

Design Space Dimensionality Reduction through Physics-based Geometry Re-Parameterization

András Sóbester · Stephen Powell

Received: date / Accepted: date

Abstract The effective control of the extent of the design space is the *sine qua non* of successful geometry-based optimization. Generous bounds run the risk of including physically and/or geometrically nonsensical regions, where much search time may be wasted, while excessively strict bounds will often exclude potentially promising regions. A related ogre is the pernicious increase in the number of design variables, driven by a desire for geometry flexibility – this can, once again, make design search a prohibitively time-consuming exercise. Here we discuss an instance-based alternative, where the design space is defined in terms of a set of representative bases (design instances), which are then transformed, via a concise, parametric mapping into a new, generic geometry. We demonstrate this approach via the specific example of the design of supercritical wing sections. We construct the mapping on the generic template of the Kulfan class-shape function transformation and we show how patterns in the coefficients of this transformation can be exploited to capture, within the parametric mapping, some of the physics of the design problem.

Keywords geometry modeling, shape description, design optimization, parametric geometry, surrogate modeling, kriging

1 Instance-Based Design Space Definition

The recent history of design optimization is characterized by an ‘arms race’ between the rapid increase in affordable computing power and a demand for increasing fidelity in the physics-based simulations such design exercises are

A. Sóbester, S. R. Powell
University of Southampton
Faculty of Engineering and the Environment
Tel.: +44 23 8059 2350
Fax: +44 23 8059 4813
E-mail: a.sobester@soton.ac.uk, s.r.powell@soton.ac.uk

based on. The fundamental constraint determining the feasibility of design searches based on computer simulations is thus still the required number of these analysis runs, which depends chiefly on the number of design variables. In fact, the cost of exploring a design space increases exponentially with the number of parameters that the objective function depends on. This *curse of dimensionality* is particularly pressing in the context of preliminary design, where the desire to explore a wide range of configurations may tempt the engineer into equipping the parametric geometry with numerous degrees of freedom. These often have a drastic effect on the complexity of the design problem. Coupled with broad ranges from which they may take values, they risk restricting any reasonable MDO (Multidisciplinary Design Optimization) process to merely scratching the surface of an unnecessarily inflated design space. The desire for new design variables with broad ranges driven by the need for flexibility must therefore be tempered by an understanding of *necessary* flexibility.

Simple parametric geometries sometimes permit design space size control via a simple adjustment of the ranges of their design variables. Increasing dimensionality and the almost inevitable accompanying increase in the complexity of variable interactions tends, however, to preclude a truly effective implementation of this straightforward approach. The resulting design space then is likely to either include regions populated by physically or even geometrically nonsensical designs (costly time wasting from an MDO point of view) or to be too restricted to yield significant performance gains.

Niche alternatives exist. For instance, the structural optimization community sees much potential in doing away with the conventional concept of design variables altogether, in favor of non-parametric, topological heuristics, typically driven by the iterative elimination of under-utilized sub-domains within the geometry [see Rozvany (2009) for a recent review of the strengths, weaknesses, future hopes and past false promises of this class of techniques]. The area of application that this study focuses on, aerodynamic shape optimization, has non-parametric approaches too. By far the most prominent amongst these is the class of methods based on a calculus of variations standpoint, where densely discretized surfaces are allowed to vary as driven by the gradients of a chosen objective. Rooted in control theory, these so-called *adjoint* methods [pioneered in an aerospace design setting by Jameson (1988)] have seen successful applications in custom-built, local search frameworks. Nevertheless, such methods are, at present, confined to the realm of relatively specialized applications. The most widely adopted schemes are still those based on an explicit choice of the (often numerous) design variables and their ranges, a process strongly intertwined with the construction of the geometry itself.

Here we advocate a substantively different approach, aimed at tackling both the variable number and the range problem, based on the following observation. Few engineers are equipped with the ability to construct the most parsimonious geometry conceivable for a given design study *and* to place appropriate bounds on the sets of design variables that define it. However, most can readily construct specific representative instances (designs) that can be

viewed as bases of a tentative design space. A generic geometry can then be built in the form of a parametric mapping between these bases and the final set of coordinates representing the new shape.

In an earlier paper [Sóbester (2009a)] we demonstrated this philosophy by constructing a Non-Uniform Rational B-Spline (NURBS) geometry, whose parametric mapping was determined by a set of bias variables, which simply positioned it in the design space with respect to a set of bases. The latter had been selected for their ‘representativeness’, that is, their ability to express, via the chosen mapping, other potential bases, which would then, of course, become redundant. Here we follow the same basic philosophy of seeking a parametric mapping between fixed bases and a new geometry, but along an entirely different route, one that enables us to construct a mapping that captures some of the physics behind the design problem too [by comparison to the purely geometrical reasoning employed in Sóbester (2009a)]. The template we use to construct this mapping is the class- and shape function transformation of Kulfan (2008) (the details of which we shall discuss in more detail in due course).

Perhaps the most germane format for the detailed presentation of the philosophy outlined above is through a specific design problem. We shall consider the case of supercritical airfoils (i.e., transonic wing sections) for transport aircraft. We endeavor to demonstrate through this example how the Kulfan transformation can be used to exploit certain ‘family traits’ amongst our chosen set of basis shapes with the ultimate goal of constructing a low dimensionality mapping.

2 An Application: Parametric Airfoils in Preliminary Design

Few would dispute that if a design brief calls for a long range airliner with a cruise Mach number of 0.8, there is little point in equipping the parametric airfoil with degrees of freedom that will enable it to reproduce, say, highly cambered sections. There is, however, a school of thought according to which it is worth adding more flexibility to a scheme that can produce suitable (in this example, supercritical) shapes (say, by inserting additional control points into a NURBS airfoil), because the new scheme will no doubt be capable of producing *additional* suitable shapes, as well as clearly inappropriate ones, which are merely seen as a byproduct of the process. Our thesis here is two-fold. On the one hand, as we hinted earlier, this is a very expensive byproduct: an automated optimization process will not ‘know’ that there is no point in running the expensive numerical multidisciplinary analysis over something that an aerodynamicist would recognize as an inappropriate, low Reynolds/Mach number section (or worse still, a completely nonsensical one) when looking for a Mach 0.8 design – therefore many evaluations may get wasted. On the other hand, once the global search is complete (on the very concise airfoil),

there is still scope for a *local* search in the vicinity of the optimum via a re-parameterized or even non-parametric model¹.

Of course, the initial, parsimonious parameterization has to be flexible enough to enable a meaningful global search. Here we argue that we can achieve this by choosing a diverse set of existing, suitable ‘training’ geometries as bases, which will also serve to limit the design space to a problem-specific ‘sensible’ region. Specifically, we shall use the SC(2) series of supercritical airfoils Harris (1990) (more on the design of which later).

The idea of exploiting the features of a well-established family of airfoils by blending them into a parametric representation is not without precedent. In fact, the orthogonal basis functions introduced by Robinson and Keane (2001) are based on the very same class of shapes we are using here: SC(2), the second generation of NASA supercritical airfoils².

Here, following on from a formulation introduced in Sóbester (2009b), we describe a recipe for building a very concise model by capturing the shapes of the members of the SC(2) family through a highly flexible approximation model (Kulfan’s class-shape function transformation – see Section III) and, by exploiting family-specific patterns in the variables of these approximations (Sections V and VI), establishing a parametric mapping between these and a new, generic shape. We show how, beyond the dimensionality reduction and implicit domain size control, as an additional benefit, the parameters of the concise airfoil can be chosen such that they are linked to the known physical properties of the members of the family. We conclude the study by reflecting on the place of these findings in the context of the overall aerodynamic design process (Section VII) and on possible future developments (Section VIII).

3 Applying the Kulfan (Class-Shape Function) Transformation to Airfoil Shapes

In what follows we shall use a coordinate system whose x axis is aligned with the chord, with the leading edge point in the origin and the trailing edge point(s) at $x = 1$. We define a universal approximation to any airfoil in the xOz plane as a pair of explicit curves $\mathcal{A} = [z^u(x, \dots), z^l(x, \dots)]$, where $x \in [0, 1]$ and the superscripts u and l distinguish between the upper and the lower surface (here and on all the symbols in the following discussion) and the dots indicate that the shape of the two curves depends on a number of parameters. \mathcal{A} becomes the approximation to a target airfoil if we determine these parameters such that they minimize some metric of difference (say, mean squared error) between \mathcal{A} and the target.

¹ We reviewed some possible schemes for such local improvement in Sóbester (2009a) – one example is the already mentioned mesh-based formulation of Jameson (1988) designed specifically for local optimization guided by adjoint flow solutions.

² See Vanderplaats (1979, 1984); Collins and Saunders (1997) for further instances of parameterisation using basis airfoils.

Here we adopt the class-shape transformation of Kulfan (2008) as the universal approximation. The main traits that make this scheme attractive for our purposes are its ability a) to approximate practically any airfoil (flexibility) and b) to require a relatively small number of design variables do so with high accuracy (conciseness) – see Kulfan (2006) for the empirical and analytical underpinning of this.

Let the generic airfoil be defined as

$$\begin{aligned} \mathcal{A}(\mathbf{V}) &= \mathcal{A}[x, v_0^u, v_1^u, \dots, v_{n_{\text{BP}}^u}^u, z_{\text{TE}}^u, v_{\text{LE}}^u, v_0^l, v_1^l, \dots, v_{n_{\text{BP}}^l}^l, z_{\text{TE}}^l, v_{\text{LE}}^l] = \\ &= [z^u(x, v_0^u, v_1^u, \dots, v_{n_{\text{BP}}^u}^u, z_{\text{TE}}^u, v_{\text{LE}}^u), z^l(x, v_0^l, v_1^l, \dots, v_{n_{\text{BP}}^l}^l, z_{\text{TE}}^l, v_{\text{LE}}^l)], \end{aligned} \quad (1)$$

where n_{BP}^u and n_{BP}^l denote the orders of sets of Bernstein polynomials that control the shape of the two curves that make up the airfoil. The upper surface of the airfoil is defined as:

$$\begin{aligned} z^u(x, v_0^u, v_1^u, \dots, v_{n_{\text{BP}}^u}^u, z_{\text{TE}}^u, v_{\text{LE}}^u) &= \underbrace{\sqrt{x(1-x)}}_{\text{class function}} \underbrace{\sum_{r=0}^{n_{\text{BP}}^u} v_r^u C_{n_{\text{BP}}^u}^r x^r (1-x)^{n_{\text{BP}}^u-r}}_{\text{scaled Bernstein partition of unity}} + \\ &+ \underbrace{z_{\text{TE}}^u x}_{\text{trailing edge thickness term}} + \\ &+ \underbrace{x\sqrt{1-x} v_{\text{LE}}^u (1-x)^{n_{\text{BP}}^u}}_{\text{supplementary leading edge shaping term}}, \end{aligned} \quad (2)$$

where $C_{n_{\text{BP}}^u}^r = \frac{n_{\text{BP}}^u!}{r!(n_{\text{BP}}^u-r)!}$. A curve built upon the same template defines the lower surface:

$$\begin{aligned} z^l(x, v_0^l, v_1^l, \dots, v_{n_{\text{BP}}^l}^l, z_{\text{TE}}^l, v_{\text{LE}}^l) &= \underbrace{\sqrt{x(1-x)}}_{\text{class function}} \underbrace{\sum_{r=0}^{n_{\text{BP}}^l} v_r^l C_{n_{\text{BP}}^l}^r x^r (1-x)^{n_{\text{BP}}^l-r}}_{\text{scaled Bernstein partition of unity}} + \\ &+ \underbrace{z_{\text{TE}}^l x}_{\text{trailing edge thickness term}} + \\ &+ \underbrace{x\sqrt{1-x} v_{\text{LE}}^l (1-x)^{n_{\text{BP}}^l}}_{\text{supplementary leading edge shaping term}}. \end{aligned} \quad (3)$$

Approximating an arbitrary smooth airfoil with these expressions amounts to finding the vectors

$$\mathbf{v}^u = \left\{ \overbrace{v_0^u, v_1^u, \dots, v_{n_{\text{BP}}^u}^u, v_{\text{LE}}^u}^{n_{\text{BP}}^u+2 \text{ design variables to define upper surface}} \right\}^T \quad (4)$$

and

$$\mathbf{v}^l = \left\{ \overbrace{v_0^l, v_1^l, \dots, v_{n_{\text{BP}}^l}^l, v_{\text{LE}}^l}^{n_{\text{BP}}^l+2 \text{ design variables to define lower surface}} \right\}^T \quad (5)$$

(note that z_{TE}^u and z_{TE}^l are simply the trailing edge ordinates of the target airfoil so they are known) which, as indicated earlier, minimize some metric of the difference between $\mathcal{A}(\mathbf{V})$ and the target airfoil.

Let us consider, say, the upper surface of a target airfoil, given as a list of n_{T}^u coordinate pairs $\{(x_{\text{T}i}^u, z_{\text{T}i}^u) \mid i = \overline{1, n_{\text{T}}^u}\}$. We can exploit the linearity (in terms of the design variables) of the Kulfan approximation by re-arranging Equation (2) in matrix form, equating each of these target points with their approximations:

$$\mathbf{B}^u \cdot \mathbf{v}^u = \mathbf{z}^u, \quad (6)$$

where $\mathbf{z}^u = \{z_{\text{T}1}^u - z_{\text{TE}}^u x_{\text{T}1}^u, z_{\text{T}2}^u - z_{\text{TE}}^u x_{\text{T}2}^u, \dots, z_{\text{T}n_{\text{T}}^u}^u - z_{\text{TE}}^u x_{\text{T}n_{\text{T}}^u}^u\}^T$ and \mathbf{B}^u is an $n_{\text{T}}^u \times (n_{\text{BP}}^u + 2)$ matrix of the class-shape function transformation terms, comprising the Bernstein polynomials

$$\mathbf{B}_{p,q} = \sqrt{x_{\text{Tp}}^u(1-x_{\text{Tp}}^u)} C_{n_{\text{BP}}^u}^{q-1} x_{\text{Tp}}^u{}^{q-1} (1-x_{\text{Tp}}^u)^{n_{\text{BP}}^u-q+1}, \quad p = \overline{1, n_{\text{T}}^u}, \quad q = \overline{1, n_{\text{BP}}^u + 1} \quad (7)$$

and the leading edge shaping terms

$$\mathbf{B}_{p, n_{\text{BP}}^u+2} = x_{\text{Tp}}^u \sqrt{(1-x_{\text{Tp}}^u)} (1-x_{\text{Tp}}^u)^{n_{\text{BP}}^u}, \quad p = \overline{1, n_{\text{T}}^u}. \quad (8)$$

Computing $\mathbf{v}^u = \mathbf{B}^{u+} \mathbf{z}^u$ (where $\mathbf{B}^{u+} = (\mathbf{B}^{uT} \mathbf{B}^u)^{-1} \mathbf{B}^{uT}$ is the Moore-Penrose pseudo-inverse of \mathbf{B}^u) will now yield the set of coefficients that correspond to a least squares fit through the points of the target airfoil. Naturally, the same procedure can be repeated for the lower surface.

The accuracy of any such approximation can be improved by increasing the orders n_{BP}^u and n_{BP}^l of the Bernstein polynomials, thus adding more shaping terms [see Kulfan (2006) for experiments illustrating this on a range of airfoils]. Generally, few applications require orders greater than about seven or eight and in many cases fewer terms are needed to approximate the upper surface of a cambered airfoil than the lower.

In what follows, when referring to the class-shape function approximation of an airfoil, we shall add the name of that airfoil to the previously introduced notation as a subscript, preceded by a '~' symbol to indicate the inexact nature of the approximation. Thus, for example, we shall refer to the the class-shape approximation of the supercritical airfoil SC(2)-0612 as

$$\begin{aligned} \mathcal{A}_{\sim\text{SC}(2)-0612} &= \mathcal{A}(\mathbf{V}_{\sim\text{SC}(2)-0612}) = \\ &= \mathcal{A}[v_{0\sim\text{SC}(2)-0612}^u, v_{1\sim\text{SC}(2)-0612}^u, \dots, v_{\text{LE}\sim\text{SC}(2)-0612}^l]. \end{aligned} \quad (9)$$

Figure 1 depicts the terms of this approximation for $n_{\text{BP}}^u = 2$ and $n_{\text{BP}}^l = 3$. As per equations (2) and (3), the total number of degrees of freedom (design variables) for this approximation is $n_{\text{BP}}^u + 2 + n_{\text{BP}}^l + 2 = 9$.

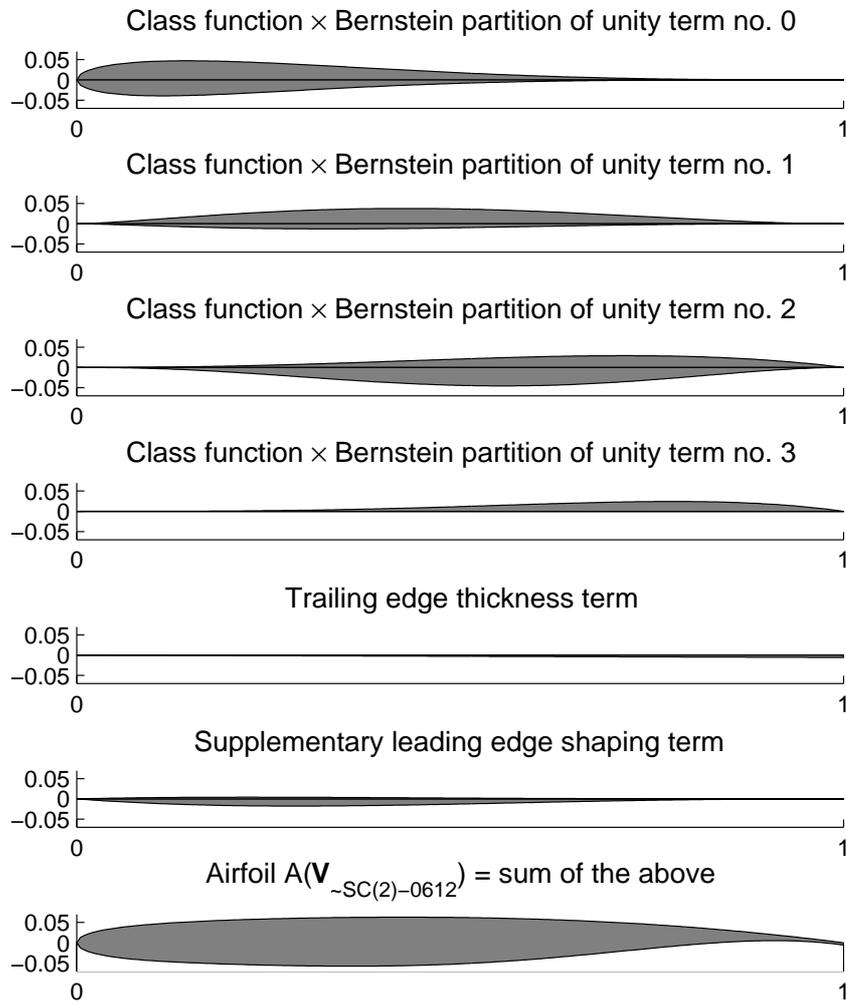


Fig. 1 The terms of equations (2) and (3) making up the class-shape approximation of the supercritical airfoil SC(2)-0612. Note that there is a single term no. 3, because we have used fewer polynomial terms to describe the upper surface. Somewhat counter-intuitively, the sole term present there is, in fact, a positive one, though it participates in the approximation of the negative, lower surface.

4 The SC(2) Family of Supercritical Airfoils – Origins and Analysis

The SC(2) Family of Supercritical Airfoils is the result of research conducted by NASA starting in the 1960s aimed at the development of “practical airfoils with two-dimensional transonic turbulent flow and improved drag divergence Mach numbers while retaining acceptable low-speed maximum lift and stall characteristics” Harris (1990). They trace their lineage back to the work of Whitcomb and Clark (1965), who noted that a three quarter chord slot between the upper and lower surfaces of a NACA 64A series airfoil gave it the ability to operate efficiently at Mach numbers greater than its original critical Mach number – hence the term ‘supercritical’, or ‘SC’ for short. The number in brackets following the ‘SC’ designation places each of these ‘family-related’ airfoils [to use the term coined by Harris (1990)] into one of three distinct phases of development through the 1970s and 1980s.

The fundamental design philosophy of the SC airfoils was to delay drag rise on the top surface through a reduction in curvature in the middle region, in order to reduce flow acceleration and thus reduce the local Mach number. This, in turn, reduces the severity of the adverse pressure gradient there and thus the associated shock is moved aft and is weakened. From a purely aerodynamic standpoint, the idea was to create a flat top pressure profile forward of the shock, obtained by balancing the expansion waves emanating from the leading edge, the compression waves resulting from their reflection off the sonic line (separating the subsonic and supersonic flow regions) back onto the surface and a second set of expansion waves associated with their reflection. Geometrically, this was achieved through a large leading edge radius (strong expansion waves) and a flat mid-chord region (reducing the accelerations that would have needed to be overcome by the reflected compression waves) Whitcomb (1974). The well-known lower surface aft-end ‘cusp’ of the SC class of airfoils is a result of efforts to increase circulation, which led to a relatively aggressive aft-loading on the airfoil, as well as to the attainment of the design lift coefficients at low angles of attack.

Of all the NASA SC airfoils the SC(2) series has shown the greatest longevity and it forms the focus of the present study. It comprises 21 airfoils of different thickness to chord ratios and design lift coefficients. The first two digits of the encoding of each airfoil represent the design lift coefficient (multiplied by ten), while the third and the fourth digit represent the maximum thickness to chord ratio (as a percentage). Thus, for instance, SC(2)-0714 is the 14% thick second series supercritical airfoil designed for a lift coefficient of 0.7.

5 Exploiting Shared Features

Let us consider six of the 21 members of the SC(2) family, all designed for transport aircraft: SC(2)-0410, SC(2)-0610, SC(2)-0710, SC(2)-0412, SC(2)-0612 and SC(2)-0712. Following the formulation described in Section III, we

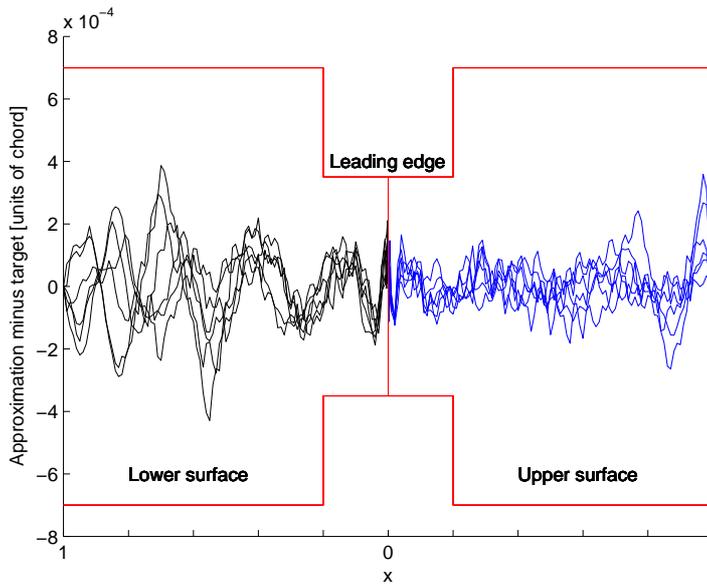


Fig. 2 Approximation errors: the differences between the six supercritical airfoils and their class-shape transformations. The horizontal lines indicate the typical tolerances of wind tunnel models (tighter within 20% chord of the leading edge).

approximate these airfoils using the class-shape function transformation based on the Bernstein partitions of unity³ of orders $n_{\text{BP}}^u = 5$ and $n_{\text{BP}}^l = 5$. This means that we have to find the 12 polynomial coefficients (6 for each surface) plus an additional leading edge shaping term for each surface, that minimizes the difference (mean squared error) between the approximation and the target. We take the trailing edge parameters z_{TE}^u and z_{TE}^l to be equal to the trailing edge thicknesses of the given target airfoil, so, of the total of 16 approximation parameters, we are left with 14 to be determined.

Figure 2 illustrates the accuracy of the approximations we have found, indicating that the approximation errors are well within the typical tolerances of wind tunnel models ($\pm 3.5 \times 10^{-4}$ units of chord within 20% of the leading edge and $\pm 7 \times 10^{-4}$ elsewhere [Kulfan and Bussoletti (2006)]).

We thus have a 16 dimensional design space inhabited by six designs with, as yet, no obvious connection between them. For a parameterization that is more useful from a preliminary design perspective, we now seek to construct a reduced dimensionality space, which we can map back into this original domain, or, more accurately, into the sub-domain delimited by the six examples. One way of achieving this is to identify common features the members of this family of six sections share.

³ So called because the terms of the series add up to one, regardless of the order n_{BP} .

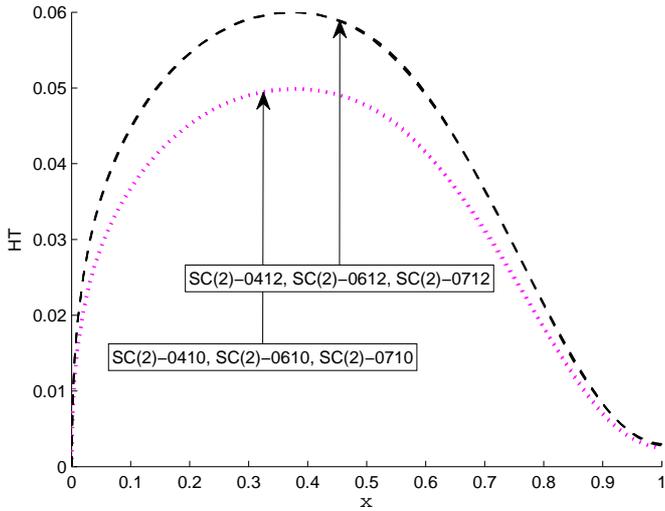


Fig. 3 The half-thickness distributions of the six SC(2) airfoils (vertical and horizontal axes are to different scales).

5.1 Divide and Conquer

Consider the thickness distributions of our six chosen airfoils. As seen in Figure 3, the airfoils with the same maximum thickness to chord ratios share, in fact, their entire thickness distributions. Thus, the different design lift coefficients are purely down to the different camber curve shapes (Figure 4) and this is good news from the perspective of mapping to a more concise description. We can apply the *divide and conquer* principle by separating, in terms of the transformation coefficients, the effects of the two features that headline each of the SC(2) airfoils, design lift coefficient (clearly determined by the shape of the camber curve) and maximum thickness to chord ratio (determined by the thickness distribution). It also gives us a strong indication that, of all the possible variables we could use, it makes most sense to define the new, concise design space in terms of maximum thickness to chord ratio and design lift coefficient – denoted as t/c and c_l respectively in what follows. The mapping we seek is therefore the first of the sequence

$$(t/c, c_l) \mapsto (v_0^u, v_1^u, \dots, v_{LE}^1) \mapsto [z^u(x), z^1(x)], \quad (10)$$

where we already have the second step in the shape of equations (2) and (3).

Turning now our attention to separating the airfoil into a camber line and a thickness distribution, the class-shape transformation gives us a compact way of writing these – manipulating equations (2) and (3) we get

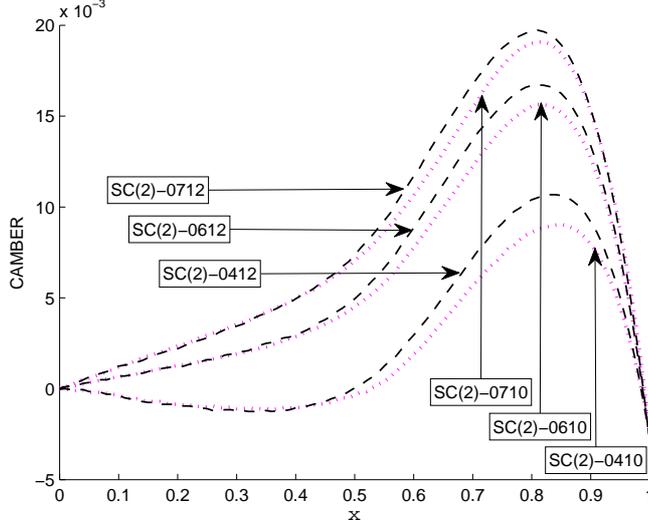


Fig. 4 The camber curves of the six SC(2) airfoils (axes to different scales).

$$\begin{aligned}
 \text{CAMBER}(x, v_0^u, \dots, v_{\text{TE}}^u, v_0^l, \dots, v_{\text{TE}}^l) = & \underbrace{\sqrt{x(1-x)}}_{\text{class function}} \underbrace{\sum_{r=0}^{n_{\text{BP}}} \frac{v_r^u + v_r^l}{2} C_{n_{\text{BP}}}^r x^r (1-x)^{n_{\text{BP}}-r}}_{\text{scaled Bernstein partition of unity}} + \\
 & + \underbrace{\frac{z_{\text{TE}}^u + z_{\text{TE}}^l}{2} x}_{\text{trailing edge thickness term}} + \quad (11) \\
 & + \underbrace{x\sqrt{1-x} \frac{v_{\text{TE}}^u + v_{\text{TE}}^l}{2} (1-x)^{n_{\text{BP}}}}_{\text{supplementary leading edge shaping term}},
 \end{aligned}$$

and

$$\begin{aligned}
 \text{HT}(x, v_0^u, \dots, v_{\text{TE}}^u, v_0^l, \dots, v_{\text{TE}}^l) = & \underbrace{\sqrt{x(1-x)}}_{\text{class function}} \underbrace{\sum_{r=0}^{n_{\text{BP}}} \frac{v_r^u - v_r^l}{2} C_{n_{\text{BP}}}^r x^r (1-x)^{n_{\text{BP}}-r}}_{\text{scaled Bernstein partition of unity}} + \\
 & + \underbrace{\frac{z_{\text{TE}}^u - z_{\text{TE}}^l}{2} x}_{\text{trailing edge thickness term}} + \quad (12) \\
 & + \underbrace{x\sqrt{1-x} \frac{v_{\text{TE}}^u - v_{\text{TE}}^l}{2} (1-x)^{n_{\text{BP}}}}_{\text{supplementary leading edge shaping term}},
 \end{aligned}$$

respectively, where HT denotes half thickness (note that in order to simplify the equations we are assuming $n_{BP}^u = n_{BP}^l = n_{BP}$). This is, in fact, a variable transformation, which gives us the possibility of breaking up the required first mapping of (10) into two more easily manageable sub-problems (divide and conquer again!), the right hand one of which we have just solved:

$$\begin{aligned} (t/c, c_l) &\mapsto \left(\frac{v_0^u + v_0^l}{2}, \frac{v_0^u - v_0^l}{2}, \dots, \frac{v_{LE}^u - v_{LE}^l}{2} \right) \\ &\mapsto (v_0^u, v_1^u, \dots, v_{LE}^l) \mapsto [z^u(x), z^l(x)], \end{aligned} \quad (13)$$

This has not reduced the dimensionality of our design space yet, but has given us intervening variables that are more useful in terms of exploiting the separation of camber and thickness distribution and have therefore taken us closer to the ultimate goal of mapping from the $(t/c, c_l)$ space. For the final remaining step we divide the problem once more and first look at the

$$t/c \mapsto \left(\frac{v_0^u - v_0^l}{2}, \frac{v_1^u - v_1^l}{2}, \dots, \frac{z_{TE}^u - z_{TE}^l}{2}, \frac{v_{LE}^u - v_{LE}^l}{2} \right) \quad (14)$$

subproblem. Having already established that the thickness distribution of the six example airfoils depends only on the maximum thickness to chord ratio t/c and noting that the relationship is clearly linear, the r^{th} half thickness term in the description of the parametric airfoil will be a function of t/c as follows:

$$\begin{aligned} \frac{v_r^u - v_r^l}{2} \Big|_{t/c} &= \\ & \frac{v_{r \sim \text{SC}(2)-0410}^u - v_{r \sim \text{SC}(2)-0410}^l}{2} + \\ & + \left(\frac{v_{r \sim \text{SC}(2)-0412}^u - v_{r \sim \text{SC}(2)-0412}^l}{2} - \frac{v_{r \sim \text{SC}(2)-0410}^u - v_{r \sim \text{SC}(2)-0410}^l}{2} \right) \times \frac{t/c-10}{t/c-12}, \\ & t/c \in [10, 12]. \end{aligned} \quad (15)$$

Note that where we used coefficients from the class-shape transformation of, say, SC(2)-0412 (for example, $v_{r \sim \text{SC}(2)-0412}^l$), we could equally have used the relevant coefficients of any of the 12 % thick airfoils, as they only appear as part of the transformation coefficients of the thickness distributions, which, as we have seen, are identical for airfoils of the same maximum thickness. This, as well as the above equation, are equally applicable to the calculation of the two remaining parameters, the additional leading edge shaping term and the trailing edge thickness term.

We now need to find a way of constructing the coefficients of the camber curve transformation of the parametric airfoil, that is, to find the

$$(t/c, c_l) \mapsto \left(\frac{v_0^u + v_0^l}{2}, \frac{v_1^u + v_1^l}{2}, \dots, \frac{z_{TE}^u + z_{TE}^l}{2}, \frac{v_{LE}^u + v_{LE}^l}{2} \right) \quad (16)$$

part of the mapping (13). This is a slightly more complicated proposition, as the shape of the camber curve, though chiefly influenced by the design c_l , varies between airfoils of different thicknesses, as shown in Figure 4. We therefore need to construct a model of each of the camber curve parameters on the right hand side of (16) as a function of design c_l and t/c , based on the six examples provided by our chosen six SC(2) sections.

5.2 A Gaussian Process Model

Considering that the sets of transformation coefficients v and z (which we have identified earlier) define *approximations* of the six ‘training’ airfoils (when inserted into equations (2) and (3)) and therefore the camber line coefficients are also approximations of the camber lines of the six airfoils, we shall build a *regression* model of (16) (as opposed to an interpolating one) to filter out the ‘noise’ in the coefficient values.

We choose to work with a Gaussian Process modeling approach – kriging – and we use the implementation described in Forrester et al (2008). The interested reader is invited to consult this reference for the details of the formulation; here we limit ourselves to a brief summary of the problem setup.

Let us, for each camber line class-shape transformation coefficient (the right hand side of (16)), consider a 6×2 matrix \mathbf{X} of the t/c ratios (column one) and design c_l values (column two) of our set of supercritical airfoils and a 6×1 vector \mathbf{y} of the corresponding values of the current camber transformation coefficient. We then construct a matrix Ψ of correlations between the 6 training points contained in \mathbf{X} , which is now a function of the correlation coefficients θ . Additionally, to account for the inexact nature of the approximations (2) and (3) constructed with the transformation variables, we add a regression parameter λ to the leading diagonal of the correlation matrix – both θ and λ are estimated subsequently via a likelihood maximization procedure.

The kriging regression model is thus given by:

$$\hat{y}(t/c, c_l) = \hat{\mu} + \psi^T (\Psi + \lambda \mathbf{I})^{-1} (\mathbf{y} - \mathbf{1}\mu), \quad (17)$$

where

$$\hat{\mu} = \frac{\mathbf{1}^T (\Psi + \lambda \mathbf{I})^{-1} \mathbf{y}}{\mathbf{1}^T (\Psi + \lambda \mathbf{I})^{-1} \mathbf{1}}, \quad (18)$$

\mathbf{I} is a 12×12 identity matrix and ψ is a vector containing the correlations between the training data and the $(t/c, c_l)$ pair, where we wish to predict the current class-shape transformation parameter.

The model (17) is an approximation of mapping (16) and thus completes the mapping (13). We therefore now have the complete route from $(t/c, c_l)$ to the explicit definition of the airfoil based on equations (2) and (3). This, then, is a parametric airfoil depending on two design variables, whose ranges are defined by the six airfoil training set: $t/c \in [10, 12]$, $c_l \in [0.4, 0.7]$.

5.3 Physical Significance

While not strictly relevant from the perspective of an automated design process, it is still natural to ask: is there a correlation between the physical properties of the new parametric airfoil we have created and the pair of design variables that control its shape?

In order to answer this question we generated 20 pairs of $(t/c, c_l)$ values, arranged in the $[10, 12] \times [0.4, 0.7]$ design space in a Latin hypercube sampling pattern [see Forrester et al (2008) for details of the formulation and the algorithm used⁴]. We then generated the corresponding airfoils using our parametric mapping and evaluated the designs in terms of their maximum thickness to chord ratios and their lift coefficients – the latter computed using the computational fluid dynamics solver FLUENT, with GAMBIT employed to create the unstructured mesh ($\sim 200,000$ cells for each mesh to obtain the required curve detail). In terms of the flow conditions, we kept the Reynolds number constant at 30×10^6 with the Mach number M and the angle of attack α allowed to float until the drag divergence Mach number ($dc_d/dM = 0.1$) was found and the pressure coefficient plot was comparable to the idealized model shown in Figure 5 (see the report by Harris (1990) for background information). More specifically, initially we computed three flow fields, one at the Mach numbers found using the Korn equation [Mason (2009)]

$$M + \frac{c_l}{10} + \frac{t}{c} = 0.95 \quad (19)$$

and the other two at $M + 0.001$ and $M - 0.001$ respectively. We then fitted a polynomial to this data in the c_d versus Mach number space, differentiation of which yielded the drag divergence Mach number (where $dc_d/dM = 0.1$). The flow field was then solved at this new condition, with the result added to the existing solutions and a polynomial fitted once more. This pattern was continued until the position of the drag divergence Mach number remained constant over consecutive iterations. The resultant pressure coefficient graph at this position was compared to the ideal graph of Figure 5, with the above heuristic repeated at a new α if the graphs were dissimilar.

Figures 6 and 7 show the results of this experiment. Correlation can be observed in both cases. In fact, the maximum thickness to chord ratio of the parametric airfoil can clearly be said to be equal, for most practical purposes, to the value of the ‘ t/c ’ design variable. Once again, this has little significance in most automated design processes, but it can be seen as a useful feature, for example, if we want to restrict the design space to, say, wings that can accommodate a certain spar depth (that is, their t/c must be greater than a certain threshold value).

Much of the reasoning behind the construction of the parametric mapping was based on the observation that we can link the design variables to easily separable elements of the airfoil shapes (crucially, we have found the thickness

⁴ Latin hypercubes have uniform projections onto all axes and are therefore ideal for correlation studies.

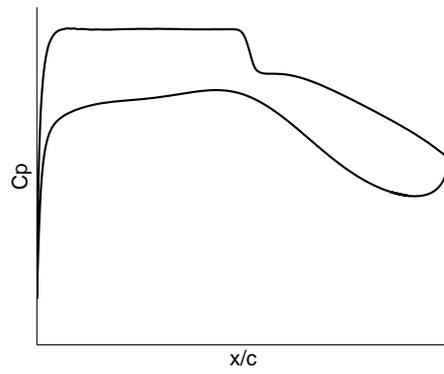


Fig. 5 Typical 'flat top' pressure distribution around an SC(2) supercritical airfoil, serving as a target for the search for the design conditions for a given supercritical airfoil.

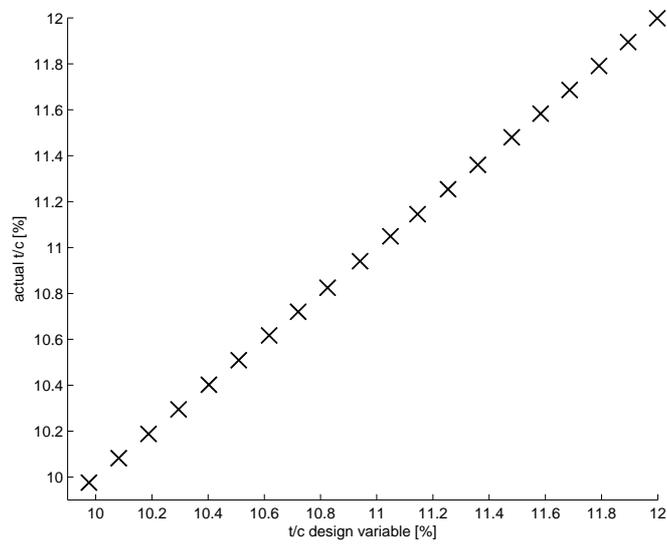


Fig. 6 Actual thickness to chord ratio versus the 't/c' design variable value at 20 airfoils spread evenly across the design space.

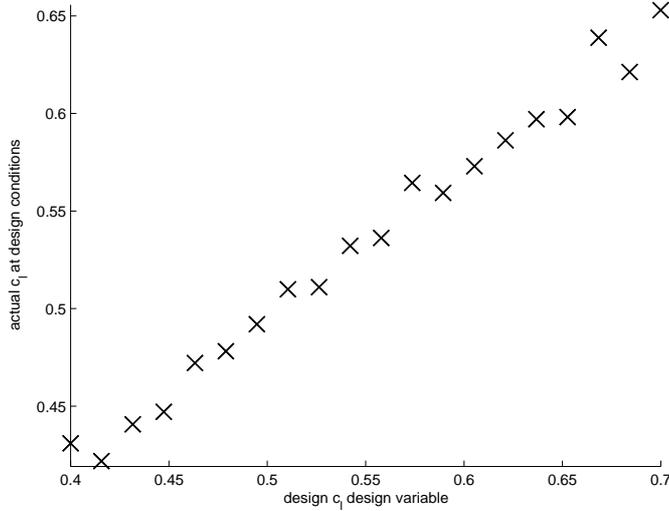


Fig. 7 Actual c_l at design M and α versus the ‘ c_l ’ design variable value at 20 airfoils spread evenly across the design space. The corresponding R^2 value is 0.9869.

distributions of the six members of the family to be connected exclusively to the maximum thickness to chord value that headlines each airfoil). We look at a more general case next, where such simplifications are no longer possible.

6 A More Diverse Family

6.1 Patterns

Consider now a larger subset of SC(2) supercritical airfoils: SC(2)-0406, SC(2)-0606, SC(2)-0706, SC(2)-0410, SC(2)-0610, SC(2)-0710, SC(2)-0412, SC(2)-0612, SC(2)-0712, SC(2)-0414, SC(2)-0614 and SC(2)-0714. These 12 sections now encompass a broader range of design c_l values and t/c ratios than the set we analyzed earlier. The crucial difference with respect to the previous family of six is that the pattern of thickness distributions and camber curve variations within the family is considerably more complicated. We shall use this broader family to illustrate a more general form of the class-shape transformation dimensionality reduction heuristic presented earlier.

Once again we begin by approximating every member of the chosen family through its class-shape transformation. This time, we set the orders of the Bernstein polynomial terms to $n_{BP}^u = 2$ and $n_{BP}^l = 3$ for the upper and lower surfaces respectively. The sets of transformation coefficients of the 12 target airfoils yielded by solving equation (6) are depicted in Figure 8.

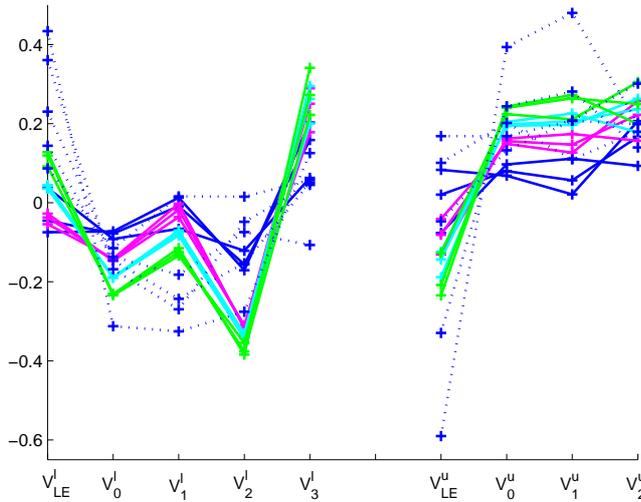


Fig. 8 Class-shape transformation coefficients of a set of well-known airfoils. The heavy, continuous lines denote the 12 SC(2) supercritical airfoils discussed here, while the dotted lines represent the approximation coefficients of NACA5410, NLR7301, RAE5215, RAE2822 and NACA24-011. Note the distinctive ‘wr’-shaped pattern of the SC(2) family.

Also shown in the same figure are the ‘coefficient-fingerprints’ of a number of additional airfoils. It is clear that the SC(2) coefficient sets form a rather obvious ‘wr’-shaped pattern, rather dissimilar to the shapes corresponding to the other airfoils whose coefficient patterns are depicted in the same figure (especially in the case of the ‘w’ corresponding to the rather typical shapes of the SC(2) lower surfaces).

If we hadn’t already studied a six airfoil subset of this family, the existence of this pattern would be our first indication that we are likely to need considerably fewer design variables to cover this restricted space than the 11 variables of the class-shape transformation itself (the nine shown in Figure 8, plus the two trailing edge thickness parameters). Essentially, we have the opportunity to trade flexibility for conciseness. Restricting any design searches to these ‘wr’-shaped coefficient sets also has the advantage of ensuring that the design space will only contain physically ‘sensible’ (and ‘supercritical’) shapes.

As before, we shall aim to map the $t/c, c_l$ pair to the space of class-shape transformation coefficients and therefore to z^u and z^l , that is, we seek to build the first part of the mapping (10) (equations (2) and (3) form the second part).

Of course, all this reasoning on patterns is based on intuition and is the expression of certain assumptions – not least that the shapes of the SC(2) airfoils are chiefly determined by design lift coefficient and thickness and that

otherwise their design generally follows the same principles across the family (this was clear in the case of our earlier, ‘separable’ set, but less obvious here). Additionally, we will assume separability, that is, that each class-shape transformation variable can be generated from a $(c_l, t/c)$ pair via a mapping that is independent of the other transformation variables. Intuitively, the similar shapes of the transformation variable patterns indicate that this is a reasonable assumption to make. In the case presented earlier we made, tacitly, a weaker form of this assumption: there we could, at least, be sure of the separability of the mappings between the half-thickness coefficients and ‘ t/c ’.

The purely intuitive nature of the above, however, is of no practical significance, as long as we manage to construct a well-posed model of the mappings and the resulting reduced dimensionality airfoil is suitable for design studies. We shall return shortly to the mathematical ‘checks and balances’ we can use to confirm the correctness of our assumptions (at least from a practical perspective) – here we merely note that an additional bonus and further confirmation of the correctness of these assumptions would be the existence, as in the previous study, of some degree of correlation between the design variable values t/c and c_l and the maximum thickness and the lift coefficient of the resulting instantiation of our parametric airfoil.

6.2 Another Kriging Model

We postulated that, within the mapping (10), the individual $(t/c, c_l) \mapsto v$ mappings are considered to be separable. We can therefore attempt to build a model of each class-shape transformation coefficient v in terms of c_l and t/c , based on the 12 known pairings resulting from our approximations of the 12 SC(2) airfoils.

We no longer have any of the handholds we took advantage of in the previous case (the six airfoil family), so we have to construct 11 such models. The process employed is much the same as before – we find the model parameters by maximizing the likelihood of the data – except that on this occasion we do not build the models in terms of the intervening camber coefficient variables, but directly in terms of the variables describing the airfoil surfaces. Figure 9 is a depiction of one such model, also showing the 12 training data points, one representing each example airfoil.

If our assumption of separability was seriously wrong, this is where that would become apparent. For instance, in the absence of a clear trend (which would imply that a third variable has a significant influence over the shapes of the airfoils) the variations within the log-likelihood landscape would be generally low and would not have clear maxima. Recall that this is a function of the the kriging model parameters (a θ per dimension and a global regression parameter λ) and the presence of significant additional factors would lead to very different combinations of these parameters being almost equally likely – clearly a sign that there are no trends in the data. Should the reader opt for other methods of determining the the θ ’s and λ , these are usually also

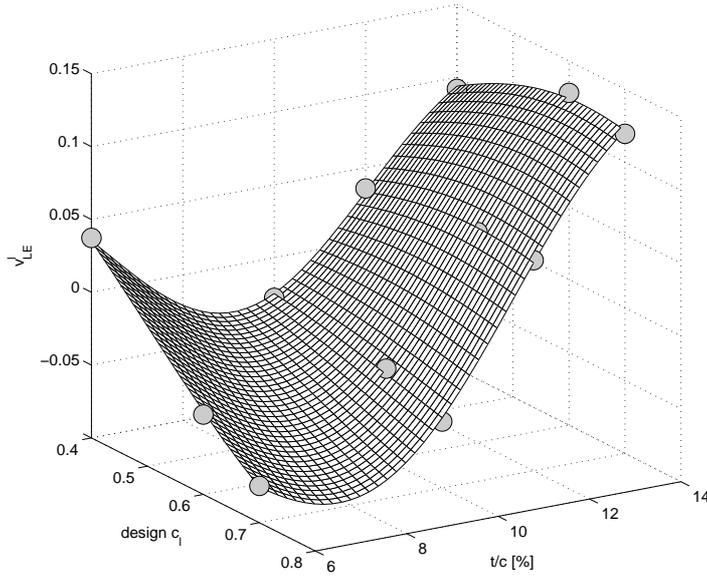


Fig. 9 Gaussian Process regression model of the value of one of the class-shape transformation parameters (v_{LE}^t), trained on the 12 values found as optimal for the set of SC(2) airfoils.

equipped with warning devices that will indicate if the initial assumptions are wrong. For example, leave- k -out cross-validation [Forrester et al (2008)] would yield cross-validation errors per data point comparable to the range of the responses – again, a sign that other factors have a significant impact on the data⁵.

6.3 Physical Relevance

As before, a space-filling set of designs was generated and tested from the point of view of the accuracy of our approximation of mapping (10). Figure 10 shows that, as before, the t/c design variable is virtually equal to the maximum thickness to chord ratio of the airfoil the mapping will generate. A weaker correlation can be observed in terms of the c_l variable (Figure 11)–

⁵ We stress the word ‘significant’ here for a good reason – in the process of tailoring the SC(2) airfoils small shape alterations *were* necessary in some cases to obtain the desired pressure profiles (in particular shock locations) and drag rise Mach numbers, but, for practical purposes, we can assume that the two major factors with consistently significant impact were t/c and the desired c_l .

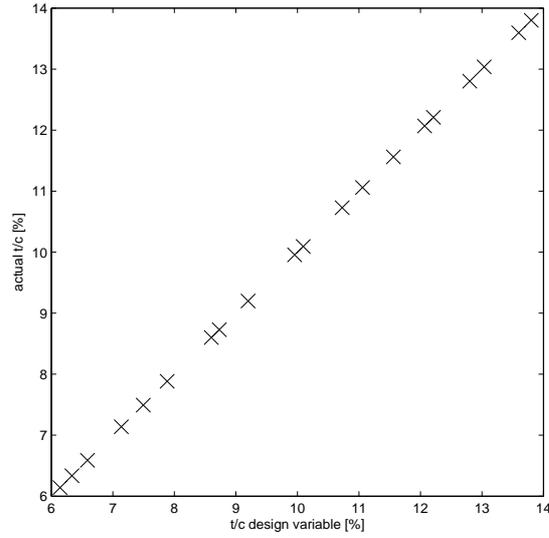


Fig. 10 Actual thickness to chord ratio versus the ' t/c ' design variable value at 20 airfoils spread evenly across the design space.

Simplex iterations	Function evaluations	Best objective
0	1	0.0156
1	17	0.0150
2	18	0.0150771
3	19	0.0150771
4	21	0.0136
5	22	0.0136
6	23	0.0136
7	24	0.0136
8	25	0.0136
9	26	0.0136
10	27	0.0136

Table 1 Simplex optimization history of the search for an 11% thick airfoil with a design c_l of 0.5. The starting point was the Kulfan transformation of SC(2)-0412.

here we can see a slight loss of approximation accuracy compared to the first case.

Figure 12 is a further illustration of the physical significance of the design variables: different values of the c_l variable produce airfoils with variable camber (left), while the camber is maintained and the thickness changes as t/c varies (right).

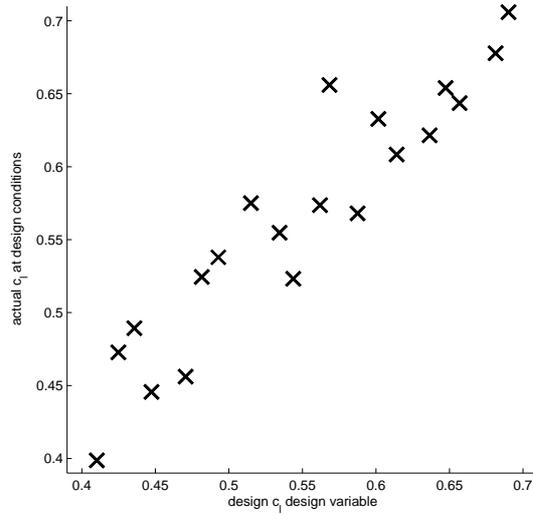


Fig. 11 Actual c_l versus the ' c_l ' design variable value at 20 airfoils spread evenly across the design space. The corresponding R^2 value is 0.8747.

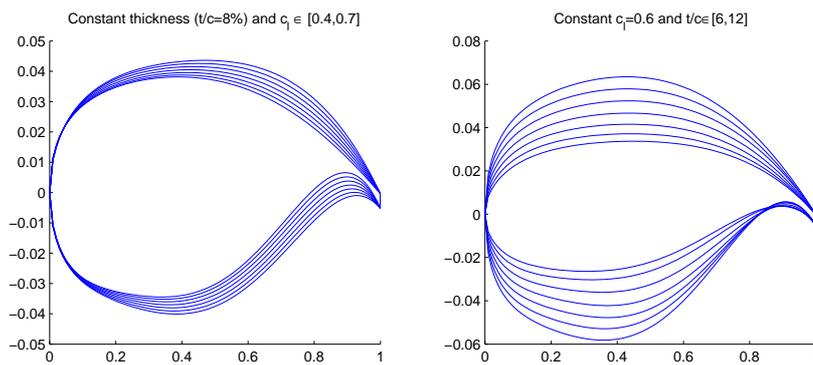


Fig. 12 Examples of instances of our two variable parametric airfoil.

7 Reflections on the Design Process

So what does all this mean from the perspective of the preliminary design process⁶? In the previous section we described the construction of a parametric

⁶ We choose to define as 'preliminary' the first phase of the design process that is centered around a geometry.

airfoil defined by two design variables – here we summarise this process as follows.

1. INPUTS:

$\mathcal{A}_i^{\text{zz}} = \{(\mathbf{x}, \mathbf{z})^{(i)}, (c_l, t/c)^{(i)}\}$, $i = 1, \dots, n_{\text{SC}}$, that is:

- n_{SC} airfoils $\mathcal{A}_i^{\text{zz}}$ given as sets of coordinate pairs
- n_{SC} corresponding pairs of design c_l and t/c values

The final goal of the algorithm is to enable the construction of new such pairs given $(c_l, t/c)$ values not included in this original set.

2. STEP I. – ENCODING OF THE INPUTS:

for $i = 1, \dots, n_{\text{SC}}$, $\mathcal{A}_i^{\text{zz}} \rightarrow \mathcal{A}_i^{\text{K}}$

We apply the Kulfan transformation to each of the n_{SC} airfoils, that is, for each $\mathcal{A}_i^{\text{zz}}$ we solve Equation (6) to obtain the corresponding set of $\mathcal{A}_i^{\text{K}} = [v_0^{\text{u}(i)}, v_1^{\text{u}(i)}, \dots, v_{n_{\text{BP}}^{\text{u}(i)}}^{\text{u}(i)}, z_{\text{TE}}^{\text{u}(i)}, v_{\text{LE}}^{\text{u}(i)}, v_0^{\text{l}(i)}, v_1^{\text{l}(i)}, \dots, v_{n_{\text{BP}}^{\text{l}(i)}}^{\text{l}(i)}, z_{\text{TE}}^{\text{l}(i)}, v_{\text{LE}}^{\text{l}(i)}]$ Kulfan coefficients.

3. STEP II. – RE-PARAMETERIZATION:

For each of the Kulfan parameters we build an approximation model of the form of Equation (17) in terms of the corresponding $(c_l, t/c)$ pairs.

We begin with v_0^{u} . For each of the $(c_l, t/c)^{(i)}$, $i = 1, \dots, n_{\text{SC}}$ pairs we have a corresponding $v_0^{\text{u}(i)}$ value from STEP I. We thus construct the kriging model $v_0^{\text{u}}(c_l, t/c)$ based on these n_{SC} data.

We repeat the above step for $v_1^{\text{u}}, \dots, v_0^{\text{l}}, v_1^{\text{l}}$ and the rest of the Kulfan parameters (Figure 9 shows an example of such a bivariate model).

4. STEP III. – CONSTRUCTION OF THE NEW PARAMETRIC GEOMETRY

We now assemble the models $v_0^{\text{u}}(c_l, t/c)$, $v_1^{\text{u}}(c_l, t/c)$, ..., $v_0^{\text{l}}(c_l, t/c)$, $v_1^{\text{l}}(c_l, t/c)$, ... from STEP II. into a parametric description of a complete Kulfan airfoil. Simply inserting these Kulfan coefficients into Equation (4) yields the sets of (\mathbf{x}, \mathbf{z}) coordinates.

The process summarised above and illustrated earlier for the family of SC(2) supercritical airfoils can be employed to exploit patterns in the class-shape transformation coefficients of other families of similar airfoils. This reduced dimensionality model can then be used for global design searches, safe in the knowledge that we have reduced the contribution of the airfoil to the overall dimensionality of the airframe geometry and that we do not need to worry about setting sensible variable bounds.

As an illustration of the time-savings afforded by this type of approach, let us consider the following design problem. The preliminary design process of an airliner requires an 11% thick airfoil with a design c_l of 0.5. We have discussed the class- shape function parameterization in great detail in this paper and we could deploy it in this context too. We can apply the Kulfan transformation to the SC(2)-0412 airfoil, which is the existing supercritical airfoil that matches our requirements most closely (having a thickness-to-chord ratio of 12% and a design lift coefficient of 0.4). We can then run a local search process starting from this airfoil (its corresponding Kulfan coefficients), which aims to minimize the drag of the airfoil, subject to the thickness constraint. Each objective call issued by the optimization algorithm (we employed a Nelder and Mead simplex search here) involves iterating through a sequence of angles of attack to attain the required lift coefficient. This is a relatively expensive process – we use a Navier-Stokes flow solver at the same level of fidelity as employed when generating the correlation plots presented earlier (e.g., Figure 11) – we ran the Nelder and Mead search with a computational budget of 10 iterations. Table 1 shows the results (the search history) of this exercise, featuring an ultimate drag coefficient value of 0.0136.

An alternative approach would be to deploy the parametric airfoil described here, using it to generate directly the target 11% thick airfoil with a design c_l of 0.5. Building this airfoil and running the same iterative analysis process (until we obtain the angle of attack that gives the required c_l) yields here a drag coefficient value of 0.01298. This is slightly better than the airfoil obtained through the rather expensive simplex optimization process started from SC(2)-0412 (see Table 1) and it involved no design search at all (the computational cost of instantiating the parametric airfoil for $c_l=0.5$ and $t/c=11\%$ is negligible).

Once this first step of the design process is complete, we are left with an airfoil expressed in the form of equations (2) and (3), i.e., as a Kulfan transformation, which can form the starting point of a subsequent local search. This second optimization procedure can then exploit the aerodynamic significance of some of the class-shape transformation variables (e.g., the first term is related to leading edge radius, number $n^{\text{BP}} + 1$ controls the boattail angle), or can simply allow an automated optimizer to exploit the current basin of attraction in terms of some design goal.

To summarise then, the preliminary designer in need of a low drag airfoil of a specific lift coefficient and thickness to chord ratio would either have to run a global search on a very flexible, generic parameterization (very expensive) or a local search from the nearest existing airfoil designed for similar flow conditions (moderately expensive). For the type of design scenario outlined here, the proposed alternative requires no optimization, indeed no analysis runs at all – a potentially significant saving at the preliminary design stage when the dimensionality of the complete airframe geometry might be quite high.

8 Conclusions and Future Work

In the above we have shown that it is possible to build a concise parametric geometry over a design space defined purely by a set of basis shapes, regardless of how these example designs are represented. At a more fundamental level, the problem of building such geometries as concisely as possible, translates into a more general question, which is worth further investigation and could be phrased as follows.

Let us consider a set of curves (or surfaces), which represent a diverse range of feasible (though not necessarily optimal) solutions to a design problem (the examples or potential bases). What is the minimum dimensionality of a parametric curve (or surface) that can reproduce all of the sample curves (or surfaces) to within a specified level of accuracy, while also creating a smooth subspace of designs defined in terms of these ‘training’ examples?

In a previous paper [Sóbester (2009a)] we have approached the problem using a NURBS description. Here we have shown the Kulfan transformation to be another feasible way of capturing the training cases and building the parametric mapping – at least for the specific case of supercritical airfoils. We have constructed two parametric airfoils that distil the aerodynamic reasoning behind the designs of their respective subsets of basis airfoils down to two design variables. Moreover, in both cases the two design variables show strong correlations with physical parameters (geometrical and aerodynamic) of the parametric airfoil, whose shape they determine, a feature that can be useful in the context of human interventions in the design process (as opposed to a purely automated search for a shape that optimizes some goal function).

Future work therefore should consider applying either strategy to broader (or different) classes of shapes. As this initial study indicates, the method has the potential to parameterize complex shapes very concisely, while constructing a design space that is relatively unlikely to include infeasible regions. Here is a summary of the key steps to be followed in order to do this successfully for other applications.

- 1) Identify the key variables. It is best to select a combination of physics- and performance-based parameters and geometrical parameters (preferably ones with intuitive engineering appeal) – a simple example, considering a whole airframe, might be a set including wing sweep angle, cruise Mach number, maximum load factor, range, aspect ratio, etc.
- 2) Identify the ‘basis geometries’, that is, the set of distinct designs for which the parameters chosen above are available.
- 3) Apply the Kulfan transformation to these basis geometries (other shape encoding methods could also be adapted, as indicated above).
- 4) Build kriging models of each element of the encoding from 3), in terms of the variables identified at 1).
- 5) Generate instances of the new parametric geometry and test their performance against what you would expect from their inputs (e.g., does the instance of the parametric airframe indeed have the range specified as the input to the parametric geometry?).

Clearly, the success of building such a geometry will depend heavily on the choice of variables at 1) and the choice of the parameterization scheme selected at 3), so in case 5) yields unsatisfactory results, the process will have to be repeated with different choices at 1) and/or 3). There is no hard and fast recipe for either and, indeed, for certain applications, finding the correct choices might be hard or impossible. It is hoped that the example shown here will encourage readers to explore these choices for their own applications.

Acknowledgements The first author's work has been supported by the Royal Academy of Engineering and the Engineering and Physical Sciences Research Council through their Research Fellowship scheme.

References

- Collins L, Saunders D (1997) Profile: Airfoil geometry manipulation and display. Contractor Report 177332, NASA
- Forrester AIJ, Sóbester A, Keane AJ (2008) Engineering Design via Surrogate Modelling. John Wiley and Sons
- Harris CD (1990) NASA supercritical airfoils – a matrix of family-related airfoils. Technical Paper 2969, NASA
- Jameson A (1988) Aerodynamic design via control theory. *Journal of Scientific Computing* 3(3):233–260
- Kulfan BM (2006) “Fundamental” parametric geometry representations for aircraft component shapes. AIAA 2006-6948 pp 1–45
- Kulfan BM (2008) Universal parametric geometry representation method. *Journal of Aircraft* 45(1):142–158, doi: 10.2514/1.29958
- Kulfan BM, Bussoletti JE (2006) “Fundamental” parametric geometry representations for aircraft component shapes. AIAA 2006-6948
- Mason W (2009) Configuration Aerodynamics. Lecture Notes, Virginia Tech University
- Robinson GM, Keane AJ (2001) Concise orthogonal representation of supercritical airfoils. *Journal of Aircraft* 38:580–583
- Rozvany GIN (2009) A critical review of established methods of structural topology optimization. *Structural and Multidisciplinary Optimization* 37:217–237
- Sóbester A (2009a) Concise airfoil representation via case-based knowledge capture. *AIAA Journal* 47(5):1209–1218
- Sóbester A (2009b) Exploiting patterns in the Kulfan transformations of supercritical airfoils. 9th AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, SC
- Vanderplaats GN (1979) Efficient algorithm for numerical optimization. *Journal of Aircraft* 16(12):842–847
- Vanderplaats GN (1984) Numerical optimization techniques for engineering design : with applications. McGraw-Hill, New York
- Whitcomb RT (1974) Review of NASA supercritical airfoils. The Ninth Congress of the International Council of the Aeronautical Sciences ICAS 74-10
- Whitcomb RT, Clark LR (1965) An airfoil shape for efficient flight at supercritical mach numbers. Technical Report TM X-1109, NASA