Masters, D. A., Taylor, N. J., Rendall, T. C. S., Allen, C. B., & Poole, D. J. (2016). A Geometric Comparison of Aerofoil Shape Parameterisation Methods. In *54th AIAA Aerospace Sciences Meeting* [AIAA 2016-0558] American Institute of Aeronautics and Astronautics Inc. (AIAA). <https://doi.org/10.2514/6.2016-0558>

Peer reviewed version

Link to published version (if available):
[10.2514/6.2016-0558](https://doi.org/10.2514/6.2016-0558)

[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 AIAA at <http://arc.aiaa.org/doi/10.2514/6.2016-0558>. 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/>

A Geometric Comparison of Aerofoil Shape Parameterisation Methods

D. A. Masters*,

Department of Aerospace Engineering, University of Bristol

N. J. Taylor[†],

MBDA UK Ltd, Filton

T. C. S. Rendall[‡], C. B. Allen[§] and D. J. Poole[¶]

Department of Aerospace Engineering, University of Bristol

A comprehensive review of aerofoil shape parameterisation methods that can be used for aerodynamic shape optimisation is presented. Seven parameterisation methods are considered for a range of design variables: CSTs; B-Splines; Hicks-Henne bump functions; a Radial Basis function (RBF) domain element approach; Bèzier surfaces; a singular value decomposition modal extraction method (SVD); and the PARSEC method. Due to the large range of variables involved the most effective way to implement each method is first investigated. Their performance is then analysed by considering the geometric shape recovery of over 2000 aerofoils using a range of design variables, testing the efficiency of design space coverage with respect to a given tolerance. It is shown that, for all the methods, between 20 and 25 design variables are needed to cover the full design space to within a geometric tolerance with the SVD method doing this most efficiently. A set transonic aerofoil case studies are also presented with geometric error and convergence of the resulting aerodynamic properties explored. These results show a strong relationship between geometric error and aerodynamic convergence and demonstrate that between 38 and 66 design variables may be needed to ensure aerodynamic convergence to within one drag and one lift count.

Nomenclature

a_i	i th design variable
$\mathbf{a}$	Vector of design variables
$B_{i,n}$	Bernstein polynomial i of n
C_D	Drag coefficient
C_L	Lift coefficient
$f(\gamma)$	Cubic spline function
$\mathbf{F}$	Generic transformation matrix
h_i	Location Hicks-Henne bump function maxima
$\mathbf{H}$	RBF full transformation matrix
i, j	Counters
k	B-Spline order
m^*	NACA 4-series maximum camber parameter
M	Number of training aerofoils
M_∞	Freestream Mach number
n, m	Number of basis functions or control points
N	Number of aerofoil coordinates
$N_{i,k}$	i th B-Spline basis function
$\mathbf{p}(X)$	RBF polynomial function
p^*	NACA 4-series maximum camber position

*Graduate Student, AIAA Student Member, dominic.masters@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

[¶]Graduate Student, AIAA Student Member, d.j.poole@bristol.ac.uk, Bristol, BS8 1TR, UK

$\mathbf{P}_i, \mathbf{P}_{ij}$	i th or i,j th control point
$S(x)$	CST shape function
S_R	RBF support radius
t^*	NACA 4-series thickness parameter
t_i	Thickness of Hicks-Henne bump function
$\mathbf{T}$	Matrix of training aerofoils
u	Scalar parameterisation variable
$\mathbf{U}$	Left singular vectors of M
$\mathbf{v}_i$	i th SVD mode
$\mathbf{V}$	Right singular vectors of M
w_i	Weighting parameters
$\mathbf{W}$	Diagonal matrix of weighting parameters
x, z	Continuous normalised coordinate directions
X	Two-dimensional position variable (x, z)
x_i, z_i	Position of i th aerofoil coordinate
$\mathbf{x}, \mathbf{z}$	Vector of all points x_i or z_i
X_i	i th vector coordinate position (x_i, z_i)
$\mathbf{X}$	Two-dimensional aerofoil $(\mathbf{x}, \mathbf{z})$
$\bar{\mathbf{z}}$	Mean $\mathbf{z}$ coordinates
$\hat{\mathbf{z}}$	Transformed $\mathbf{z}$ coordinates
α	Angle of attack
$\boldsymbol{\beta}_{DE}$	RBF coefficient vector
γ	Cubic spline parameterisation
ξ_i	Arc length from X_{le} to X_i
ρ	Cubic spline smoothing parameter
$\boldsymbol{\Sigma}$	Diagonal matrix of singular values of M
$\phi_i(x)$	i th basis function

Subscripts

le/te	Leading edge / Trailing edge
min/max	Minimum or maximum
$upper$	Indicates the upper aerofoil surface
$lower$	Indicates the lower aerofoil surface

Superscripts

$approx$	Least Squares approximation of target aerofoil
$error$	Positive aerofoil approximation error
$initial$	Initial aerofoil to be deformed
$target$	Target aerofoil to approximate

I. Introduction

With optimisation becoming more common in aerodynamic design, a significant effort is being made to improve the effectiveness and efficiency¹⁻⁵ of the full optimisation process. A typical CFD-based aerodynamic optimisation consists of four stages: shape parametrisation/control; mesh creation/deformation; flow solution; and optimisation. Within this context the efficiency is typically measured by the time or resources needed to complete the optimisation and the effectiveness by the quality of the final solution. This paper focuses on how the shape parameterisation can influence these properties for two-dimensional aerofoil design.

Shape parameterisation concerns the way the geometry is handled and deformed by the optimisation algorithm, determining both the fidelity and range of control available. An effective and efficient parameterisation method is characterised by the ability to cover a large design space with a limited set of design variables. This paper presents a direct comparison of widely-used shape parameterisation methods, by considering geometric shape recovery and analysing how accurately each can recover a large range of different aerofoils. A further study is then performed on a reduced set of aerofoils to assess the influence of the geometric accuracy on the aerodynamic properties. For each method a variety of techniques for implementing them have also been considered in order to identify the optimum conditions for their use. This allows a fair final comparison to be made between all parameterisation methods considered.

II. Background

A wide range of methods have previously been used for aerofoil geometry representation. These vary from general geometry representation techniques such as B-Splines to aerofoil specific methods such as PARSEC⁶ and all have a wide range of characteristics and attributes. For the purpose of this study the various aerofoil parameterisation methods will be categorised as either constructive or deformative based on how the design variables/parameters influence the aerofoil surface. Constructive methods represent the aerofoil surface based purely on the values of the parameters specified; examples of this include polynomials and splines^{7,8}, partial differential equation methods (PDE)⁹ and CSTs^{10,11}. Deformative methods take an existing aerofoil then deform it to create the new shape; these include discrete¹², analytical¹³, basis vector¹⁴ and free-form deformation (FFD)^{15,16} methods.

One of the simplest forms of parameterisation is the discrete (or free-surface) method. This directly uses the surface points of a discretely defined aerofoil as the design variables¹². The benefit of this is that it allows extremely fine control over the shape with absolutely no restriction on the design space. The size and complexity of the resulting optimisation problem can however cause significant difficulties. For example, as all the point displacements are considered independently, the resulting sensitivities are often not smooth, which can present difficulties for flow solvers if not appropriately handled¹⁷. The large number of design variables involved can also lead to slow convergence rates and extremely expensive finite-difference gradient calculations. For these reasons more robust and efficient parameterisation methods are usually favoured over the discrete method.

Hicks and Henne's¹³ early analytical approach based on bump functions represents one alternative. It takes a base aerofoil and then adds a linear combination of single-signed sine functions to deform its upper and lower surfaces to create a new aerofoil shape. This concept of adding a linear combination of simple basis functions to a base shape has also been used in a constructive manner, for instance, Kulfan's CST¹⁰ method which adds a combination of Bernstein polynomials to a simple, analytical 'aerofoil class' shape. Both of these methods have seen frequent use within the framework of aerodynamic optimisation^{18–24}.

Other, more classical methods, such as B-Splines or polynomial fitting are also commonly used. B-Splines represent a class of versatile, piecewise polynomial, control point based curves with variable continuity and support. Due to their intuitiveness and flexibility they have been applied to a wide range of applications with extensive use throughout shape optimisation including a range of aerofoil specific cases^{1,25–29}. Sobieczky's⁶ PARSEC (Parameterised Sections) method is also popular, approximating each surface by a 6th order polynomial. One positive feature of this method is that it uses real geometric properties such as the aerofoil's crest location and curvature as the design variables, allowing more intuitive control of the shape. However, as the method is limited to only 12 design variables it does not provide the range or flexibility in fidelity made available by many of its alternatives.

Attempts have also been made to mathematically derive a set of orthogonal modes to represent an aerofoil. This is typically done through proper orthogonal decomposition (POD) of a set of training aerofoils which will create a set of optimal orthogonal shape modes based on a range of training data. Studies of this nature have been produced by Toal *et al.*³⁰, Ghoman *et al.*³¹ and then by Poole *et al.*³² who used a large, varied collection of aerofoils and singular value decompositions to produce a universal set of modes representing the deformation of aerofoil shapes.

Another approach to shape parameterisation is to use FFD which is a method typically used in soft object animation. This creates a smooth continuous volume transformation based on the change in position of a series of control points. This volume transformation can also be used to deform computational volume meshes seamlessly with the aerofoil. This can have significant cost benefits particularly in three dimensions. The two principal FFD techniques in use are radial basis functions (RBFs)¹⁵ applied on an arbitrary domain element (a series of user positioned initial control points), and Bèzier surfaces³³ (often referred to as just 'FFD') which use a structured lattice of initial control points. These control points are commonly grouped together to create global transformations such as thickness and camber, or twist and sweep in three dimensions, to reduce the total number of design variables. The RBF domain element method also allows the local fidelity of movement to be controlled through the proximity and density of the point distribution. Both of these methods have shown promising optimisation results^{16,34–36}.

A range of other comparative shape parameterisation studies have been previously investigated^{37–40}. Castonguay and Nadarajah³⁷ considered the discrete method, Hicks-Henne, B-Splines and PARSEC for inverse design and drag minimisation on an ONERA M6 (D section) aerofoil. They found, for the inverse design test case, that the B-spline needed 32 control points (64 design variables) to obtain a satisfactory fit for the geometry but only 16 control points (32 design variables) to get a satisfactory fit on the pressure distribution. They used a fixed bump position, fixed bump width configuration of Hicks-Henne and found they needed 32 bumps to get a satisfactory fit for both the geometry and the pressure. B-Splines, for equal design variables, always out-performed Hicks-Henne, and PARSEC failed to approximate the ONERA M6 to a useful degree of accuracy. For the drag minimisation test case both the B-Splines and Hicks-Henne bumps gave comparable results.

This study was extended by Mousavi, Castonguay and Nadarajah³⁸ who also performed an inverse design case on

the ONERA M6 aerofoil and a drag minimisation (though on a NACA0012) but instead used the discrete, B-spline and CST parameterisation methods. They found very similar results for B-Splines as Castonguay and Nadarajah³⁷ and found that the CST method gave good results with 22 design variables (comparable to the best B-spline cases) but then decreased in accuracy with additional design variables due to the influence of high frequency oscillations. The drag optimisation showed comparable results across both methods and all orders of accuracy with the CST method using only 10 design variables attaining a 40% drag reduction.

The objective of this study is to present a comprehensive comparison of the shape parameterisation methods available for aerodynamic optimisation. This adds to the previous studies in this area by including a larger range of methods and testing them over a significantly broader range of aerofoils. Furthermore each parameterisation scheme is tested under a variety of conditions in order to identify best practice for implementing them; this allows a fair and effective comparison to be made between all the methods investigated. Selected aerofoil case studies are also included to provide insight into the relationship between the geometric accuracy and their aerodynamic properties.

III. Methodology

Seven parameterisation methods have been considered in this work; B-Splines, CST, SVD, PARSEC, Hicks-Henne bump functions, Bezier surface FFD and RBF domain elements. Each method has first been presented individually with varying configurations to determine the optimal method of implementation. A further comparison of all the methods has then been completed where the success of each parameterisation method is evaluated on its ability to geometrically reconstruct a library of aerofoils to a prescribed tolerance. The percentage of these aerofoils that can be reconstructed to this tolerance can then be quantified to give a good estimate of the available design space for each method. This was then calculated for a sweep in design variables, giving a perspective on the efficiency of the total design space coverage.

A further set of case studies has also been performed on five common aerofoils. For each of these investigations a full set of aerofoil approximations has been constructed for each parameterisation method along with aerodynamic solutions calculated using the full potential flow solver VGK⁴¹. The aim of this is to compare how the aerodynamic quantities of the aerofoils (C_L , C_D) converge as the number of design variables increases. Before performing these tests the associated methodology must first be examined.

A. Aerofoil Libraries

For the large database geometry tests two distinct aerofoil libraries have been used; a library of NACA 4-series aerofoils and a library of aerofoils taken from the UIUC aerofoil database^a. A further ‘combined’ library that consists of the union of these two sets is also considered. In order to accommodate the broad range of parameterisation methods considered, the geometry of all the aerofoils has been normalised and discretized such that they satisfy all the constraints imposed by the methods themselves and provide an impartial and consistent testing platform. Further explanation of this can be found in appendix A.

The NACA aerofoil library has been created by taking all the possible integer combinations of the three input parameters t^* , m^* and p^* for the NACA 4-series definition (Appendix B) in the ranges $t^* \in [6, 24]$, $m^* \in [0, 9]$ and $p^* \in [3, 7]$. The aerofoils were then normalised and discretized as described in appendix A producing a final library of 874 unique aerofoils.

The original UIUC aerofoil library^a consists of raw data for 1517 aerofoils taken from a wide range of aerospace applications. The aerofoils are discretized with between 24 and 205 points with varying levels of precision and no information as to how they were formed. For these reasons care has been taken to appropriately refine them and normalise them to the specifications outlined in appendix A. It was also found that if an interpolating cubic spline was used, small oscillations between the raw data points were created for some of the aerofoils, particularly for those that are coarsely defined. For this reason an energy minimizing smoothing spline⁴² was employed to re-sample the original aerofoil data. This is defined as the spline f such that the function

$$\rho \sum_i |z_i - f(\gamma_i)|^2 + (1 - \rho) \int |f''|^2 \quad (1)$$

is minimised for some parameter ρ , taken as the optimum determined by de Boor⁴², and a spline parameterisation γ_i .

^ahttp://aerospace.illinois.edu/m-selig/ads/coord_database.html accessed 5th June 2014.

The spline parameterisation was chosen to be

$$\gamma_i = \begin{cases} +\sqrt{\xi_i} & \text{for } i \in \text{upper surface} \\ -\sqrt{\xi_i} & \text{for } i \in \text{lower surface} \end{cases} \quad (2)$$

where ξ_i is the arc length between points X_i and X_{le} . This was used as it helped increase the smoothness around the high curvature (often densely discretized) leading edge.

Given an appropriate spline f the aerofoil could then be sampled and normalised to satisfy the geometry normalisation criteria. A final visual check was performed on the set of aerofoils where any aerofoils that were not deemed to be suitable for the tests were discarded. Aerofoils were mainly discarded due to errors in the raw data or unsatisfactory surface smoothness, giving a final database of 1300 aerofoils for testing.

In previous work by the authors⁴³ a similar study to this was carried out, though it used a simple smoothing spline for the re-sampling procedure opposed to the smoothing spline described above. This had a negative influence on the final results and is thought to be the major contributing factor to the variation between the results presented in⁴³ and this work.

B. Initial Aerofoil for Deformative Methods

For all of the deformative parameterisation methods an initial aerofoil shape must be provided which is then deformed to provide the final aerofoil; for tests in this work a NACA0012 has been used. This choice was made due to its analytic description and its smooth shape, which should provide a well-rounded starting point for deformation to the full range of aerofoils provided in both the NACA and UIUC libraries. It should however be noted that the NACA0012 (as defined by equation 37) has a sharp trailing edge yet the final UIUC database consists of aerofoils with both blunt and sharp trailing edges. As the deformative methods will be unable to overcome these topological difficulties a slightly modified NACA0012 shape has been used. It has been defined as

$$NACA0012^*(x) = 0.6 \left(0.2969\sqrt{x} - 0.126x - 0.3516x^2 + 0.2843x^3 - 0.1036x^4 \right) \pm x \cdot z_{te}^{target} \quad (3)$$

where the target aerofoil half trailing edge thickness z_{te}^{target} is provided as a known quantity. This value is also supplied to all the constructive methods in order to provide a fair and unbiased testing environment.

C. Solution Calculation and Error Tolerance

For each individual test the best approximation (or reconstruction) has been calculated through a least squares solution, from which the errors in the geometry can then be calculated.

The least squares solution $\mathbf{X}^{approx} = (\mathbf{x}^{approx}, \mathbf{z}^{approx})$ is the solution that minimises the expression

$$\frac{1}{N} \sum_{i=1}^N w_i (z_i^{target} - z_i^{approx})^2 \quad (4)$$

for some set of weights w_i given that $x_i^{target} = x_i^{approx} = x_i$ for $i = 1 : N$.

For each of the parameterisation methods, $\mathbf{z}^{approx}$ can be expressed in the form of a linear combination

$$\mathbf{F}\mathbf{a} = \mathbf{z}^{approx} \quad (5)$$

where $\mathbf{F}$ is the transformation matrix defined by the shape parameterisation method and $\mathbf{a}$ is a column vector of the design variables. The weighted least squares solution can then be calculated as

$$\mathbf{a} = (\mathbf{W}\mathbf{F})^+ \mathbf{W}\mathbf{z}^{target} \quad (6)$$

where the superscript $+$ denotes the Moore-Penrose pseudoinverse and $\mathbf{W}$ is a diagonal matrix where $W_{ii} = w_i \forall i$.

To obtain measurable, comparable results it is important that the approximation data is processed into a well defined error metric that reflects the aims of the experiment; in this case the ability of shape parameterisation methods to replicate different aerofoil shapes. In a study by Kulfan¹⁰, confined to the CST method, the geometric errors in the solution aerofoils were frequently compared to a ‘typical wind tunnel tolerance’ defined as

$$z^{error} < \begin{cases} 4 \times 10^{-4} & \text{if } x/c < 0.2 \\ 8 \times 10^{-4} & \text{if } x/c > 0.2 \end{cases} \quad (7)$$

with solutions deemed acceptable if they were within this range. For a metre long chord this tolerance is equivalent to 0.4mm for the leading edge and 0.8mm everywhere else. Kulfan then showed for two cases, the RAE2822 and NASA NSC 2-0714, that L/D converged to within $\approx 2\%$ for errors within this bound.

This tolerance (equation 7) has been used throughout the testing in this paper as a benchmark for a successful approximation. Each method is then typically assessed on the percentage of the aerofoil library approximated to this geometric tolerance versus the number of design variables used. Further tolerances have also been considered which have been expressed in terms of weighted z error where

$$z^{weight} = \begin{cases} 2 \times z^{error} & \text{if } x/c < 0.2 \\ z^{error} & \text{if } x/c > 0.2. \end{cases} \quad (8)$$

In this notation, Kulfan's geometric tolerance is equal to $z^{weight} < 8 \times 10^{-4}$. These tolerances motivate using the least squares weighting

$$w_i = \begin{cases} 2 & \text{if } x_i < 0.2 \\ 1 & \text{if } x_i > 0.2 \end{cases} \quad (9)$$

to provide a tighter fit in the leading edge area.

D. Large Database Tests

The large database tests performed in this work consist of testing whether a range of aerofoils can be successfully reconstructed for each parameterisation method across a large range of design variables. Each individual test is implemented by applying a least squares approximation to a given target aerofoil for a particular parameterisation. If the resulting parameterised aerofoil is within the a prescribed tolerance of the target it is deemed successful. For the majority of the tests this tolerance is defined as Kulfan's geometric tolerance (equation 7), though the influence of scaling this has also been investigated. The success of this test therefore defines if the particular target aerofoil is included in the design space of the parameterisation method tested. This process is then repeated for all the aerofoils in the UIUC and NACA aerofoil libraries and for between 1 and 100 design variables. This enables the percentage of the testing database successfully reconstructed to be calculated for each design variable interval and therefore give an estimation of the total aerofoil design space covered. The resulting data is then used to quantify the efficiency of the aerofoil design space coverage for the particular parameterisation method.

E. Case Studies

In addition to the large database tests, five case studies have also been considered. For each case study aerofoil, a full range of reconstructions have been calculated for each parameterisation method at each design variable interval. An aerodynamic flow solution is then calculated for each of these approximations using the Euler solver SU²⁴⁴. This allows the behaviour of the aerodynamic properties and the max weighted z error with respect to the number of design variables to be analysed.

IV. Parameterisation Methods

The following section outlines each of parameterisation methods tested in this papers and investigates the various options for their implementation. The methods have been categorised as either a constructive method or a deformative method based on the relationship between the design variables/parameters and the aerofoil surface. Constructive methods define the aerofoil surface geometry purely based on the magnitude of the design variables whereas deformative methods take an existing aerofoil shape and deform it to create the new shape. The constructive methods investigated in this paper are: B-Splines, CSTs, and the SVD method and the deformative methods are: Hicks-Henne bump functions, Bèzier surface FFD and the RBF domain element method. For each method, its formal definition is described and then an investigation into the best technique for its implementation is performed. This has been done in an attempt to provide the best results possible for each method and therefore a fair and unbiased comparison of their performance.

A. Constructive Methods

1. B-Splines

B-Splines are a widely used family of method used for producing multi-dimensional smooth curves that are defined based on the product of a set of basis functions and a set of spatially defined discrete control points $\mathbf{P}_i \in \mathbb{R}^3$. Typically

they have up to three degrees of freedom based on spatial position for every control points used plus, in the case of non-uniform rational B-Splines (NURBS), additional degrees of freedom in the form of a weighting for each basis function and variable knot positions. Though these additional parameters do increase the flexibility of the B-Splines the increase in the design variables and complexity is not deemed beneficial for this study. Therefore only the case of uniform, non-rational B-Splines will be considered in this work.

Given these assumptions the B-Spline curves are parametrised by the scalar $u \in [0, 1]$, and are defined by the curve

$$X(u) = \sum_{i=0}^{n-1} N_{i,k}(u) \mathbf{P}_i \quad (10)$$

for n basis functions $N_{i,k}$ of order $k < n$ and control points $\mathbf{P}_i$. Additionally if $k = n - 1$ the B-Splines are called ‘Bèzier curves’ and the basis functions are the Bernstein polynomials of order k . A full and detailed description of B-Spline curves can be found in Piegel and Tiller⁴⁵.

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.

B-Splines can be used to represent aerofoils in a variety of different ways but in order to satisfy the constraints imposed by the other parameterisation methods the following configuration has been used. Each aerofoil is represented by two B-Splines; one for each of the upper and lower surfaces. For each B-spline, $\mathbf{P}_0$ is fixed at the leading edge $(0, 0)$, $\mathbf{P}_{n-1}$ is fixed at the trailing edge and $\mathbf{P}_1$ is aligned vertically with the leading edge. The movement of each point has also been restricted to just the vertical z -axis. This constrains each control point to just a single design variable and allows a simple linear least squares solution to be used. The distribution of the control points in the chord-wise direction, however, is not restricted by this. Two possible distributions will be investigated: a uniform and a cosine distribution.

For the uniform distribution the positions of the control points are formally described as

$$\mathbf{P}_0 = (0, 0), \quad \mathbf{P}_i = \left(\frac{i-1}{n}, a_i \right), \quad \mathbf{P}_{n-1} = (1, z_{te}), \quad (11)$$

whereas for the cosine distribution they are defined as

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

where a_i denotes a design variable and z_{te} the z -wise component of the trailing edge position.

Each of these configurations has been tested using the methodology for the ‘large database tests’ outlined in section III.D for increasing orders of B-Spline (up to a maximum of 15). It was found that the influence of the order on results was significantly different for the two chord-wise control point configurations. For the uniformly distributed case increasing the order of the spline significantly improves the accuracy whereas increasing the number of control points has less influence. For the cosine distributed case improvements in accuracy are primarily driven by increasing the number of control points used, with only marginal gains made by increasing the order of the spline.

For the uniformly distributed case the best results are obtained when the order is at its maximum value for the given number of control points. This relates to an order of $k = \min(n - 1, 15)$ where n is the number of control points in the spline. For the cosine distributed configuration it is less clear which order to use, however it was found that using $k = \min(n - 2, 15)$ gave consistently good results. Comparing the results for these best choices of order it was found that both configurations provided very similar performance. The dependence of the uniform distribution on using a high order is however not considered desirable as the computation time for initialising the B-spline is exponentially related to the order. For this reason the cosine distributed configuration will be considered hereafter.

2. Class Function/ Shape Function Transformations (CST)

The CST method was developed by Kulfan^{10,11} primarily as a method of analytically 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. Each aerofoil surface is defined by

$$z(x) = \left(x^{0.5} \cdot (1 - x) \right) \cdot S(x) + x \cdot z_{te} \quad (13)$$

where z_{te} defines the trailing edge half thickness, $S(x)$ is a ‘shape function’ and $x \in [0, 1]$.

Kulfan^{11,46} suggested defining $S(x)$ as the linear combination of Bernstein polynomials i.e.

$$S(x) = \sum_{i=0}^n a_i B_{i,n}(x) + \underbrace{a_{n+1} x^{0.5} (1-x)^{n-0.5}}_{LEM}. \quad (14)$$

where $B_{i,n}(x)$ is a Bernstein polynomial of degree n and coefficient a_i . The second term is an optional leading edge modification (LEM) presented later⁴⁶ to improve the fidelity at the leading edge - thought at the cost of an additional design variable.

Large database tests were performed using the CST method both with and without the LEM. It was found that using the LEM significantly improved the over all results. For example, in order to successfully reconstruct 80% of the ‘combined’ library to the geometric tolerance without the LEM, 22 design variables were needed, opposed to just 16 design variables when it was used. For this reason the CST method will be implemented with the LEM for the rest of this work.

3. Singular Value Decomposition (SVD)

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.*³⁰ then by Ghoman *et al.*³¹ and Poole *et al.*³². Ghoman *et al.*³¹ used a series of supercritical aerofoils to derive the modes and showed that other supercritical aerofoils could efficiently be reconstructed. Poole *et al.*³² 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.

To formulate these modes the $\mathbf{z}$ coordinates of the training aerofoils, translated around the mean $\bar{\mathbf{z}}$ such that $\hat{\mathbf{z}} = \mathbf{z} - \bar{\mathbf{z}}$, are used to form the rows of the matrix

$$\mathbf{T} = \begin{bmatrix} \hat{z}_1^1 & \hat{z}_2^1 & \cdots & \hat{z}_N^1 \\ \hat{z}_1^2 & \hat{z}_2^2 & \cdots & \hat{z}_N^2 \\ \vdots & \vdots & \ddots & \vdots \\ \hat{z}_1^M & \hat{z}_2^M & \cdots & \hat{z}_N^M \end{bmatrix} \quad (15)$$

for M training aerofoils each of length N .

This matrix is then decomposed into the singular value decomposition

$$\mathbf{T} = \mathbf{U} \cdot \mathbf{\Sigma} \cdot \mathbf{V} \quad (16)$$

where columns of $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{\min(N,M)}]$ represent the ordered, orthogonal aerofoil modes.

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

$$\mathbf{z} = \bar{\mathbf{z}} + \sum_i a_i \mathbf{v}_i + z_{te} \mathbf{x} \quad (17)$$

where a_i represents the design variables and z_{te} represents the trailing edge half thickness.

The design space of the SVD method is totally dependent on the library of aerofoils used to make the modes. For this reason three different mode sets have been investigated in this work; one that has been created using the UIUC library, one that has been created using the NACA library and another that used the combined UIUC and NACA libraries. Figure 1 shows the mean aerofoil shape and first five mode shapes created using the combined UIUC and NACA libraries. One feature of modes created in this fashion is that their frequency increases linearly. Figure 2 then compares the large database results for these three different cases when tested on each of the three libraries. Unsurprisingly, for each case the best results came when the testing library and the mode creation library were the same. The results when the modes are created with the NACA library show a large variation in accuracy. When tested on the NACA library over 99% of the library are approximated to Kulfan tolerance with just 6 modes, however when tested on the UIUC library the results are very poor. This indicates the consequences of creating the mode library from

a restricted set of aerofoils. In contrast it can be seen that when the modes are created using the broad UIUC library the results are fairly consistent across the three testing libraries. Due to the large variation in these results, all three SVD mode libraries will be tested in the method comparison.

Figure 1. The mean aerofoil shape and the first 5 modes created using the combined UIUC and NACA library.

Figure 2. Library reconstruction results for the SVD method for modes created with three different libraries (different lines types) tested on the three different libraries (a)-(c).

4. PARSEC

The PARSEC method was developed by Sobieczky⁶ as a system to create analytically defined aerofoils based on meaningful properties such as the upper/lower crest position, max thickness, leading edge radius and boat-tail angle. This was done by proposing that the upper and lower surfaces should each be defined by the following 6th order polynomial

$$z(x) = \sum_{i=1}^6 a_i x^{i-0.5}. \quad (18)$$

Twelve equations, subject to twelve free parameters that define the characteristics of an aerofoil, were then defined with the resulting system solved to form the PARSEC parameters. It should however be noted that as the trailing edge

positions are taken to be a known quantity in these tests (see section III.B) the number of free design variables has been reduced from 12 to 10.

B. Deformative Methods

1. Hicks-Henne Bump Functions

Hicks-Henne bump functions¹³ use a base aerofoil definition plus a linear combination of a set of n basis (or bump) functions. The bumps are defined using an augmented sine function so each surface is defined by

$$\mathbf{z} = \mathbf{z}^{initial} + \sum_{i=0}^n a_i \sin^{t_i} \left(\pi \mathbf{x}^{\ln(0.5)/\ln(h_i)} \right) \quad (19)$$

for coefficients a_i for $i = 1, \dots, n$ where h_i determine location of the bump function maxima and t_i their width.

Each bump function is therefore defined by three variables, each of which can be optimised or fixed. Various combinations of fixing and optimising these variables have been performed, for example Wu¹⁸ opted to fix $t_i = 4$ and let

$$h_i = \frac{1}{2} \left[1 - \cos \left(\frac{i\pi}{n+1} \right) \right], \quad i = 1, \dots, n. \quad (20)$$

whereas Khurana⁴⁷ varied all three variables.

For this study the initial aerofoil has been taken as the modified NACA0012 defined in equation 3, a_i have been varied and represent the design variables, h_i have been defined as equation 20 and the thickness parameters t_i have been taken as

$$t_i = 2 \left(\frac{n-i}{n-1} \right)^3 + 1 \quad \text{for } i = 1, \dots, n. \quad (21)$$

This gives the forward most basis functions, $t_1 = 3$, and the rear most basis function, $t_n = 1$. This configuration has been found to give the most efficient design space coverage.

2. Bèzier Surface FFD

The Bèzier surface FFD technique used here follows the work originally proposed by Sederberg and Parry⁴⁸. This deforms a two-dimensional shape by embedding it with a Bèzier surface constrained to a plane then controlling the associated control points. The method used here is outlined in detail in Sederberg⁴⁹. This describes the deformation of an initial aerofoil $\mathbf{X}^{initial} = [\mathbf{x}^{initial}, \mathbf{z}^{initial}]$ aerofoil with respect to a set of $m \times n$ Bèzier surface control points $\mathbf{P}_{ij}$ by

$$\mathbf{X} = \sum_{j=0}^n \sum_{i=0}^m B_{i,m}(\mathbf{x}^{initial}) B_{j,n} \left(\frac{\mathbf{z}^{initial} - z_{min}}{z_{max} - z_{min}} \right) \mathbf{P}_{ij}. \quad (22)$$

where $B_{i,m}$ are the set of Bernstein polynomials of degree m . The initial control point positions are then given by

$$\mathbf{P}_{ij}^{initial} = \left(x_{min} + \frac{i}{m} (x_{max} - x_{min}), z_{min} + \frac{j}{n} (z_{max} - z_{min}) \right). \quad (23)$$

For this analysis, in agreement with the other methods, the modified NACA0012 (equation 3) has been chosen as the initial aerofoil and the control points have been constrained to just move in the z direction. This leaves the configuration of the control point lattice as the main option to be determined. To investigate the various options the large database tests have been performed on a range of different row and column combinations. Interestingly the results for the two different libraries convey different trends with the ‘3 row’ configuration showing to be the most efficient the NACA library whereas for the UIUC results the ‘4 row’ configuration was the best. This could be due to the fact that the NACA library is formed based on a symmetric thickness projection from a camber line. This means that for Bèzier surface lattices with an odd number of rows the central row can approximate the camber line, and the outer rows the thickness projection. The required thickness profile would also be embedded in the initial NACA0012 shape used, providing further benefits when approximating the NACA library. It seems, however, for the UIUC aerofoil library, where the upper and lower aerofoil surfaces may have unrelated shapes, that a ‘4 row’ lattice is most accurate. Due to this variation in results, both the ‘3 row’ and ‘4 row’ configurations will be considered hereafter. The comparison of these two configurations can be seen in the final results graphs, figures 5-7.

Figure 3. Example deformation of a NACA0012 using a 3×4 Bèzier surface control point lattice.

3. Radial Basis Function Domain Elements

The RBF domain element 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, 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 Wendland⁵⁰ and Buhmann⁵¹; the formulation used here is then presented extensively in Rendall and Allen¹⁵ and its use as a parametrisation technique in Morris *et al.*³⁴. They define that an aerofoil $\mathbf{X}^{initial}$, deformed relative to the position of a set of RBF domain element control points $\mathbf{P}_i$, is given by

$$\mathbf{X} = \sum_{i=1}^n \beta_{DE_i} \phi(\|\mathbf{X}^{initial} - \mathbf{P}_i\|/S_R) + \mathbf{p}(\mathbf{X}^{initial}) \quad (24)$$

where β_{DE} is coefficient vector dependant on the initial control point positions, ϕ is the radial basis function, S_R is the support radius and $\mathbf{p}(\mathbf{X}^{initial})$ is a linear polynomial used to ensure that translation and rotation are captured without added shape deformation. It is then further shown that this can be decomposed into a simple matrix transformation

$$\mathbf{X} = \mathbf{H}\mathbf{P}. \quad (25)$$

where $\mathbf{H}$ is invariant of the current control point positions so is only needed to be calculated once.

There are a few factors that affect the use of the RBF 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 function⁵⁰. Similarly to Bèzier Surface FFD, the modified NACA0012 (equation 3) will be used as the initial aerofoil and the control points will be constrained to the z direction.

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 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 challenging⁵². For this reason two simple control point schemes have been proposed and optimised under a series of constraints, with both the unoptimised and optimised forms tested on the full aerofoil database.

The two initial control point schemes that have been considered are, an 'off-surface' and an 'on-surface' configuration. The off-surface configuration is specified symmetrically about the x -axis on an ellipse around the initial aerofoil with the first, 6 point, configuration having points located at $x = -0.1, 0.5$ and 1.1 . The on-surface configuration is specified on the surface where the initial 4 points are specified at the leading and trailing edges as well as at the half chord points on the upper and lower surface. Note the leading and trailing edge points for this configuration will always be held fixed.

Additional points are then placed using either a 'standard' or 'optimised' procedure. For the 'standard' method additional points are placed symmetrically with x locations bisecting two points from the previous configuration; points are placed alternately at the largest interval towards the leading edge and then the trailing edge (figure 4). For the 'optimised' method the bisection of each possible interval from the previous configuration is tested, with points again added symmetrically about the x -axis. The interval which provides the best results, where 'best' is defined as the largest percentage of the full aerofoil library (both NACA and UIUC) reconstructed to within the geometric tolerance,

is then used. Both of these methods are repeated iteratively to provide a range of control point configurations of increasing density.

Large database tests were then performed for the on-surface and off-surface control point schemes for both their standard and optimised configurations. It was found that the on-surface RBF configuration performed the best with the optimisation procedure significantly improving the efficiency of the design space coverage. This is selected here, however for practical optimisation manual placement may sometimes be preferable.

Figure 4. Initial control point positions for the standard on-surface (●) and standard off-surface (■) configurations with 10 design variables.

V. Results

A. Large Database Tests

Figures 5-7 show the results of the large database tests on the UIUC and NACA aerofoil libraries plus the union of the two, the combined library. For each library the seven parameterisation methods have been used to approximate all of the aerofoils across a wide range of design variables. The results graphs then show, for each design variable interval, the percentage of the aerofoil library that has been successfully reconstructed to within Kulfan's geometric tolerance (equation 7). This gives an estimation of the coverage of the design space represented by that method.

Figure 5 shows the results when the parameterisation methods are tested on the UIUC library. It shows that, in these circumstances, the SVD method built from the UIUC library gives the most efficient coverage, closely followed by the SVD method built from the combined library. It also shows that the Bèzier Surface method with 3 rows and the SVD method built from the NACA aerofoils perform the worst. However, discounting these two methods, it can be seen that all the results show a similar trend with roughly a five design variable gap in performance across the spectrum. Approximately 13-18 design variables are needed to approximate 80% of the library to the geometric tolerance or 20-25 for the full library. There also seems to be a distinction between the results from the constructive methods (CST, B-Spline and SVD) and the deformative methods (Bèzier Surface, RBF and Hicks-Henne) where the constructive methods seem to give slightly better coverage across the library, in particular for fewer design variables.

Figure 6 shows the results of the large database tests for the NACA library. It shows that the SVD method built from the NACA library gives by far the most efficient coverage under this configuration, successfully approximating 99% of the library for just 6 design variables. The success of this can be attributed to the SVD method constructing the design space such that it covers just the relatively narrow range represented by the NACA library. This is reinforced by the poor performance of this method in figure 5. It should however be noted that the NACA library was originally defined by varying just three parameters so six design variables does not represent the absolute best possible coverage (though may present the best linear reconstruction). The other parameterisation methods also perform better in general, with a particular improvement for the lower design variable cases. For example between 8 and 10 design variables are needed to reconstruct 50% of the NACA library successfully opposed to 11-15 for the UIUC. This is most likely due to the relative simplicity of the NACA library; the deformative methods may also be positively influenced by using the NACA0012 as the starting aerofoil.

The large database results for the full set of 2174 available aerofoils is shown in figure 7. Here, in agreement with the other two libraries, the SVD method built from the testing library performs best, this is followed by the other constructive methods (excluding the SVD built from NACA case) and then the deformative methods. It should be noted that the steep rise for the SVD built from NACA case with between 3-6 design variables represents the approximation the NACA library aerofoils ($\approx 40\%$ of the combined library) and coincides with figure 6.

Figures 8 (a)-(f) show how the large database results (tested on the combined library) are affected by changing the

Figure 5. Large database results for all the parameterisation methods when tested on the UIUC library.

Figure 6. Large database results for all the parameterisation methods when tested on the NACA library.

Figure 7. Large database results for all the parameterisation methods when tested on both libraries.

error tolerance. The weighted z tolerance values used are scaled such that errors in the leading 20% of the chord are doubled to coincide with Kulfan's geometric tolerance (equation 7). The tolerance used for the previous tests therefore represents a weighted z tolerance of 8×10^{-4} , this is presented as a solid black line. Each colour band then represents a change in the tolerance by a factor of two.

For all the methods, as the tolerance is decreased from the original geometric tolerance (in black) the shape of the contours stay reasonably consistent, with between 5 and 10 extra design variables initially needed to decrease the maximum weighted errors by half. This behaviour then drops off as tolerance continues to decrease, with the gradient of the contours gradually reducing. This represents the methods struggling to accurately reconstruct the more difficult aerofoils in the library to a tighter tolerance. The rate at which the contours change shape gives a good indication as to how the performance of each method translates when the required level of accuracy is increased.

The SVD method (figure 8f) shows the best performance for the lower tolerance bounds, with about 65% of the UIUC library successfully reconstructed to with a weighted z tolerance of 3.1×10^{-6} for 80 design variables. In contrast, the other methods approximate between 10-20% of the library under similar conditions. The B-Spline, CST, Bèzier surface and RBF methods (figures 8 (a)-(f)) show broadly comparable results through the design variable spectrum, all struggling to match many aerofoils to the lowest tolerance bound even for 100 design variables (despite still showing steady improvements for higher tolerances). The Hicks-Henne method (figure 8e) shows similar behaviour up until approximately 70 design variables where it struggles to significantly improve results for increased fidelity.

Figure 9 then shows similar results for the SVD method built from the NACA library tested on just the NACA library. It can be seen that the accurate approximation of the library at original tolerance tolerance shown in figure 6 is continued as the weighted z tolerance is decreased. In particular 99% of the library can be approximated to the lowest tolerance of 7.8×10^{-7} for just 34 design variables.

B. Case Studies

In addition to the large database tests a series of five case study aerofoils have been investigated: RAE 2822, NACA 4412, ONERA M6, NLR 7301 and NACA 66(3)-418. For each of these aerofoils a full range of reconstructions have been calculated using the same method as for the large database tests i.e. least squares approximation using each parameterisation methods for between 1 and 100 design variables. For these case studies however, only the best implementation of each parameterisation method is used. An aerodynamic solution has then been calculated for each reconstruction, as well as for the exact aerofoil, using the Euler solver SU^2 ⁴⁴. Each simulation was run on a 513×257 O-mesh, generated using a conformal mapping approach where all surface cells have aspect ratio one with a farfield distance of 50 chord lengths. The flow conditions used, as well as aerodynamic loads calculated on the exact aerofoils are displayed in table 1. These results enabled the convergence of the lift and drag coefficient (C_L and C_D) with respect

Figure 8. Percentage of combined aerfoil database approximated to varying weighted geometric tolerances.

Figure 9. Percentage of NACA aerofoil database approximated to varying weighted geometric tolerances for the SVD method with modes built from the NACA library.

to the geometric error and number of design variables to be analysed. A tolerance of one drag count (10^{-4}) and one lift count (10^{-3}) is used throughout this analysis to represent a benchmark for convergence. The geometric accuracy of the approximation is again presented in the form of the ‘max weighted z error’ (equation 8).

Aerofoil	M_∞	α ($^\circ$)	C_L (counts)	C_D (counts)
RAE2822	0.7	3.2	1029.7	89.0
NACA4412	0.7	0.0	766.5	142.4
ONERA M6 (D section)	0.7	3.0	538.2	90.2
NLR 7301	0.7	1.0	770.8	53.3
NACA 66(3)-418	0.7	0.0	544.7	181.3

Table 1. Case study flow conditions and aerodynamic loads.

For the RAE2822 and NACA 66(3)-418 case studies full convergence results are presented in figures 10 and 11^b. For both of these cases it can be seen that the lift and drag coefficients converge to the target value for all the parameterisation methods. It can however be seen that it takes many more design variables to do this for the NACA 66(3)-418. For example the RAE 2822 case requires between 28 and 44 design variables to converge both the force coefficients to within the one count tolerance whereas it requires between 39 and 72 to do this for NACA 66(3)-418. An additional star (★) symbol is also presented on the aerodynamic convergence figures; this represents the first point at which Kulfan’s geometric tolerance is satisfied. For both cases it can be seen that the aerodynamic properties have not sufficiently converged for any of parameterisation methods at this point; this is an indication that for these cases Kulfan’s geometric tolerance is not tight enough to ensure aerodynamic convergence. In fact, in both instances aerodynamic convergence to within a lift or drag count requires a geometric error that is more than an order of magnitude lower than Kulfan’s geometric tolerance.

A summary of the number of design variables required to converge the lift and drag to a single count for the five case studies is presented in table 2. For the different parameterisation methods it was found that the SVD method converged both of the aerodynamic coefficients for the fewest design variables; 40 on average. The CST and B-Spline methods required 42, followed by the RBF method with 44 and the Bezier Surface method with 50, though this is significantly influenced by difficulties in converging the lift for the NLR 7301. The Hicks-Henne method then represented the worst aerodynamic performance, requiring on average 60 design variables to converge both the lift and drag to the one count tolerance. In addition, it can be seen that for NACA4412, NLR 7301 and NACA 66(3)-418 cases it required significantly more design variables than any of the other methods. Interestingly across all the cases it can be seen that the Hicks-Henne method gives a very similar geometry error to the RBF method (figures 10c, 11c and supplementary figures S1c, S2c and S3c), yet performs significantly worse aerodynamically.

^bFull results for the NACA 4412, ONERA M6 and NLR 7301 can be found in supplementary figures S1 - S3

a) Lift Coefficient

b) Drag Coefficient

c) Max Weighted z Error

Figure 10. Convergence of aerodynamic and geometric properties for increasing design variables for each parameterisation method for the RAE 2822 case study; ★ represents first point Kulfan's geometry tolerance is satisfied.

a) Lift Coefficient

b) Drag Coefficient

c) Max Weighted z Error

Figure 11. Convergence of aerodynamic and geometric properties for increasing design variables for each parameterisation method for the NACA 66(3)-418 case study; ★ represents first point Kulfan's geometry tolerance is satisfied.

	RAE 2822		NACA4412		ONERA M6		NLR 7301		NACA 66(3)-418	
	Lift	Drag	Lift	Drag	Lift	Drag	Lift	Drag	Lift	Drag
Bezier Surface (4 row)	36	32	44	32	28	32	92	56	44	32
CST w/ LEM	28	26	34	28	24	40	58	44	50	52
Hicks Henne	42	30	68	52	26	36	82	58	72	68
B-Spline (cos)	36	32	40	34	24	32	58	56	44	42
RBF on surface (opt)	44	32	40	22	16	32	44	50	54	32
SVD (built from combined)	30	24	25	16	45	36	63	46	39	37

Table 2. Number of design variables required to converge lift and drag to within one count (10^{-3} for lift, 10^{-4} for drag) of the exact aerofoil.

Comparing the different cases it can be seen that for the RAE 2822, NACA4412 and ONERA M6, the number of design variables required for convergence is broadly similar (between 25 and 45, omitting the Hicks-Henne results); for the NLR 7301 and NACA 66(3)-418 however, more are required.

Finally the relationship between the aerodynamic convergence and geometric error was also investigated. This compared the convergence of lift and drag against the max weighted z error for all 1600 case study CFD simulations. The convergence of C_L is represented as

$$C_L^{conv} = \max_{DV \geq DV^{current}} |C_L^{DV} - C_L^{target}| \quad (26)$$

i.e. the maximum C_L error for equal or more design variables for a given parameterisation and test case; an analogous definition also holds for C_D .

It was found that a strong positive linear correlation exists between the convergence of the aerodynamic forces and the geometric error where C_D^{conv} and $C_L^{conv}/10$ are of equivalent order to the max weighted z error (full dataset is presented in supplementary figure S4). This implies that geometry error may indeed be a suitable metric for approximating aerodynamic error in Euler flow. It was however found that Kulfan's geometric tolerance may only indicate aerodynamic convergence to within approximately 10 counts. Consequently a reduced tolerance of 5×10^{-5} is required to ensure convergence of the lift and drag to the single count benchmark proposed. This reduction does however represent a 16 fold increase in accuracy which will require significantly more design variables to achieve. The dashed lines on figures 8 and 9 show the influence of using this 'reduced tolerance' on the large database tests. It shows that to approximate 80% of the combined aerofoil library with this tolerance between 38 and 66 design variable are required, compared to 13-18 for Kulfan's. This is approximately three times the number of design variables.

VI. Conclusions

In this work a detailed comparison of seven aerofoil parameterisation methods has been presented. The large database tests provided a good overview of the general design space coverage achieved by the shape parameterisation methods. It has been shown that to cover the full set of 2174 aerofoils to Kulfan's geometric tolerance between 20 and 25 design variables were required or to cover 80%, between 13 and 18 were required. Of all the methods the SVD method was found to give the most efficient coverage of the aerofoil design space.

The effect of changing the success tolerance was investigated and, in general, it was found that as the tolerance was reduced, an increasing number of design variables were needed to increase the design space coverage. For example, to increase the design space coverage from 80% to 90% using Kulfan's geometric tolerance approximately two additional design variables were needed, though for a weighted tolerance of 5×10^{-5} this increased to as many as 40 design variables. This degradation in performance however varied between the methods, with SVD method showing the least degradation and the Hicks-Henne the worst. As a result of this the SVD method shows significantly better design space coverage than the other methods for lower tolerances.

Investigation of specific transonic aerofoils case studies helped assess the link between the geometric convergence, driven by the increase in design variables, and the convergence of the aerodynamic properties. These results show that the convergence of the geometric position correlates well with the convergence of the aerodynamic properties. They decrease at a very similar rate such that an order of magnitude decrease in geometry error should correspond to an order of magnitude reduction in the aerodynamic error. It was found that Kulfan's geometric tolerance ensured aerodynamic convergence to approximately 10 lift or drag counts and to reduce this to a single count a max weighted z error of approximately 5×10^{-5} was required. It was also shown that to cover 80% of the full design space to this tolerance requires between 38 and 66 design variables. This indicates that a significant number of design variables

may be required in order to cover the aerofoil design space to a fine aerodynamic accuracy.

VII. Acknowledgements

The authors 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.

References

- [1] 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, doi: 10.2514/6.2014-0568.
- [2] 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, doi: 10.2514/6.2014-0413.
- [3] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., "Impact of Shape Parameterisation on Aerodynamic Optimisation of Benchmark Problems," *54rd AIAA Aerospace Sciences Meeting*, Jan 2016.
- [4] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., "Progressive Subdivision Curves for Aerodynamic Shape Optimisation," *54rd AIAA Aerospace Sciences Meeting*, Jan 2016.
- [5] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., "A Locally Adaptive Subdivision Parameterisation Scheme for Aerodynamic Shape Optimisation," *34th AIAA Applied Aerodynamics Conference*, 2016.
- [6] Sobieczky, H., "Parametric airfoils and wings," *Recent Development of Aerodynamic Design Methodologies*, Springer, 1999, pp. 71–87.
- [7] Braibant, V. and Fleury, C., "Shape optimal design using B-splines," *Computer Methods in Applied Mechanics and Engineering*, Vol. 44(3), 1984, pp. 247–267, doi: 10.1016/0045-7825(84)90132-4.
- [8] Sobieczky, H., "Geometry generator for CFD and applied aerodynamics," *Courses and Lectures-International Centre for Mechanical Sciences*, 1997, pp. 137–158.
- [9] Anderson, W. K., Karman, S. L., and Burdyslaw, C., "Geometry parameterization method for multidisciplinary applications," *AIAA Journal*, Vol. 47, 2009, pp. 168–1578, doi: 10.2514/1.41101.
- [10] Kulfan, B. M. and Bussioletti, J. E., "Fundamental parametric geometry representations for aircraft component shapes," *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, September 2006.
- [11] Kulfan, B. M., "Universal parametric geometry representation method," *Journal of Aircraft*, Vol. 45, No. 1, 2008, pp. 142–158, doi: 10.2514/1.29958.
- [12] Jameson, A., "Aerodynamic design via control theory," *Journal of Scientific Computing, also ICASE Report No.88-64*, Vol. 3, 1988, pp. 233–260, doi: 10.1007/BF01061285.
- [13] Hicks, R. M. and Henne, P. A., "Wing design by numerical optimization," *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412, doi: 10.2514/3.58379.
- [14] Pickett, J. R., Rubinstein, M., and Nelson, R., "Automated structural synthesis using a reduced number of design coordinates," *14th Structures, Structural Dynamics, and Materials Conference*, 1973.
- [15] 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, doi: 10.1002/nme.2219.
- [16] Chauhan, D., Praveen, C., and Duvigneau, R., "Wing shape optimization using FFD and twist parameterization," *12th Aerospace Society of India CFD Symposium*, 2010.
- [17] Jameson, A., "Advances in aerodynamic shape optimization," *Computational Fluid Dynamics 2004*, Springer, 2006, pp. 687–698.
- [18] 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.
- [19] Kim, S., Alonso, J., and Jameson, A., "Two-dimensional high-lift aerodynamic optimization using the continuous adjoint method," *Multidisciplinary Analysis Optimization Conferences*, American Institute of Aeronautics and Astronautics, Sept. 2000.

- [20] Bogue, D. and Crist, N., “CST Transonic Optimization using Tranair++,” 2008.
- [21] Lane, K. and Marshall, D., “A Surface Parameterization Method for Airfoil Optimization and High Lift 2D Geometries Utilizing the CST Methodology,” *Aerospace Sciences Meetings*, American Institute of Aeronautics and Astronautics, Jan. 2009.
- [22] Ceze, M., Hayashi, M., and Volpe, E., “A study of the CST parameterization characteristics,” 2009.
- [23] Lane, K. A. and Marshall, D. D., “Inverse Airfoil Design Utilizing CST Parameterization,” *48th AIAA Aerospace Sciences Meeting*, 2010.
- [24] Hewitt, P. and Marques, S., “Aerofoil Optimisation Using CST Parameterisation in SU2,” *Royal Aeronautical Society Biennial Applied Aerodynamics Research Conference, 2014*, 2014.
- [25] Bisson, F. and Nadarajah, S., “Adjoint-Based Aerodynamic Optimization of Benchmark Problems,” *52nd Aerospace Sciences Meeting*, Jan 2014.
- [26] 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.
- [27] Song, W. and Keane, A. J., “A study of shape parameterisation methods for airfoil optimisation,” *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2004, pp. 2031–2038.
- [28] Sherar, P., Thompson, C., Xu, B., and Zhong, B., “An Optimization Method Based On B-spline Shape Functions & the Knot Insertion Algorithm,” *World congress on engineering*, Citeseer, 2007, pp. 862–866.
- [29] Lépine, J., Trépanier, J.-Y., and Pépin, F., “Wing aerodynamic design using an optimized NURBS geometrical representation,” *The 38th AIAA Aerospace Sciences Meeting and Exhibit*, No. AIAA-2000-699, Reno, Nevada, 2000.
- [30] 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, doi: 10.2514/1.41420.
- [31] 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.
- [32] 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, doi: 10.2514/1.J053427.
- [33] Sóbester, A. and Forrester, A. I., *Aircraft Aerodynamic Design: Geometry and Optimization*, John Wiley & Sons, 2014.
- [34] 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, doi: 10.1002/flid.1769.
- [35] Morris, A., Allen, C., and S. Rendall, T., “Domain-Element Method for Aerodynamic Shape Optimization Applied to Modern Transport Wing,” *AIAA journal*, Vol. 47, No. 7, 2009, pp. 1647–1659.
- [36] Allen, C. B. and Rendall, T. C. S., “CFD-based optimization of hovering rotors using radial basis functions for shape parameterization and mesh deformation,” *Optimization and Engineering*, Vol. 14, No. 1, Mar 2013, pp. 97–118, doi: 10.1007/s11081-011-9179-6.
- [37] Castonguay, P. and Nadarajah, S. K., “Effect of shape parameterization on aerodynamic shape optimization,” *45th AIAA Aerospace Sciences Meeting and Exhibit*, 2007.
- [38] 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.
- [39] Sripawadkul, V., Padulo, M., and Guenov, M., “A comparison of airfoil shape parameterization techniques for early design optimization,” *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, 2010.
- [40] Selvan, K. M., “On The Effect of Shape Parameterization on Aerofoil Shape Optimization,” *International Journal of Research in Engineering and Technology*, 2015, doi: 10.15623/ijret.2015.0402016.
- [41] Freestone, M., “VGK method for two-dimensional aerofoil sections Part 1: principles and results,” April 2004, ESDU 96028.
- [42] de Boor, C., *A Practical Guide to Splines*, Vol. 27 of *Applied Mathematical Sciences*, Springer-Verlag New York, 1978.

- [43] 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, doi: 10.2514/6.2015-0761.
- [44] 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 (SU²): An open-source integrated computational environment for multi-physics simulation and design,” *Aerospace Sciences Meetings*, American Institute of Aeronautics and Astronautics, Jan. 2013.
- [45] Piegl, L. A. and Tiller, W., *The NURBS book*, Springer, 1997.
- [46] Kulfan, B. M., “Modification of CST airfoil representation methodology,” Retrieved from <http://www.brendakulfan.com/docs/CST8.pdf>.
- [47] Khurana, M., Winarto, H., and Sinha, A., “Airfoil geometry parameterization through shape optimizer and computational fluid dynamics,” *46th AIAA Aerospace Sciences Meeting and Exhibit*, 2008.
- [48] Sederberg, T. W. and Parry, S. R., “Free-form deformation of solid geometric models,” *ACM Siggraph Computer Graphics*, Vol. 20, ACM, 1986, pp. 151–160.
- [49] Sederberg, T. W., “Computer Aided Geometric Designs,” Published online: <http://tom.cs.byu.edu/~557/text/cagd.pdf>, Oct 2014.
- [50] Wendland, H., *Scattered data approximation*, Cambridge University Press Cambridge, 2005.
- [51] Buhmann, M. D., “Radial basis functions,” *Acta numerica*, Vol. 9, 2000.
- [52] 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, doi: 10.2514/6.2015-0762.
- [53] Abbott, I. and Von Doenhoff, A., *Theory of Wing Sections*, Dover Publications, Inc., 1949.

A. Geometry Normalisation

All aerofoils in this work have been normalised such that the leading edge is situated at [0,0] and the trailing edge is defined symmetrically about the point [1,0], and have been discretized by 301 half cosine squared distributed points for $x \in [0, 1]$. To ensure this the leading edge has been defined as the point situated furthest away from the trailing edge. This normalisation procedure satisfies the following equations:

$$X_i = (x_i, z_i), \quad \text{for } i = 1 : 301 \quad (27)$$

$$X_{le} = X_{151} = (0, 0), \quad (28)$$

$$X_{te} = (X_1 + X_{301})/2 = (1, 0), \quad (29)$$

$$x_1 = x_{301} = 1, \quad (30)$$

$$z_1 = -z_{301} = -z_{te} \leq 0, \quad (31)$$

$$x_i = \left(1 - \cos\left(\frac{(i - 151)\pi}{150}\right)\right)^2 / 4, \quad (32)$$

$$\|X_{le} - X_{te}\| > \|X_i - X_{te}\| \quad \forall i \neq le, \quad (33)$$

where X_i denotes the i th point of the aerofoil and the subscripts le and te denote the leading and trailing edges, respectively.

B. NACA 4-series Definition

The NACA 4-series aerofoils used throughout this paper are defined by

$$x_{upper} = x - z_t \sin \theta, \quad z_{upper} = z_c + z_t \cos \theta, \quad (34)$$

$$x_{lower} = x + z_t \sin \theta, \quad z_{lower} = z_c - z_t \cos \theta, \quad (35)$$

where

$$z_t = 500t^* \left(0.2969\sqrt{x} - 0.126x - 0.3516x^2 + 0.2843x^3 - 0.1036x^4\right), \quad (36)$$

$$z_c = \begin{cases} 100m^* \frac{x}{p'^2} (2p' - x) & \text{for } 0 \leq x \leq p' \\ 100m^* \frac{1-x}{(1-p')^2} (1 + x - 2p') & \text{for } p' \leq x \leq 1 \end{cases} \quad (37)$$

and

$$\theta = \arctan\left(\frac{dz_c}{dx}\right) \quad (38)$$

for given values of the maximum thickness parameter t^* , maximum camber parameter m^* , and maximum camber location $p^* = 10p'$. This definition is equivalent to the original definition⁵³ apart from a common modification to the x^4 coefficient has been applied to ensure a sharp trailing edge thickness.