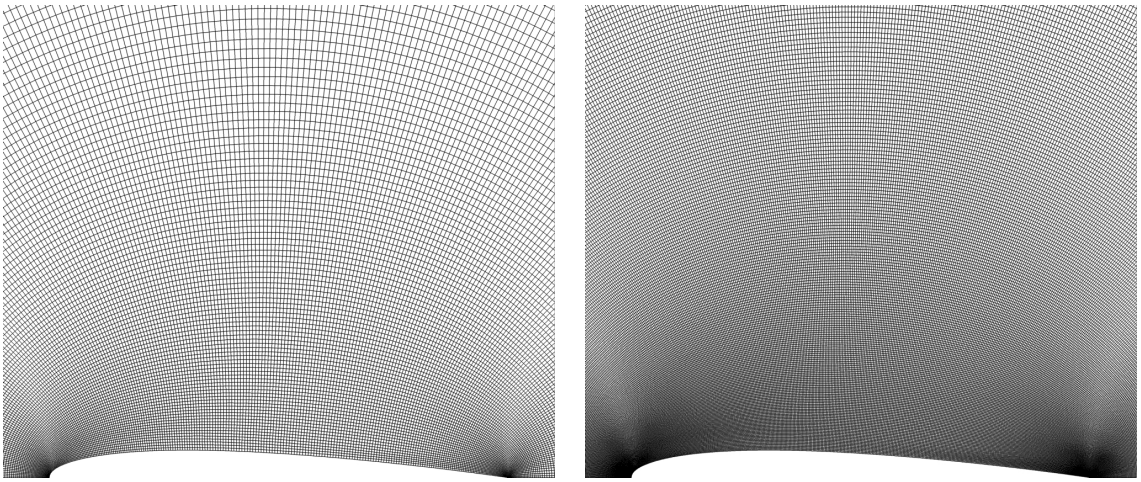

a) 257×257 **Optimisation Mesh** **b)** 513×513 **Fine Mesh**

Figure 7. Mesh resolutions used for optimisation, (a), and final tests, (b).

The optimisation mesh is used for each optimisation procedure throughout with the surface and volume mesh deformed using RBFs for each flow solve. This is done by applying the same method as for the RBF domain element parameterisation with 'control points' placed on each surface point though without the use of the 'polynomial' terms. This is applied with a support radius of 10 chord lengths and Wendland's C4 RBF,

$$\phi(x) = (1 - x)^6(35x^2 + 18x + 1), \tag{36}$$

to ensure the surface deformations are dissipated smoothly across the volume to maintain mesh quality.

The design variable gradients are calculated using the adjoint method which allows all the gradients to be calculated for a computational cost in the order of one flow solve. This is done by solving the adjoint equations[37] to calculate the sensitivity of drag coefficient with respect to the unit normal at each surface mesh point, i.e.

$$\frac{\partial C_D}{\partial \hat{\boldsymbol{x}}_i} \tag{37}$$

for unit surface normal $\hat{\boldsymbol{x}}_i = [\hat{x}_i, \hat{z}_i]$. Then as the surface perturbations in this work will be applied in the z direction only, the required surface sensitivities are given by

$$\frac{\partial C_D}{\partial z_i} = \hat{z}_i \frac{\partial C_D}{\partial \hat{\boldsymbol{x}}_i}. \tag{38}$$

The required gradients of C_D with respect to the design variables, a_i, are then calculated by multiplying this with the

Jacobian of the surface geometry with respect to the chosen design variables,

$$\underbrace{\begin{bmatrix} \frac{\partial C_D}{\partial a_1} \\ \frac{\partial C_D}{\partial a_2} \\ \vdots \\ \frac{\partial C_D}{\partial a_n} \end{bmatrix}}_{\text{Gradients}} = \underbrace{\begin{bmatrix} \frac{\partial z_1}{\partial a_1} & \cdots & \frac{\partial z_m}{\partial a_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial z_1}{\partial a_n} & \cdots & \frac{\partial z_m}{\partial a_n} \end{bmatrix}}_{\text{Geometric Sensitivities}} \underbrace{\begin{bmatrix} \frac{\partial C_D}{\partial z_1} \\ \frac{\partial C_D}{\partial z_2} \\ \vdots \\ \frac{\partial C_D}{\partial z_m} \end{bmatrix}}_{\text{Surface Sensitivities}} . \tag{39}$$

The optimisations were then performed using the multi-purpose large-scale optimiser SNOPT[17]. This is a gradient-based sequential-quadratic programming (SQP) method that employs a reduced-Hessian BFGS search direction and a cubic line-search technique. The surface profile restriction was applied as 'linear' constraint and was enforced at every fifth percentile for $x/c \leq 20\%$ and $x/c \geq 40\%$ and every second percentile for $30\% < x/c < 40\%$ as well as at 99% chord. This ensures strict compliance in the maximum thickness region.

Convergence of the optimisation was typically determined through SNOPT which terminated from one of two conditions. Either the KKT optimality condition[17] satisfied the tolerance of 1×10^{-6} or the optimiser was unable to improve the objective function after multiple attempts. In some cases the optimisations were also stopped manually if no significant improvement was made after a large number of iterations.

V. Optimisation Results

A. Initial Optimisation Results

For an initial investigation, design variables sweeps from 4 to 40 were conducted for CST, Bèzier Surface, Cubic B-Spline, Hicks-Henne, RBF-DE on surface and SVD (unit scaled) methods. Figure 8 shows the optimisation results for these cases. It can be seen that there is a clear distinction between the performance of the methods for larger numbers of design variables. The Bèzier surface, CST and Hicks-Henne methods seem to give reasonably consistent results for greater than 20 design variables, whereas the B-Spline, RBF-DE and SVD methods get consistently worse over this period. For all these methods it is clear that better solutions exist than those found by the optimiser; this suggests that the optimiser has been unable to fully exploit the available design space. Figure 9 then shows the optimiser convergence history for the six methods for 36 design variables. At this design point it can be seen that the convergence path for the RBF-DE, B-spline and SVD methods is significantly worse than for the Bèzier surface, CST and Hicks-Henne methods.

To investigate this, the parameterised surface sensitivity was calculated for each method at the optimisation starting point with 40 design variables. This was calculated by multiplying the geometric sensitivities by the design variable gradients. For an unconstrained optimisation this is equivalent to the surface deformation direction, however due to the active profile constraints in this problem this deformation could provide an infeasible shape so merely represents an 'ideal' step direction. The parameterised sensitivities for all of the methods have been plotted in figure 10. This figure shows a clear distinction between the Bèzier surface, CST and Hicks Henne methods, and the B-Spline, RBF-DE and SVD methods. The B-Spline, RBF-DE and SVD methods follow the exact adjoint sensitivity much closer than the other methods and as a consequence have significantly sharper, higher frequency curves. For the RBF-DE and B-Spline methods this is because the support of the basis functions reduces as the fidelity increases whereas for the SVD method it is as result of the increasing frequency of the modes. The impact of this is that the resulting surface perturbation may reduce the smoothness of the aerofoil. It has previously been identified that ensuring the smoothness in the aerofoil optimisation sequence is crucial to the success of aerofoil optimisation[38, 39] which suggests that this capacity to create unsmooth aerofoils may be negatively influencing the optimisation process.

To ensure smoother sensitivities were created, two different approaches have been taken. For the RBF-DE and B-spline methods, configurations have been chosen that create larger smoother basis functions. For the RBF-DE method this is done by moving the initial control points positions away from the aerofoil surface and for the B-Spline method a larger polynomial order, equal to half the number of control points, is used. For the SVD method a different approach was taken due to the form of the basis functions; in this case the modes were scaled by the modal energy such that the later, higher frequency modes had a lower maximum value.

Figure 13 shows the impact this has on the parameterised surface sensitivities; it can be seen that for the new configurations they are significantly smoother. The impact of this on the optimisation results and convergence history at 36 design variables is then shown in figures 11 and 12. Both of these figures show a significant improvement in the results. This suggests that the smoothness of the surface sensitivities has a large impact on the robustness and rate of convergence of this optimisation procedure.

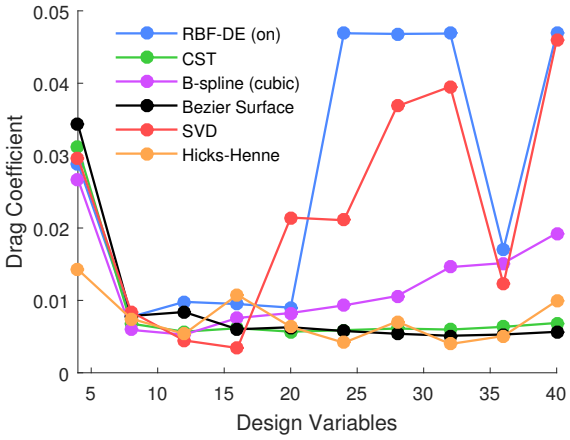

Figure 8. Optimisation results for the initial set of parameterisation configurations.

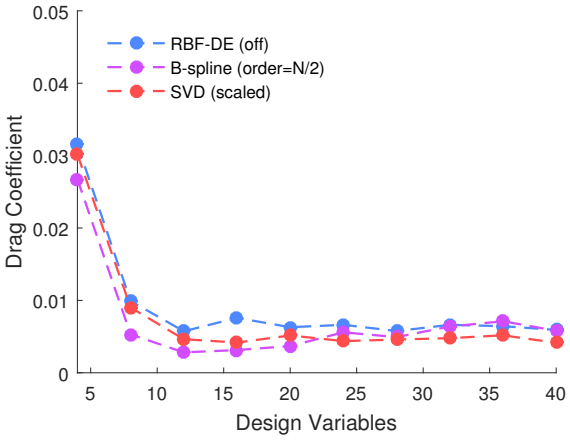

Figure 11. Optimisation results for the improved set of parameterisation configurations.

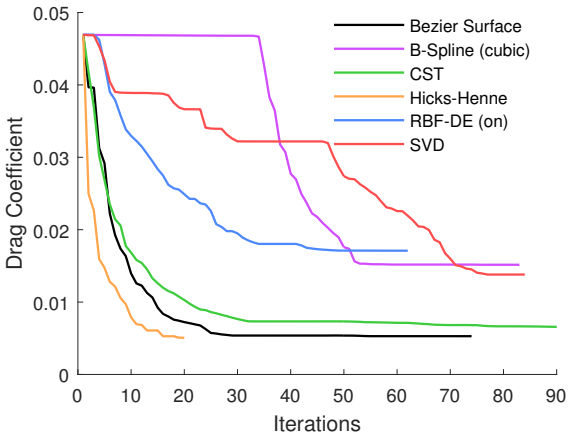

Figure 9. Convergence history for the initial set of parameterisation configurations with 36 design variables.

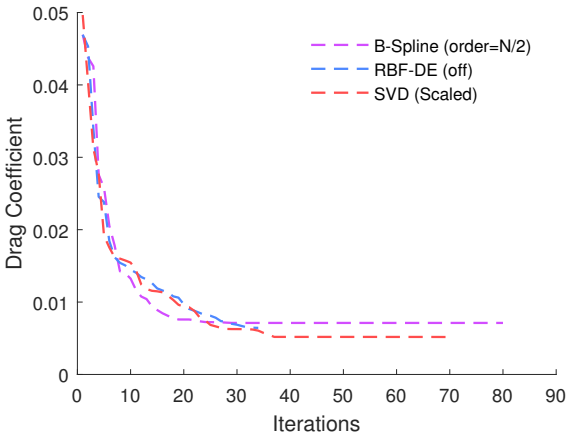

Figure 12. Convergence history for the improved set of parameterisation configurations with 36 design variables.

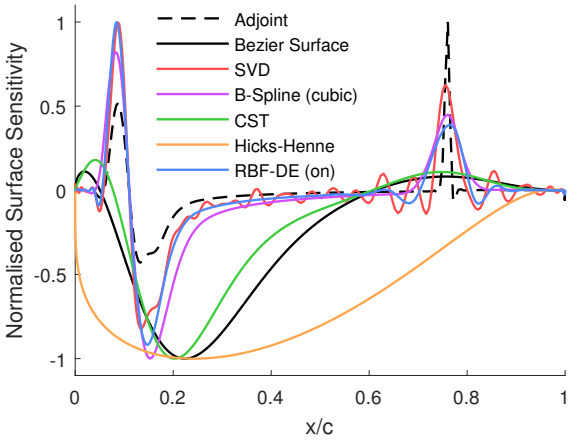

Figure 10. Surface sensitivities on NACA0012 for the initial set of parameterisation configurations for 40 design variables.

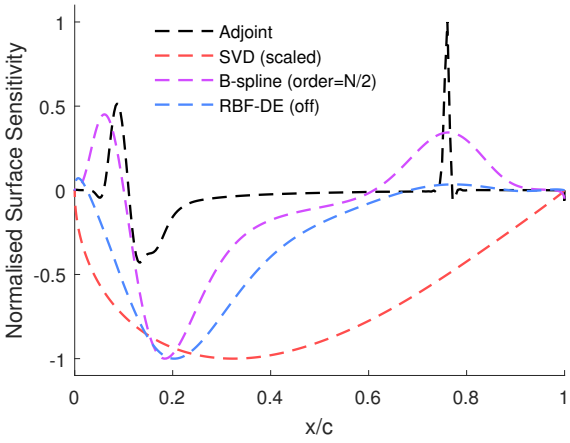

Figure 13. Surface sensitivities on NACA0012 for the improved set of parameterisation configurations for 40 design variables.

B. Improved Optimisation Results

Figure 14 shows the final optimisation results achieved, with the inclusion of the improved methods for the B-Spline, RBF-DE and SVD methods. These show that the B-Spline method provides the best result with 12 design variables achieving a drag count of 28.2; this equates to a 94% reduction in drag. The other methods, bar Hicks-Henne, seem to produce very similar, consistent, results across the full design variable spectrum investigated, with the SVD method generally attaining a slightly better result than the rest. The Hicks-Henne method on the other hand shows the most variation in results. This may be due to the formation of the basis functions used, as an increase in the number of design variables does not necessarily strictly increase the design space[21]. This can mean that lower design variable solutions are not included in higher design variable space.

Each of the optimum solutions was then re-meshed and run at the 'fine', 513×513, mesh resolution to confirm the results and reduce grid dependence. These results can be seen in figure 15. It was expected that all of the results would either decrease or stay the same based on previous work[5, 14]. It can however be seen that for some of the results (indicated by a ▲), a large increase in drag is observed. Figure 16 shows the pressure distributions for two of these cases for the two different mesh resolutions tested. It can be seen that they have produced completely different shock structures near the trailing edge. This is evidence of non-unique solutions for these aerofoils and is similar to the behaviour identified by Meheut *et al.*[20]. Further analysis of this hysteretic behaviour is presented in the next section. Ignoring these results it can be seen that figures 14 and 15 show very similar trends with a general reduction in drag shown for the finer mesh.

The best results obtained by each method at the fine mesh resolution are then compared in table 1, from this it can be seen that the B-Spline method gives the best overall result with 16 design variables attaining a drag count of 25.1 counts, which equates to a 95% reduction in drag. The aerofoil shapes and pressure distributions associated with these best results are then shown in figures 17 and 18. From the aerofoil shapes it can be seen that there are two distinct surface results. The two best results, for the B-spline and Hicks-Henne methods, can be seen to have a thicker trailing edge region and an increase in thickness around 50% chord. The other results however maintain the maximum thickness of the original NACA0012 and have slightly thinner shape around the trailing edge; these results produced between 5 and 35 counts more drag. The C_P plots (figure 18) show a similar pattern with the B-Spline and Hicks-Henne giving a different trailing edge shock structure to the other methods. These, lower drag, results display a single normal shock at around 97% chord whereas the other cases appear to form a partially reflected oblique shock at 92% chord followed by a normal shock at 94%.

Method	**# of DVs**	C_D **(counts) (257x257 mesh)**	C_D **(counts) (513x513 mesh)**
NACA0012	-	*469.3*	*469.4*
CST	20	56.8	56.0
Bèzier Surface	32	51.2	50.0
B-Spline (Order = N/2)	16	31.4	25.1
Hicks-Henne	32	39.8	31.9
RBF-DE (off)	28	57.7	56.4
SVD (scaled)	24	38.6	37.2

Table 1. Table showing the best results attained by each method on the 257×257 optimisation mesh and 513×513 refined mesh.

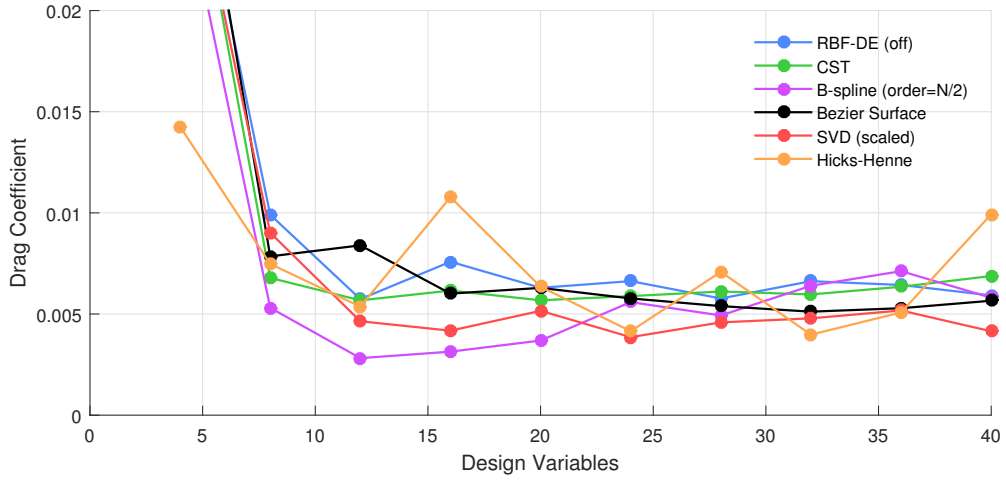

Figure 14. Final optimisation results for each parameterisation method on the 257×257 optimisation mesh.

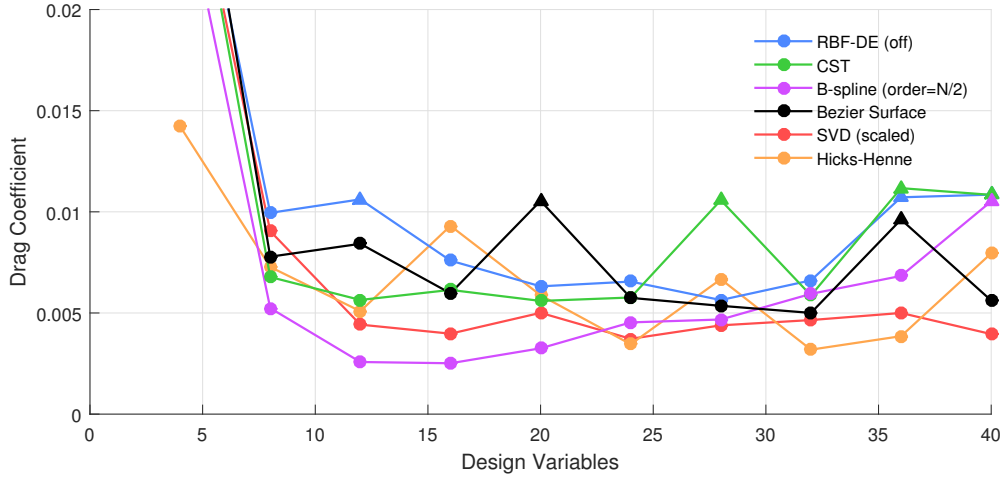

Figure 15. Final optimisation results for each parameterisation method run on the 513×513 fine mesh. ▲ indicates a solution that may have converged to the upper branch of a hysteresis loop.

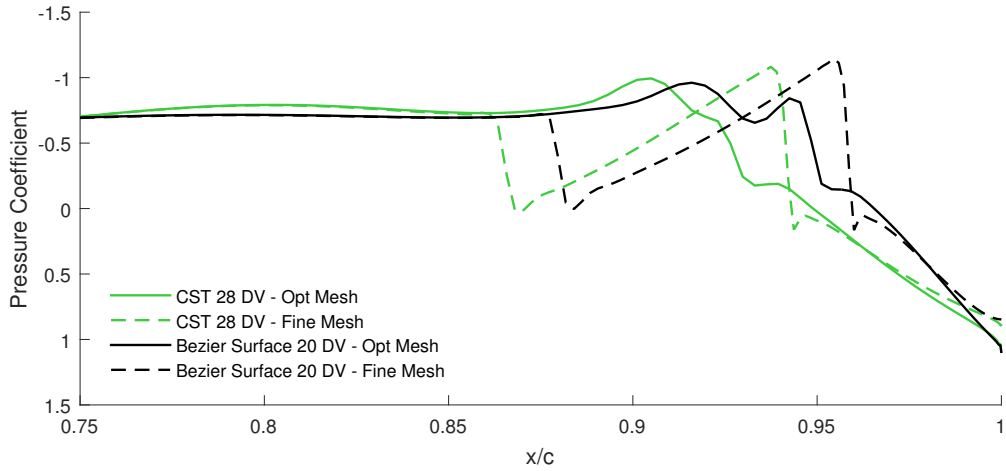

Figure 16. Comparison of trailing edge pressure distributions for two aerofoils at the two mesh resolutions.

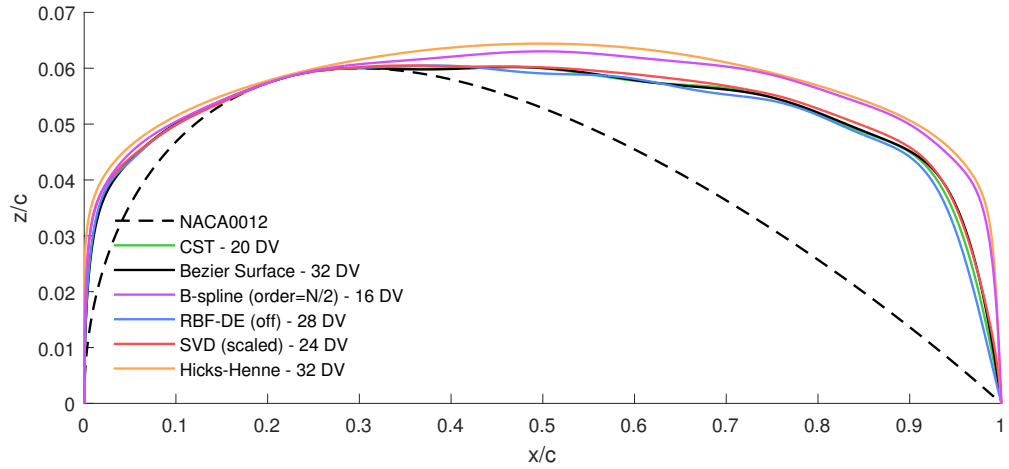

Figure 17. Comparison of the aerofoil shapes for the optimum result achieved with each method.

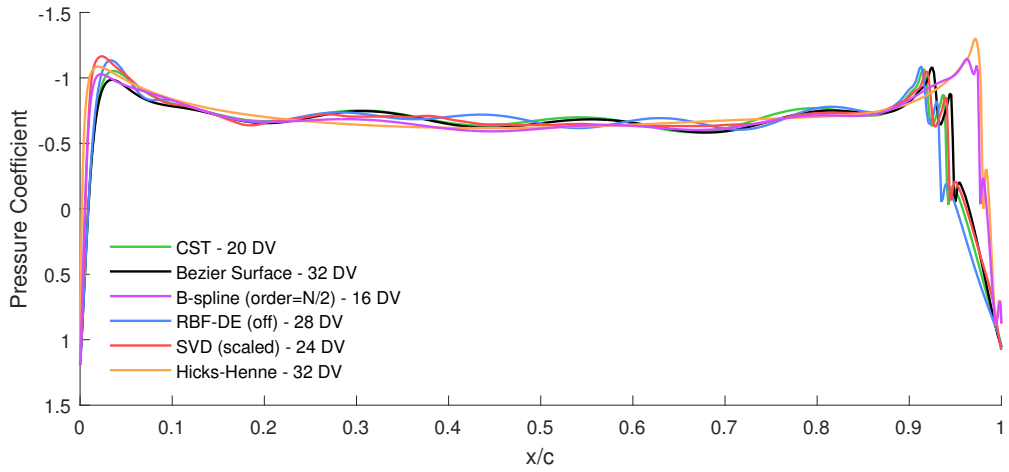

Figure 18. Comparison of the aerofoil shapes for the optimum result achieved with each method

VI. Validation and Hysteresis Study

It is clear from figure 17 that the six different parameterisation methods tested produce two distinct optimised shapes. The CST, Bèzier Surface, RBF-DE and SVD methods produce a shape with a lower thickness over the aft half of the aerofoil chord with a shallower boat-tail angle, whilst the B-spline, Hicks-Henne and SVD methods produce a thicker aerofoil with a higher boat-tail angle. The thicker of the two distinct shapes produced drag values, according to SU^2, of around 25-30 counts on the fine mesh. The other, slightly thinner shape, produces drag values between 35 and 55 counts on the fine mesh.

A. Mesh Convergence Study

To validate the results from SU^2, a mesh convergence study was performed using another solver on an example of each of the two distinct designs. The best results for the B-spline method (order=N/2, 16 DV) and Bèzier surface method (32 DV) were the two aerofoils tested. The flow-solver used was the structured multiblock finite-volume, upwind code of Allen[40], which uses the flux vector splitting of van Leer[41]. Convergence acceleration is achieved through multigrid [42]. The meshes used were the 'fine' 513×513 meshes used in SU^2, which were then mirrored to produce full, 1025×513, meshes. Each of these was then coarsened twice in each direction to produce three distinct mesh levels for the mesh convergence study: 1025×513, 513×257 and 257×129. It should be noted that to avoid lifting solutions (a hysteresis study is presented below to explain this more fully) a symmetric wake was enforced. The results of the mesh convergence studies are shown in table 2. The mesh convergence study for the NACA0012 is also shown. It can

be seen that the results from the mesh convergence study on the structured multi-block solver correlate well with those produced on SU^2.

	C_D (counts)		
Mesh Density	**NACA0012**	**B-spline**	**Bèzier Surface**
257×129	487.5	38.5	60.8
513×257	471.7	27.8	51.2
1025×513	469.6	25.0	49.8

Table 2. Table of mesh convergence studies on structured multi-block code.

B. Hysteresis Study

As noted above, to avoid lifting solutions, which were otherwise very common, a symmetric wake was enforced in the mesh convergence study above. The presence of lifting solutions means that multiple solutions are possible for these types of aerofoils at this design point. This is a result that has also been found by other researchers in the ADODG [5, 7, 20], and has also been subject to research by Jameson *et al.* [43, 44, 45].

To investigate hysteresis in the two aerofoil sections highlighted (those designed using B-spline and Bèzier surface parameterisations), firstly a large sweep in Mach number up and down is performed. The sweep is performed on both the 257×129 mesh and the 513×257 mesh, and the results are shown in figure 19. As has previously been found by other researchers, these types of aerofoils exhibit multiple solutions at the optimization point.

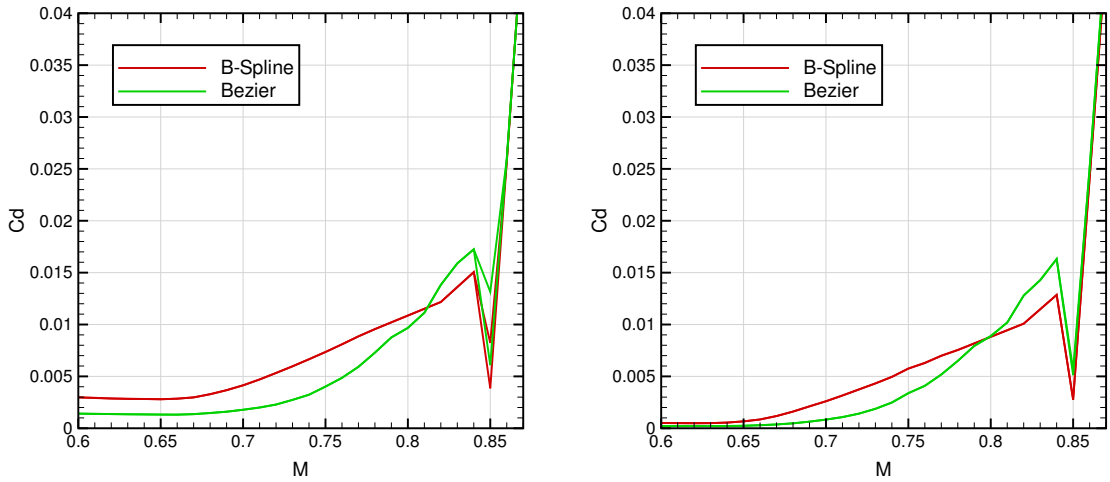

a) 257×129 **b) 513×257**

Figure 19. Large Mach sweeps

To further investigate this point, a smaller Mach sweep was run on a smaller deviation around $M = 0.85$ on the same meshes. The results for these are shown in figure 20. It is interesting to note that the hysteresis occurs at a very small window of freestream Mach number and is highly mesh dependent. For example, for the B-spline parameterised aerofoil on the 513×257 mesh, there is no hysteresis at the design point, and instead this occurs at a lower Mach number. The Mach flow contours are shown in figures 21 and 22 for both aerofoils on the 257×129 mesh (both the upper and lower branches are shown on the same plot). The upper branch of both of the aerofoils has a double shock flow structure with drag values approximately double the lower branch, which has a single shock. This is in agreement with the original problem set by Vassberg *et al.*[2], where it was believed that this is the lowest Mach number at which a shock free optimized solution is not possible.

VII. Conclusions

In this work the NACA0012 symmetric, inviscid drag optimisation benchmark case has been run with six different parameterisation methods (B-splines, Bèzier surfaces, CST, Hicks-Henne, RBF-DE and SVD) with between 4 and 40 design variables. Firstly it was found that the smoothness had a significant impact on the robustness and rate of convergence of the optimisations, with the capacity for some of the parameterisation methods to create un-smooth

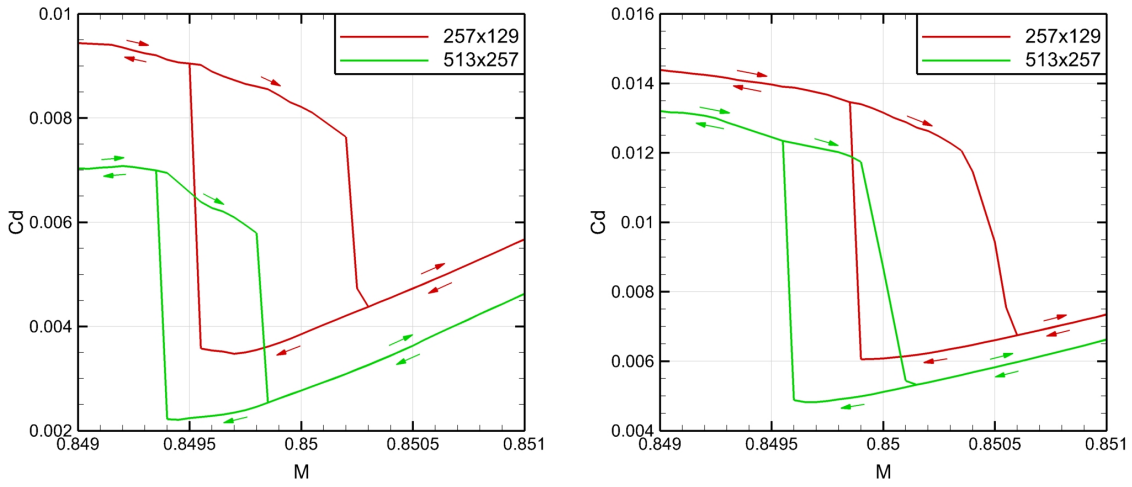

a) B-spline **b) Bèzier Surface**

Figure 20. Hysteresis loops produced for small Mach sweeps. Arrows indicated sweep direction.

aerofoils negatively impacting their results for higher design variables. This lead to improvements for the the B-spline and RBF-DE methods, by increasing the support of the basis functions and, for the SVD method, by reducing the influence of the higher frequency modes.

A set of final optimisation results were then presented at two mesh resolutions; a 257×257 mesh used for the optimisations themselves and a fine 513×513 mesh. They showed that the best result was achieved by the B-spline method with 16 design variables, producing a result with 25.1 drag counts on the 513×513 mesh. Two distinct surface solutions were found with the B-spline and Hicks-Henne methods producing a thicker, larger boat-tail angle compared to those produced by the Bèzier surface, CST, RBF-DE and SVD methods. The thicker shapes correlate with lower drag results and different shock structures.

A further investigation of the hysterestic behaviour encountered was then presented for the best results achieved with the B-spline and Bèzier surface methods. This showed that for both aerofoils hysteresis loops could be formed by sweeping up and down in Mach number within a range of ± 0.001 of the $M = 0.85$ design point.

VIII. Acknowledgements

This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol - http://www.bris.ac.uk/acrc/. The authors also wish to acknowledge the financial support provided by Innovate UK: the work reported herein has been undertaken in GHandI (TSB 101372), a UK Centre for Aerodynamics project.

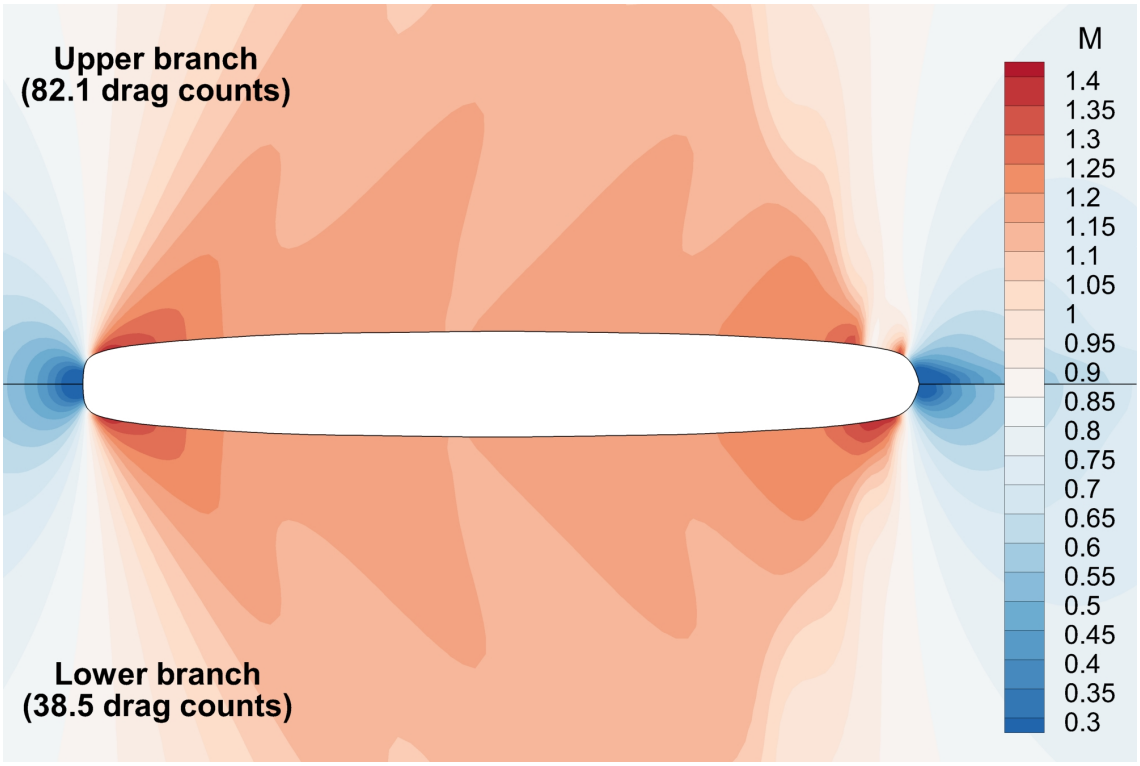

a) B-spline

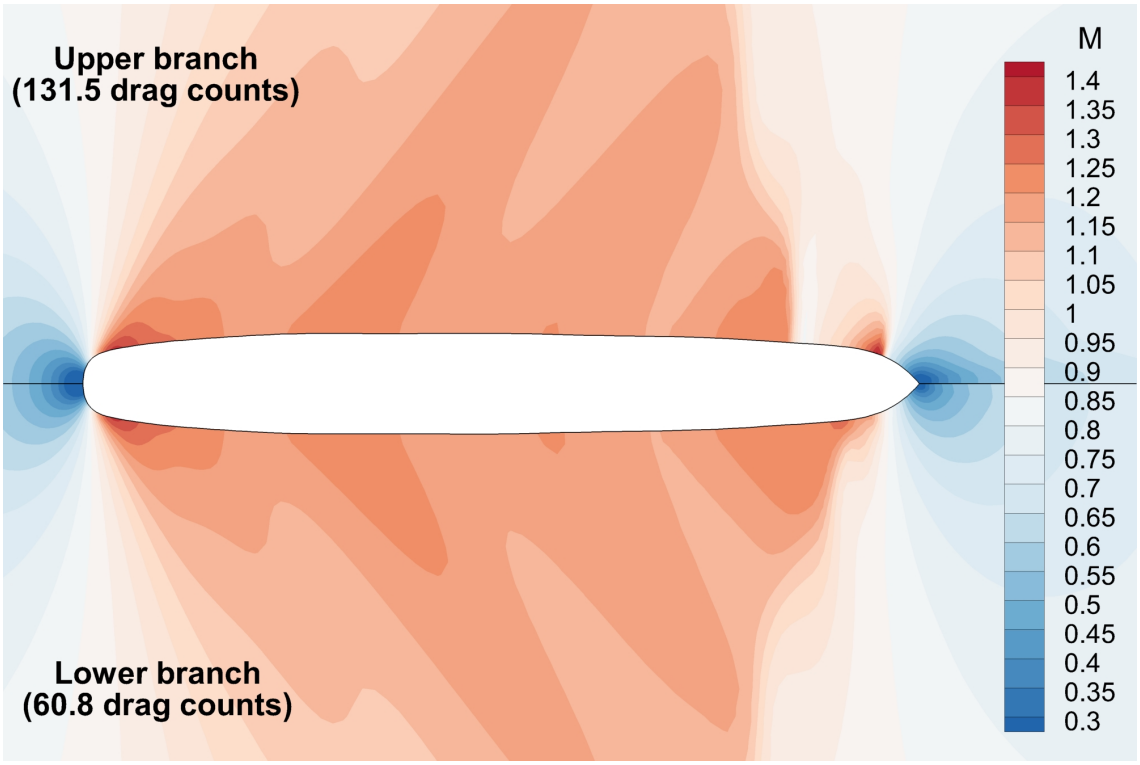

b) Bèzier Surface

Figure 21. Flowfield plots at $M = 0.85$ at upper and lower flow branches on 257×129 mesh.

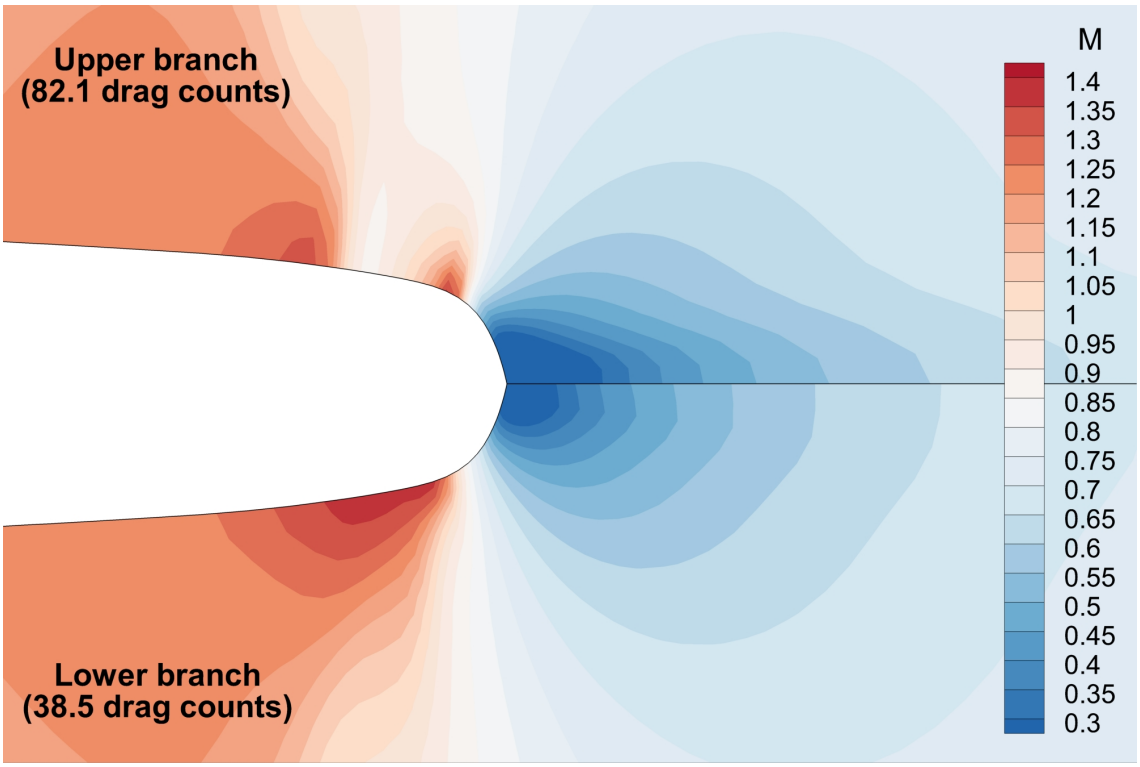

a) B-spline

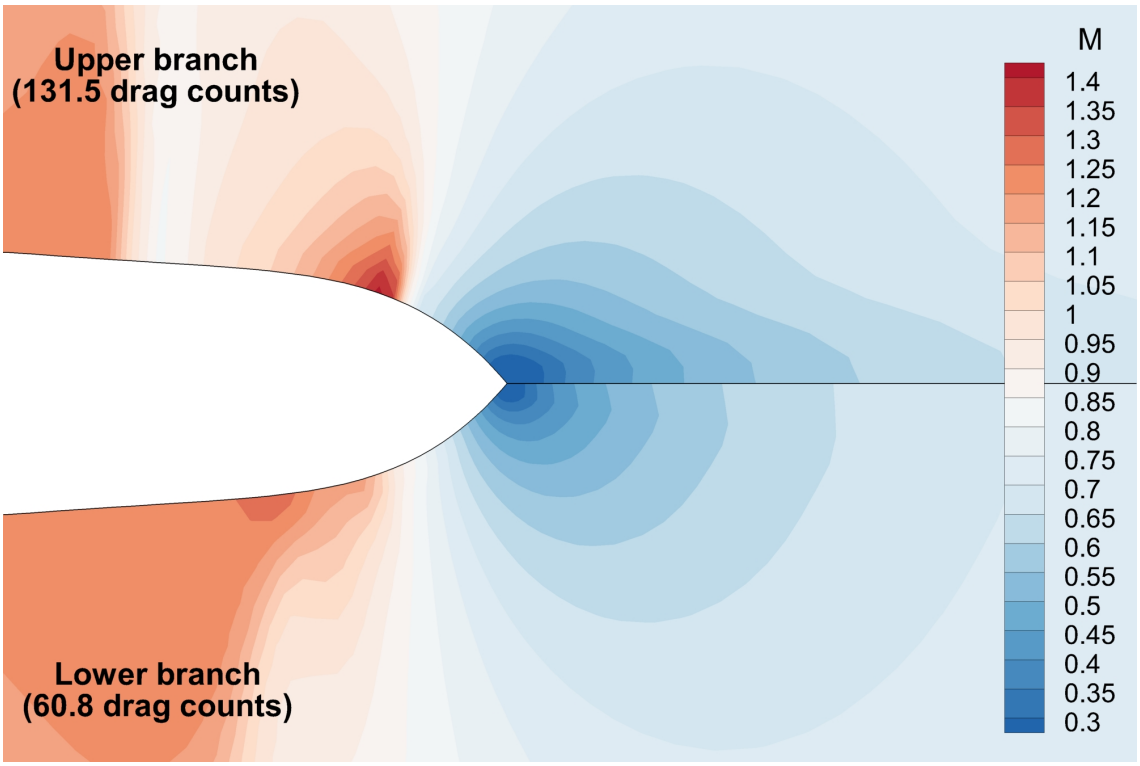

b) Bèzier Surface

Figure 22. Trailing edge flowfield plots at $M = 0.85$ at upper and lower flow branches on 257×129 mesh.

References

[1] Zhang, M., Rizzi, A. W., and Nangia, R. K., "Transonic Airfoils and Wings Design Using Inverse and Direct Methods," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[2] 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*, Jun 2011.

[3] Bisson, F. and Nadarajah, S., "Adjoint-Based Aerodynamic Optimization of Benchmark Problems," *52nd Aerospace Sciences Meeting*, Jan 2014.

[4] 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 Aerospace Sciences Meeting*, Jan 2014, 10.2514/6.2014-0568.

[5] Nadarajah, S., "Adjoint-Based Aerodynamic Optimization of Benchmark Problems," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[6] Telidetzki, K., Osusky, L., and Zingg, D. W., "Application of Jetstream to a Suite of Aerodynamic Shape Optimization Problems," *52nd Aerospace Sciences Meeting*, Jan 2014.

[7] Lee, C., Koo, D., Telidetzki, K., Buckley, H., Gagnon, H., and Zingg, D. W., "Aerodynamic Shape Optimization of Benchmark Problems Using Jetstream," *53rd AIAA Aerospace Sciences Meeting*, 2015.

[8] Gariepy, M., Trepanier, J.-Y., Malouin, B., and Tribes, C., "Direct Search Airfoil Optimization Using Far-Field Drag Decomposition Results," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[9] Fusi, F., Quaranta, G., Guardone, A., and Congedo, P. M., "Drag minimization of an isolated airfoil in transonic inviscid flow by means of genetic algorithms," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[10] Leifsson, L. T., Koziel, S., Tesfahunegn, Y. A., Hosder, S., and Gramanzini, J.-R., "Application of Physics-Based Surrogate Models to Benchmark Aerodynamic Shape Optimization Problems," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[11] Amoignon, O., Hradil, J., and Navratil, J., "Study of parameterizations in the project CEDESA," *52nd Aerospace Sciences Meeting*, Jan 2014.

[12] Leifsson, L. T., Koziel, S., and Hosder, S., "Aerodynamic Design Optimization: Physics-based Surrogate Approaches for Airfoil and Wing Design," *52nd Aerospace Sciences Meeting*, Jan 2014.

[13] Anderson, G. R., Nemec, M., and Aftosmis, M. J., "Aerodynamic Shape Optimization Benchmarks with Error Control and Automatic Parameterization," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[14] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Control Point-Based Aerodynamic Shape Optimization Applied to AIAA ADODG Test Cases," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[15] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Application of control point-based aerodynamic shape optimization to two-dimensional drag minimization," *52nd Aerospace Sciences Meeting*, 2014.

[16] LeDoux, S. T., Vassberg, J. C., Young, D. P., Fugal, S., Kamenetskiy, D., Huffman, W. P., Melvin, R. G., and Smith, M. F., "Study Based on the AIAA Aerodynamic Design Optimization Discussion Group Test Cases," *AIAA Journal*, Feb 2015, pp. 1–26.

[17] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM journal on optimization*, Vol. 12, No. 4, 2002, pp. 979–1006.

[18] Palacios, F., Alonso, J., Duraisamy, K., Colonno, M., Hicken, J., Aranake, A., Campos, A., Copeland, S., Economon, T., Lonkar, A., Lukaczyk, T., and Taylor, T., "Stanford University Unstructured (SU^2): An open-source integrated computational environment for multi-physics simulation and design," *Aerospace Sciences Meetings*, American Institute of Aeronautics and Astronautics, Jan. 2013.

[19] Nadarajah, S., "Aerodynamic Design Optimization: Drag Minimization of the NACA 0012 in Transonic Inviscid Flow," Retreived from https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG/Test%20Cases/ADODG%20Case%201%20and%202%20NACA0012%20and%20RAE%202822.pdf.

[20] Meheut, M., Destarac, D., Carrier, G., Anderson, G., Nadarajah, S., Poole, D., Vassberg, J., and Zingg, D. W., "Gradient-Based Single and Multi-points Aerodynamic Optimizations with the elsA Software," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[21] Masters, D. A., Taylor, N. J., Rendall, T. C. S., Allen, C. B., and Poole, D. J., "A Geometric Comparison of Aerofoil Shape Parameterisation Methods," *54rd AIAA Aerospace Sciences Meeting*, Jan 2016.

[22] Kulfan, B. M., "Universal parametric geometry representation method," *Journal of Aircraft*, Vol. 45, No. 1, 2008, pp. 142–158.

[23] Kulfan, B. M. and Bussoletti, J. E., "Fundamental parametric geometry representations for aircraft component shapes," *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, September 2006.

[24] Kulfan, B. M., "A universal parametric geometry representation method-CST," *45th AIAA Aerospace Sciences Meeting and Exhibit*, January 2007.

[25] Kulfan, B. M., "Modfication of CST airfoil representation methodology," Retrieved from http://www.brendakulfan.com/docs/CST8.pdf.

[26] Hicks, R. M. and Henne, P. A., "Wing design by numerical optimization," *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412.

[27] Wu, H.-Y., Yang, S., Liu, F., and Tsai, H.-M., "Comparison of three geometric representations of airfoils for aerodynamic optimization," *16th AIAA Computational Fluid Dynamics Conference, Orlando, Florida*, 2003.

[28] Masters, D. A., Taylor, N. J., Rendall, T. C. S., Allen, C. B., and Poole, D. J., "Review of Aerofoil Parameterisation Methods for Aerodynamic Shape Optimisation," *53rd AIAA Aerospace Sciences Meeting*, Jan 2015.

[29] Wendland, H., *Scattered data approximation*, Cambridge University Press Cambridge, 2005.

[30] Buhmann, M. D., "Radial basis functions," *Acta numerica*, Vol. 9, 2000.

[31] 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, pp. 1519–1559.

[32] 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, No. 8, Nov 2008, pp. 827–860.

[33] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Optimal Domain Element Shapes for Free-Form Aerodynamic Shape Control," *53rd AIAA Aerospace Sciences Meeting*, 2015.

[34] Toal, D. J. J., Bressloff, N. W., Keane, A. J., and Holden, C. M. E., "Geometric filtration using proper orthogonal decomposition for aerodynamic design optimization," *AIAA Journal*, Vol. 48, 2010, pp. 916–928.

[35] Ghoman, S. S., Wan, Z., Chen, P. C., and Kapania, R. K., "A POD-based reduced order design scheme for shape optimization of air vehicles," *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2012.

[36] 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.

[37] Jameson, A., "Advances in aerodynamic shape optimization," *Computational Fluid Dynamics 2004*, Springer, 2006, pp. 687–698.

[38] Jameson, A., "Aerodynamic shape optimization using the adjoint method," *Lectures at the Von Karman Institute, Brussels*, 2003.

[39] Schmidt, S., Ilic, C., Gauger, N., and Schulz, V., "Shape gradients and their smoothness for practical aerodynamic design optimization," *Optim. Eng.(20 08) Preprint-Number SPP1253-10-03*, 2008.

[40] Allen, C. B., "Parallel Universal Approach to Mesh Motion and Application to Rotors in Forward Flight," *International Journal for Numerical Methods in Engineering*, Vol. 69, No. 10, 2007, pp. 2126–2149.

[41] Parpia, I. H., "Van-Leer Flux Vector Splitting in Moving Coordinates," *AIAA Journal*, Vol. 26, 1988, pp. 113–115.

[42] Allen, C. B., "Multigrid Convergence of Inviscid Fixed- and Rotary-Wing Flows," *International Journal for Numerical Methods in Fluids*, Vol. 39, No. 2, 2002, pp. 121–140.

[43] Jameson, A., "Airfoils Admitting Non-unique Solutions to the Euler Equations," *AIAA 22nd Fluid Dynamics, Plasmadynamics and Lasers Conference*, Honolulu, Hawaii, 1991, AIAA Paper 91–1625.

[44] Jameson, A., Vassberg, J. C., and Ou, K., "Further Studies of Airfoils Supporting Non-Unique Solutions in Transonic Flow," *AIAA Journal*, Vol. 50, No. 12, 2012, pp. 2865–2881.

[45] Ou, K., Jameson, A., and Vassberg, J. C., "Airfoils Supporting Non-unique Transonic Solutions for Unsteady Viscous Flows," *7th AIAA Theoretical Fluid Mechanics Conference*, Atlanta, Georgia, 2014, AIAA Paper 2014–2927.