

# Blade Envelopes Part I: Concept and Methodology

Chun Yui Wong<sup>†\*</sup>, Pranay Seshadri<sup>‡\*</sup>, Ashley Scillitoe<sup>\*</sup>,  
Andrew B. Duncan<sup>‡\*</sup>, Geoffrey Parks<sup>†</sup>

<sup>†</sup>Department of Engineering, University of Cambridge, U.K.

<sup>‡</sup>Department of Mathematics, Imperial College London, U.K.

<sup>\*</sup>Data-Centric Engineering, The Alan Turing Institute, U.K.

Blades manufactured through flank and point milling will likely exhibit geometric variability. Gauging the aerodynamic repercussions of such variability, prior to manufacturing a component, is challenging enough, let alone trying to predict what the amplified impact of any in-service degradation will be. While rules of thumb that govern the tolerance band can be devised based on expected boundary layer characteristics at known regions and levels of degradation, it remains a challenge to translate these insights into quantitative bounds for manufacturing. In this work, we tackle this challenge by leveraging ideas from dimension reduction to construct low-dimensional representations of aerodynamic performance metrics. These low-dimensional models can identify a subspace which contains designs that are invariant in performance—the inactive subspace. By sampling within this subspace, we design techniques for drafting manufacturing tolerances and for quantifying whether a scanned component should be used or scrapped. We introduce the blade envelope as a visual and computational manufacturing guide for a blade. In this paper, the first of two parts, we discuss its underlying concept and detail its computational methodology, assuming one is interested only in the single objective of ensuring that the loss of all manufactured blades remains constant. To demonstrate the utility of our ideas we devise a series of computational experiments with the Von Karman Institute’s LS89 turbine blade.

## 1 INTRODUCTION

Manufacturing variations and in-service degradation have a sizeable impact on aerodynamic performance of a jet engine (see Figure 1). Gauging the aerodynamic repercussions of such variability prior to manufacturing a component is challenging enough, let alone trying to predict what the amplified impact of any in-service degradation might be. In a bid to reduce losses and mitigate the risks in Figure 1, designers today pursue a two-pronged approach. First, components are being designed to operate over a range of conditions (and uncertainties therein) via robust optimization techniques [10, 11] as well as more traditional design guides such as loss buckets—i.e., loss across a range of positive and negative incidence angles [4]. In parallel, there has been a growing research effort to assess 3D manufacturing




Fig. 1. Impacts associated with manufacturing variations in a jet engine; see [1–9].

variations and in-service degradation by optically scanning (via GOM) the manufactured blades, meshing them, and running them through a flow solver [12]. Both approaches, while useful in extracting aerodynamic inference, are limiting. One of the key bottlenecks is the cost of evaluating flow quantities of interest via computational fluid dynamics (CFD), as the dimensionality of the space of manufactured geometries is too large to fully explore, even with an appropriately tailored design of experiments (DoE). To reduce the dimensionality, some authors [3, 12] use principal components analysis (PCA) to extract a few *manufacturing modes*, which correspond to modes of largest manufacturing deviation observed in the scanned blades. One drawback of this approach is that the PCA model is not performance-based, i.e., the mode of greatest geometric variability need not correspond to the mode of greatest performance scatter, a point raised by Dow and Wang [13]. Additionally, GOM scans can only be carried out on cold and manufactured components, ignoring the uncertainty in performance associated with in-service operating conditions. Finally, through neither of these paths are manufacturing engineers offered a set of pedigree rules or guides on manufacturing for an individual component, prior to actually manufacturing the component. This motivates some of the advances in this paper. We argue that challenges associated with both manufacturing

\*Address all correspondence to Chun Yui Wong, cyw28@cam.ac.uk

variability and in-service deterioration can be adequately addressed by introducing the concept of a *blade envelope*.

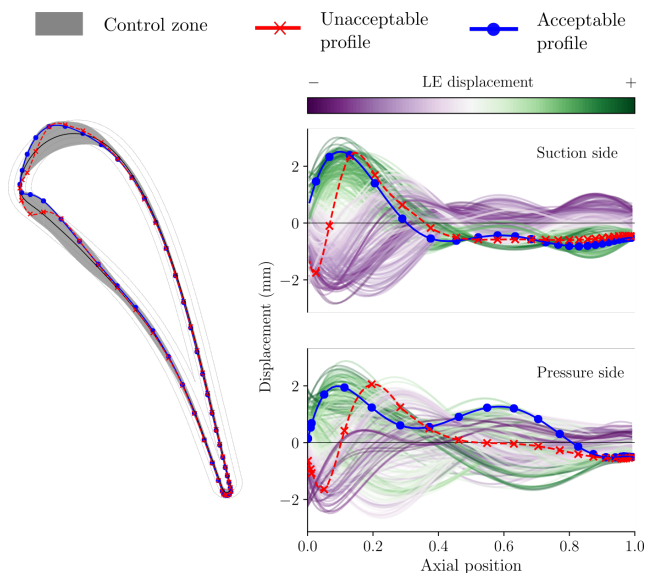


Fig. 2. A blade envelope. It comprises a control zone (grey region) and a tolerance covariance (characterized by the airfoil displacement plots). Dots denote CMM measurements. Note: in displacement plots, the axial distance is normalized by the axial chord.

A blade envelope (see Figure 2) is a visual and computational manufacturing guide for a blade that instructs manufacturers whether a scanned component should be used or scrapped. It does this by identifying a *control zone* (shown in grey) within which all scanned blades that adhere to a certain *tolerance covariance* are guaranteed to offer near identical performance. One way to visualize the tolerance covariance is through the leading edge (LE) displacement—invariant profiles which are positively displaced near the LE tend also to be displaced positively in the adjacent region along the chord. In broad terms, the tolerance covariance captures the characteristics of the geometric variations of invariant profiles. We offer a rigorous definition of this concept in the subsequent sections of this paper. As an illustrative demonstration of this idea, consider two scanned profiles shown as red and blue markers in Figure 2. Each marker corresponds to airfoil profile measurements taken from a coordinate measurement machine (CMM). The measurements are interpolated using B-splines to obtain the airfoil profiles—both of which lie within the control zone. However, as can be expected, it is the profile variation within the control zone that truly dictates how adverse the change in performance will be. It is clear that the red profile is unacceptable and therefore should be scrapped, while the blue one can be used in an engine.

The utility of a blade envelope, however, does not end here: it has profound consequences for the way we think about tolerances too. Consider the work of Lobato et al. [14] who offer a recipe for acceptance and rejection of airfoils based on the curvature variation within a chord-wise interval. Provided the curvature of the measured profile lies within a computed upper

and lower tolerance band—neither of which are necessarily inferred from CFD—the profile is deemed acceptable. With blade envelopes, profile tolerance constraints on displacement and curvature are supported by computational and/or experimental aerodynamic insight. This helps avoid unnecessary scrapping due to potentially overly conservative bounds.

Beyond the binary scrap-or-use decision, blade envelopes can provide clarity on where tolerances need to be tightened and where they can be relaxed—bringing forth the possibility of not only reducing performance variability, but also reducing manufacturing scraps. This may engender transformative cost-saving manufacturing strategies for certain components, underpinned by greater predictive confidence in the tolerances.

So, how are these tolerances obtained, and how can performance guarantees for designs manufactured within the envelope be offered? This two-part paper presents a series of computational methods to rigorously answer these questions. In this first part, we introduce blade envelopes for loss, focusing primarily on the governing methodology. In the second part [15], we study how one can accommodate multiple objectives—both scalar-valued and vector-valued—when constructing such envelopes, and also elucidate their parallels to *inverse design*.

## 2 TURBINE TEST CASE

In this section, the test case used throughout this paper is introduced, along with details of its *manufacturing space* and the strategy for obtaining numerical solves. However, it should be emphasized that the concept of a blade envelope and its generation is not limited to any particular profile nor restricted to any manufacturing space.

### 2.1 Geometry and design space

We select the Von Karman Institute LS89 linear turbine cascade [16]. This transonic, highly-loaded blade serves as a rich experimental (and subsequently computational) repository with Schlieren flow visualizations, blade static pressure measurements, exit flow angle measurements, and even blade convective heat transfer values. An experimental campaign was carried out by Arts et al. [16] at a range of different exit Mach numbers ranging from 0.70 to 1.10, exit Reynolds numbers ranging from  $0.5 \times 10^5$  to  $2.0 \times 10^6$  and freestream turbulence intensities of 1% to 6%. The overall chord of the tested profile was approximately 67 mm with a stagger angle of  $55^\circ$ .

To simulate deformations arising from manufacturing variations, a large design space around the baseline profile is defined. This space, notionally denoted by  $\mathcal{X}$ , is parameterized by  $d = 20$  independent free-form deformation (FFD) design variables, *scaled* to lie within the range  $-1$  to  $1$ , such that each design vector  $\mathbf{x} \in [-1, 1]^d$ . Thus, one can think of  $\mathcal{X}$  as being a  $d$ -dimensional hypercube. The FFD nodes are constrained to move only in the direction perpendicular to the inflow; some possible deformations are captured in Figure 3. Note that the magnitude of geometric variations imposed here grossly exaggerates the typical scales of deviations seen in real turbine blades.




Fig. 3. The FFD box (black dots) with some possible deformations within the design space (grey curves) and the baseline profile (black curve).

Table 1. Flow properties used in this study.

Flow property	Symbol	Value
Inlet stagnation pressure	$P_{01}$	$1.1 \times 10^6 \text{ N m}^{-2}$
Inlet stagnation temperature	$T_{01}$	592.295 K
Inlet density	$\rho$	$1.2866 \text{ kg m}^{-3}$
Freestream Reynolds number	$Re$	$6.0 \times 10^5$
Exit static pressure	$p_2$	$5.23 \times 10^5 \text{ N m}^{-2}$
Heat capacity ratio	$\gamma$	1.4

## 2.2 Computational fluid dynamics

The open-source CFD suite  $SU^2$  [17] is used for the mesh deformation of the various geometries as well as their subsequent numerical solves. The latter are obtained by solving the Reynolds-averaged Navier-Stokes (RANS) equations with a Spalart-Allmaras turbulence closure. No transition model is used in our work and all the boundary layers are assumed to be turbulent. A stagnation pressure driven inlet and an exit static pressure outlet were used to set the passage pressure ratio. Periodic boundary conditions were imposed along the pitch-wise direction, and the freestream turbulent viscosity ratio was set to 100. The exact pressures, temperatures and densities adopted for our studies are shown in Table 1, and fall within the regime of the conditions tested during the experimental campaign. All the files and codes used to generate numerical solves can be found online at [https://github.com/ncywong/VKI\\_CFD](https://github.com/ncywong/VKI_CFD). We omit a detailed experimental-CFD validation here as prior comparisons of the  $SU^2$  suite on this blade are available in literature [18].

## 2.3 Key quantities of interest

Our interest in this paper lies in quantifying and controlling the performance variation associated with the manufacturing process. We focus on both scalar-valued (1D) quantities of interest (qoi) and vector-valued qois as functions of the profile geometry. The first performance metric we study is the stagnation pressure loss coefficient given by

$$Y_p = \frac{p_{02} - p_{01}}{p_{02} - p_2}, \quad (1)$$

where  $p_{02}$  is the circumferentially mass-averaged exit stagnation pressure. We also monitor the flow capacity or *exit mass flow function*

$$f_m = \frac{\dot{m} \sqrt{T_{01}}}{p_{01}} \times 10^4, \quad (2)$$

where  $\dot{m}$  is the mass flow rate and a scale factor of  $10^4$  is included for convenience. The units for  $f_m$  are  $\text{s} \cdot \text{m} \cdot \text{K}^{-1}$ . Our final qoi is the *isentropic Mach number distribution*, defined as

$$M(s) = \sqrt{\frac{2}{\gamma - 1} \left( \left( \frac{p_{01}}{p(s)} \right)^{\frac{\gamma - 1}{\gamma}} - 1 \right)}, \quad (3)$$

where  $p(s)$  is the local static pressure at location  $s$  on the surface. This profile qoi is discretized into  $N = 240$  discrete measurements, implying that the distribution can be defined by a *vector-valued* function,

$$\begin{aligned} \mathbf{M} &= [M(s_1), M(s_2), \dots, M(s_N)] \\ &= [M_1, M_2, \dots, M_N]. \end{aligned} \quad (4)$$

We create a DoE with  $K$  uniformly distributed random (Monte Carlo) samples generated within our design space  $\mathcal{X}$ . For each design vector generated, the mesh is parameterically deformed using the FFD parameters described above and run using the  $SU^2$  flow solver. The key qois are saved for each design generated, along with their corresponding design vectors.

## 3 STATISTICAL METHODOLOGIES

The input-output database from the CFD forms the bedrock of our methodology for prescribing blade envelopes. In this section, we detail some of the statistical techniques that will be applied to this database with the aim of extracting the control zone, the tolerance covariance, and for determining whether a scanned component should be used or scrapped.

### 3.1 Active and inactive subspaces

The formulation of a blade envelope relies on the ability to identify and generate blade designs that are invariant in one

or multiple objective(s). To this end, we partition the design space into two *subspaces*: the *active* subspace and the *inactive* subspace. Here a *subspace* is formed from a linear combination of all the parameters. The construction of these subspaces should satisfy the following: along the active subspace, qois vary significantly; while along the inactive subspace, qois are predominantly invariant to changes in the parameters. To clarify this breakdown, let  $\mathbf{Q} \in \mathbb{R}^{d \times d}$  be *any* orthogonal matrix, implying that  $\mathbf{Q}\mathbf{Q}^T = \mathbf{I}$ , where  $\mathbf{I}$  is the identity matrix. Now, consider the following decomposition

$$\begin{aligned} \mathbf{x} &= \mathbf{I}\mathbf{x} \\ &= \mathbf{Q}\mathbf{Q}^T \mathbf{x} \\ &= [\mathbf{W} \ \mathbf{V}] [\mathbf{W} \ \mathbf{V}]^T \mathbf{x} \\ &= \mathbf{W}\mathbf{W}^T \mathbf{x} + \mathbf{V}\mathbf{V}^T \mathbf{x}, \end{aligned} \quad (5)$$

where  $\mathbf{W} \in \mathbb{R}^{d \times r}$  and  $\mathbf{V} \in \mathbb{R}^{d \times (d-r)}$  are columns of  $\mathbf{Q}$ . The matrix  $\mathbf{V}$  can also be obtained by determining the null space of  $\mathbf{W}^T$ . The value of  $r$  used in practice will be dictated by how many *effective dimensions* one needs to capture the variation in the qoi(s). Ideally, rather than using *any* orthogonal matrix, our goal is to identify a specific  $\mathbf{Q}$  that affords a partition of the design space (5) such that  $r$  is as small as possible. In other words, the qoi(s) varies significantly along the directions  $\mathbf{W}^T \mathbf{x}$  but are invariant to changes in  $\mathbf{V}^T \mathbf{x}$ . So, how do we find a  $\mathbf{Q}$  such that its columns admit such a dimension-reducing structure?

Samarov [19] and Constantine [20] draw our attention to a covariance matrix that is the averaged outer product of the gradient

$$\mathbf{C} = \int_{\mathcal{X}} \nabla_{\mathbf{x}} f(\mathbf{x}) \nabla_{\mathbf{x}} f(\mathbf{x})^T \tau \, d\mathbf{x}, \quad (6)$$

where  $f$  represents the qoi, such as the loss coefficient, and  $\tau = 2^{-d}$  defines a uniform distribution over the  $d$ -dimensional hypercube  $\mathcal{X}$ . The symbol  $\nabla_{\mathbf{x}} f(\mathbf{x}) \in \mathbb{R}^d$  denotes the gradient of  $f$  with respect to its shape parameters  $\mathbf{x}$ . As  $\mathbf{C}$  is a symmetric matrix, one can write its eigendecomposition as

$$\mathbf{C} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T \quad (7)$$

where  $\mathbf{Q}$  is an orthogonal matrix of  $\mathbf{C}$ 's eigenvectors and  $\mathbf{\Lambda}$  a diagonal matrix of its eigenvalues. Indeed, *this is the  $\mathbf{Q}$*  we shall use to partition our design space as per (5). The value of  $r$  is set by the spectrum of eigenvalues in  $\mathbf{\Lambda}$ . We consider the separation

$$\mathbf{C} = [\mathbf{W} \ \mathbf{V}] \begin{bmatrix} \mathbf{\Lambda}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{\Lambda}_2 \end{bmatrix} \begin{bmatrix} \mathbf{W}^T \\ \mathbf{V}^T \end{bmatrix}, \quad (8)$$

where  $\mathbf{\Lambda}_1$  is a diagonal matrix containing the largest eigenvalues, and  $\mathbf{\Lambda}_2$  the smallest eigenvalues, both sorted in descending order. This partition should be chosen such that there is a large gap between the last eigenvalue of  $\mathbf{\Lambda}_1$  and the first eigenvalue of  $\mathbf{\Lambda}_2$  [20]. Thus, this partitioning of  $\mathbf{Q}$  yields the active subspace matrix  $\mathbf{W}$  and the inactive subspace matrix  $\mathbf{V}$ .

### 3.2 Sparse polynomials for gradient estimation

One challenge with the subspace idea detailed above is the necessity for gradient evaluations of the qoi (see (6)). Gradient evaluations can be efficiently computed using the adjoint method. However, automatic differentiation tools may not always be available, particularly in older industrial codes. In this paper, we estimate the gradients through a global polynomial approximation over the design space

$$f(\mathbf{x}) \approx \sum_{\mathbf{i} \in I} a_{\mathbf{i}} \Psi_{\mathbf{i}}(\mathbf{x}) \quad (9)$$

where the polynomial coefficients are denoted by  $a_{\mathbf{i}}$  and the basis terms are  $\Psi_{\mathbf{i}}(\mathbf{x})$ . These basis terms are formed from the product of univariate orthogonal polynomials and the *multi-index set*  $I$  in (9) dictates which  $O$  permutations of univariate polynomials of varying degrees are present. More specifically, one can think of each multivariate polynomial as

$$\Psi_{\mathbf{i}}(\mathbf{x}) = \prod_{j=1}^d \psi_{i_j}^{(j)}(x^{(j)}), \quad \text{where } \mathbf{i} = (i_1, \dots, i_d). \quad (10)$$

The superscript ( $j$ ) denotes the univariate polynomial along dimension  $j$ , which is only a function of the univariate parameter  $x^{(j)}$ . The degree for each univariate polynomial  $\psi_{i_j}^{(j)}$  is specified in the *multi-index*  $\mathbf{i}$ . All the different multi-indices lie in the finite multi-index set  $I$ , mentioned above. To qualitatively draw parallels between different multi-index sets, Figure 4 shows typical isotropic multi-index permutations for  $d = 2$  where the highest degree is restricted to  $p = 3$ . The number of squares in each of these subfigures corresponds to  $O$ , i.e., the number of coefficients that need to be computed. Selecting

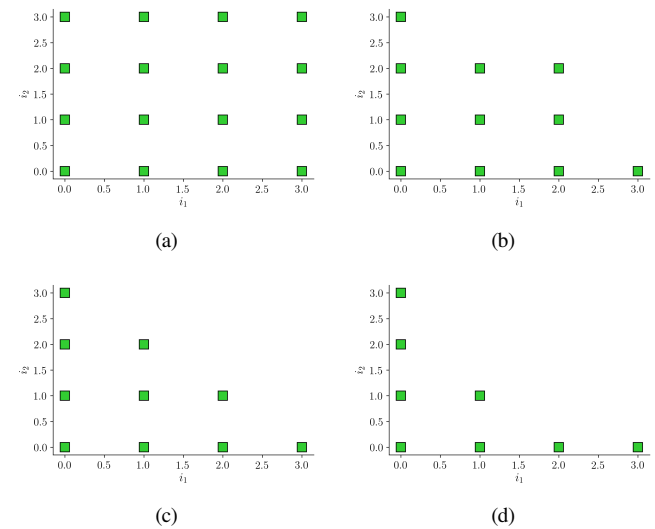


Fig. 4. Standard multi-index set  $I$  options: (a) tensorial; (b) Euclidean degree; (c) total order; and (d) hyperbolic basis for  $d=2$  and  $p=3$ .

the appropriate multi-index set is an important aspect of approximating multivariate functions via polynomials (see [21]). Decisions should be made depending on whether one would expect higher order interactions among the mixed terms or whether most of the large coefficients are associated with lower order principal terms. The total order basis, shown in Figure 4(c), is often a good compromise and ameliorates problems related to the exponential scaling in the number of associated coefficients with increasing  $d$ —an artifact of the tensorial basis in (a). The number of coefficients for a total order basis is given by

$$O = \binom{p+d}{p} = \frac{(p+d)!}{p!d!}. \quad (11)$$

Now, to understand the relationship between the number of coefficients  $O$  and the number of CFD evaluations  $K$  required to obtain an approximation, let the input-output qoi database be given by

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1^T \\ \vdots \\ \mathbf{x}_K^T \end{bmatrix}, \quad \mathbf{f} = \begin{bmatrix} f_1 \\ \vdots \\ f_K \end{bmatrix}, \quad (12)$$

where  $\mathbf{X} \in \mathbb{R}^{K \times d}$  is the matrix of inputs and  $\mathbf{f}$  is a vector of the qois across the  $K$  CFD evaluations. We can write the coefficients and discretized basis terms in (9) as

$$\mathbf{a} = \begin{bmatrix} a_1 \\ \vdots \\ a_o \end{bmatrix}, \quad \Psi = \begin{bmatrix} \Psi_1(\mathbf{x}_1) & \dots & \Psi_O(\mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ \Psi_1(\mathbf{x}_K) & \dots & \Psi_O(\mathbf{x}_K) \end{bmatrix}. \quad (13)$$

Once the coefficients  $\mathbf{a}$  are known, gradient approximations can be readily computed via

$$\nabla_{\mathbf{x}} f(\mathbf{x}) \approx \sum_{i \in I} a_i \nabla_{\mathbf{x}} \Psi_i(\mathbf{x}). \quad (14)$$

This idea of extracting gradients for estimating  $\mathbf{C}$  using a polynomial approximation was previously considered in [1] where the authors constructed a global least squares quadratic fit. The trouble with a least squares approach is that it requires  $K \geq O$ , i.e., we need as many CFD evaluations as basis terms (and ideally more), which can be prohibitive in high-dimensional problems.

To abate this computational cost, we propose the use of compressed sensing for estimating the coefficients; for an introductory background see [22]. This involves the solution to the optimization problem

$$\begin{aligned} & \underset{\mathbf{a}}{\text{minimize}} \quad \|\mathbf{a}\|_1 \\ & \text{subject to} \quad \|\Psi\mathbf{a} - \mathbf{f}\|_2 \leq \varepsilon, \end{aligned} \quad (15)$$

where the notation  $\|\cdot\|_1$  denotes the  $L_1$  norm, i.e., the sum of the absolute values of the arguments;  $\|\cdot\|_2$  denotes the familiar

$L_2$  or Euclidean norm, i.e., the square root of the sum of the squares of the argument. Here the small positive constant  $\varepsilon$  is used to account for the mis-match between the polynomial approximation and the true function evaluations  $\mathbf{f}$ .

Central to the compressed sensing paradigm is the notion of *sparsity* in the coefficients. In other words, to solve (15) we assume that a fraction of the coefficients are close to zero, although we do not need to know precisely which ones are. This enables us to solve for  $\mathbf{a}$  even when  $\Psi$  has more columns than rows, i.e., we have more unknown coefficients than CFD evaluations. More broadly, the use of compressed sensing for polynomial approximations has received considerable attention in literature [23], particularly for applications focused on uncertainty quantification.

To summarize, the salient advantage of solving (15) is the reduction in the number of CFD evaluations required. We can extract solutions for  $\mathbf{a}$  even when  $K \ll O$ , provided that  $\mathbf{a}$  is sparse. This is very useful as we would like to reduce the number of function evaluations required for estimating the gradients and subsequently the active and inactive subspaces.

### 3.3 Sampling from the inactive subspace

Equipped with a covariance matrix  $\mathbf{C}$  that facilitates dimension reduction, our goal is to generate blade profiles that will be provably invariant. One way to do this is to generate samples from the inactive subspace  $\mathbf{V}$  for a constrained value along the active subspace  $\mathbf{W}$ . Building upon (5), we write

$$\begin{aligned} \mathbf{x} &= \mathbf{W}\mathbf{W}^T \mathbf{x} + \mathbf{V}\mathbf{V}^T \mathbf{x} \\ &= \mathbf{W}\mathbf{u} + \mathbf{V}\mathbf{z}, \end{aligned} \quad (16)$$

where  $\mathbf{u} := \mathbf{W}^T \mathbf{x}$  is the active coordinate and  $\mathbf{z} := \mathbf{V}^T \mathbf{x}$  the inactive coordinate. Next, the task is to find points  $\mathbf{x}$  with a fixed  $\mathbf{u}$  but different  $\mathbf{z}$ , while obeying the constraint  $-\mathbf{1} \leq \mathbf{x} \leq \mathbf{1}$ . The bounds imply that


$$\begin{aligned} \mathbf{V}\mathbf{z} &\leq \mathbf{1} - \mathbf{W}\mathbf{u}, \\ -\mathbf{V}\mathbf{z} &\leq \mathbf{1} + \mathbf{W}\mathbf{u}. \end{aligned} \quad (17)$$

These linear constraints define a *polytope* in the inactive subspace (see Figure 5 I).


To generate samples from this polytope, we adopt the hit-and-run method [24] based on the implementation by Constantine et al. [25]. In this method, one first identifies a feasible point by locating the Chebyshev centre of the polytope<sup>1</sup> at the specified active coordinate  $\mathbf{u}$ , by solving a linear program. Then, starting from the Chebyshev centre, a random direction is selected and, in this direction, the longest line segment that lies within the polytope (and goes through the centre) is identified. Following this, a point on this line segment is selected at random yielding the first sample. This procedure is repeated by starting with a new random direction and identifying the longest line segment along that line. These steps are captured in Figure 5 II to VI.

<sup>1</sup>The Chebyshev centre is the centre of the largest hypersphere that can fit within the polytope (see page 418 in [26]).


I. Find the Chebyshev centre of the polytope, which is the centre of the largest inscribed circle




II. Find a random direction and the longest line segment within the polytope along it passing through the Chebyshev centre




III. Sample a random point on this line segment



IV. Find a random direction and the longest line segment within the polytope along it passing through the previous point



V. Sample a random point on this line segment



VI. Repeat until enough points are sampled




Fig. 5. Schematic showing the hit-and-run sampling method for generating samples in the inactive subspace.

### 3.4 Machine learning blade envelopes

The statistical methodologies presented thus far set the stage for generating both the *control zone* and the *tolerance covariance*—the two key elements of the blade envelope. To formalize these two, let  $H$  be the number of inactive samples generated via the hit-and-run method for the nominal value of  $\mathbf{u}$  at the origin—corresponding to the undeformed geometry. Furthermore, let  $\mathbf{s} \in \mathbb{R}^N$  denote the vertical coordinates of each airfoil corresponding to these inactive samples, where we assume that the horizontal coordinates are the same over all the samples. In other words, the airfoil coordinates  $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_H\}$  represent designs that are invariant in the selected  $\text{qoi}(\mathbf{s})$ . The bounds of the control zone are then specified by the upper and lower limits  $\mathbf{c}_u, \mathbf{c}_l \in \mathbb{R}^N$  where

$$\begin{aligned} \mathbf{c}_u &= \text{maximum} \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_H\}, \\ \mathbf{c}_l &= \text{minimum} \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_H\}, \end{aligned} \quad (18)$$

where the maximum and minimum are taken coordinate-wise. The control zone by itself is very useful to offer some insight into where tolerances need to be tightened and where they may potentially be relaxed.

The *tolerance covariance* matrix can notionally be interpreted as a constraint on the curvature variation along both the suction and pressure sides of the airfoil, as mentioned before. In Figure 2 we show numerous  $\mathbf{s}$  profiles to capture what this covariance structure represents. More formally, we define it as the matrix  $\mathbf{S} \in \mathbb{R}^{N \times N}$  where each element is given by

$$\mathbf{S}_{ij} = \text{cov}(s^{(i)}, s^{(j)}), \quad (19)$$

where the superscript  $(i)$  denotes the  $i$ -th discrete measurement coordinate.

As stated earlier, when a manufactured component is scanned, we wish to determine whether it complies with a computed blade envelope for the design in question. This necessitates a machine learning paradigm where we can map the CMM readings of a blade onto its blade envelope and distill a binary scrap-or-use output. We will assume that, although the number of CMM points may be less than  $N$ , one can linearly interpolate them to estimate the vertical displacements (or surface normal displacements) at the same horizontal coordinates (or chord-wise locations) as the blade envelope. Alternatively, one can do the opposite and downsample the coordinates associated with the blade profile, which are obtained solely from CFD.

To do this, we propose calculating the Mahalanobis distance between the scanned blade and the distribution associated with the blade envelope. Although we can generate infinitely many airfoil profiles (see Figure 2), we do not know the true distribution from which these samples are obtained. However, we can still estimate the *ensemble mean*

$$\boldsymbol{\mu} = \frac{1}{H} \sum_{i=1}^H \mathbf{s}_i. \quad (20)$$

To determine whether a new scanned blade  $\tilde{\mathbf{s}}$  lies within the blade envelope, we compute the Mahalanobis distance

$$\zeta(\tilde{\mathbf{s}}) = \sqrt{(\tilde{\mathbf{s}} - \boldsymbol{\mu}) \mathbf{S}^{-1} (\tilde{\mathbf{s}} - \boldsymbol{\mu})^T}. \quad (21)$$

If the distance is very small, then the new profile  $\tilde{\mathbf{s}}$  is said to lie within the blade envelope and is thus deemed acceptable. On the other hand, if the distance is large, the new profile clearly lies far away from the distribution associated with the blade envelope and therefore should not be used.

We can convert this distance—a relative measure of acceptance and rejection—to an absolute one via a logistic function

$$g(\tilde{\mathbf{s}}) = \frac{\beta_1}{1 + \exp\{-\beta_2(\zeta(\tilde{\mathbf{s}}) - \beta_3)\}} \quad (22)$$

where model parameters  $\beta_1, \beta_2$  and  $\beta_3$  have to be determined by a training-testing criterion. Generating training and testing samples requires no additional CFD evaluations, and the parameters themselves can be estimated via standard gradient-based

optimization. Once determined for a specific blade, expression (22) can be used on the factory floor to render a scrap-or-use judgment for a manufactured blade.

We remark that here that one alternative to training the logistic function in (22) is to estimate the critical values of the Mahalanobis distance, given many scanned blades. As these distance metrics yield a chi-squared distribution, one can assign a *significance level* to the distribution to deliver a scrap-or-use verdict on all the blades scanned.

### 3.5 Summarizing the key steps

We conclude this section with a summary of the main techniques involved in generating a blade envelope and exploiting it to query whether a scanned blade should be used; these steps are captured in Figure 6.

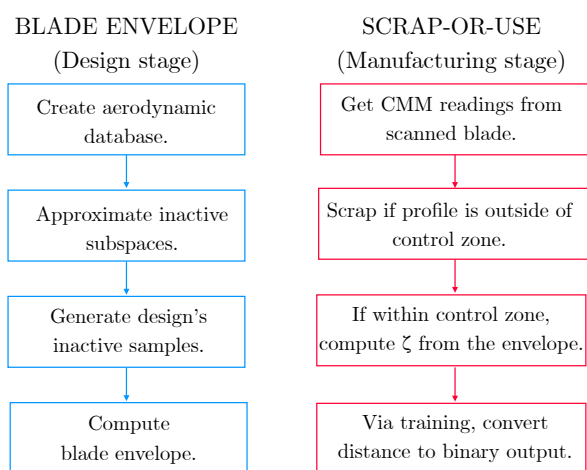


Fig. 6. Flowchart summarizing the key steps involved in generating a blade envelope and querying it to ascertain whether a scanned component should be used.

For generating a blade envelope, the first requirement is a well-sampled aerodynamic database that comprises a collection of geometries around the baseline profile and the output performance of each geometry (2.3). Using this input-output database, we can approximate the inactive subspaces (3.1 and 3.2) and generate sample designs that will admit the same performance, but have very different geometries (3.3). By following this workflow, we can generate the blade envelope, shown in Figure 2.

The mean and covariance of the *inactive samples* associated with this blade envelope can then be used to ascertain whether a scanned blade’s CMM coordinates are likely to lie within this envelope (3.4). A logistic function can be trained to yield a binary result that can be deployed in various blade manufacturing sites.


## 4 BLADE ENVELOPES FOR 1D QUANTITIES

In this section, we use the techniques discussed above to identify and sample from the inactive subspace associated with loss. Our focus is to deliver a scrap-or-use decision for blade


profiles that could have been manufactured for the LS89 profile, if near-design loss was the only performance consideration.

### 4.1 Inactive subspace for loss

To compute the inactive subspace, we must first find the active subspace. Our DoE comprises 1000 randomly generated designs in  $d=20$  which we split into  $M=800$  training and 200 testing samples. For each design, we store the CFD yielded loss, mass flow function and isentropic Mach number, although in this paper we focus our attention on the loss only. We fit a global sparse polynomial with order  $p=3$  that has a total of  $O=1771$  basis terms (see (11)). We use the Effective Quadratures package [27] (see <https://www.effective-quadratures.org>) to construct the sparse polynomial approximation. The resulting coefficients are shown in Figure 7(a), where an automated value of  $\epsilon=0.001$  was found to be suitable to solve (15). Deploying our model on the testing data-set, we obtain an  $R^2$  value of 0.957, which gives us confidence in its predictive capabilities (see Figure 7(b)). As mentioned in 3.2, we can use




(a)




(b)

Fig. 7. Global sparse polynomial approximation: (a) sorted coefficients; (b) validation of the model using the testing data.

this approximation to estimate the gradients and thus construct



(a)



(b)

Fig. 8. Active subspace computation: (a) eigenvalues of the covariance matrix for loss; (b) sufficient summary plot of CFD values of loss on its first active coordinate.

the covariance matrix associated with active subspaces for loss  $\mathbf{C}_{Y_p}$ . Computing its eigendecomposition yields

$$\mathbf{C}_{Y_p} = [\mathbf{W}_{Y_p} \ \mathbf{V}_{Y_p}] \begin{bmatrix} \Lambda_{Y_p,1} & \mathbf{0} \\ \mathbf{0} & \Lambda_{Y_p,2} \end{bmatrix} \begin{bmatrix} \mathbf{W}_{Y_p}^T \\ \mathbf{V}_{Y_p}^T \end{bmatrix}. \quad (23)$$

We report the eigenvalues of this covariance matrix in Figure 8(a) and, based on this decay, set  $\mathbf{W}_{Y_p} \in \mathbb{R}^{d \times 1}$ . In other words, we define our active subspace for loss as being spanned by the first column in  $\mathbf{W}_{Y_p}$ , and thus its remaining 19 columns span the inactive subspace for loss. For completeness, we plot the loss for all 1000 geometries on the *active* coordinate of  $\mathbf{W}_{Y_p}$  in Figure 8(b). Here we define the active coordinate corresponding to design  $\mathbf{x}$ ,

$$u_{Y_p,1} = \mathbf{x}^T \mathbf{W}_{Y_p}[:,1], \quad (24)$$

where the notation  $[:,1]$  denotes the first column of the matrix.




Fig. 9. Loss invariant designs generated by sampling from the inactive subspace.

## 4.2 Generating training and testing samples

Endowed with the inactive subspace matrix for loss  $\mathbf{V}_{Y_p} \in \mathbb{R}^{d \times 19}$ , we can now generate infinitely many blade designs that have approximately the same loss as the nominal LS89. The degree of approximation is dictated entirely by the cut-off value we assume for partitioning the eigenvalues. Here, we generate  $H = 5000$  new geometries using the hit-and-run sampling strategy detailed in 3.3. We plug in the matrices obtained for  $\mathbf{W}_{Y_p}$  and  $\mathbf{V}_{Y_p}$  into the bounds in (17) and set  $u = 0$ , which corresponds to the *active coordinate* associated with the nominal blade profile (a scalar since the active subspace is 1-dimensional). The values for  $\mathbf{z}$  obtained via hit-and-run sampling are then used to arrive at design vectors  $\mathbf{x}$ . To verify that these 5000 additional geometries have similar  $Y_p$  losses compared to the datum, we ran 500 of these designs through the CFD solver. The results are reported in Figure 9 and shown alongside the original training data used to fit the sparse polynomial. This graph gives us confidence that our methodology has successfully generated designs that have comparable loss values to the LS89. Moreover, it also serves to show that our global sparse polynomial approximation offers an acceptable characterization of the loss. One important point to emphasize here is that the only computational overhead associated with generating 5000 geometries is the cost of deforming the nominal mesh and generating a new mesh, and thus airfoil coordinates. While it is recommended to run a few CFD evaluations to confirm the veracity of the subspaces, 500 is perhaps two orders of magnitude more than required in practice.

The steps laid out thus far permit us to generate a blade envelope for loss; see Figure 10. Any blade that was manufactured within the control zone that satisfies the tolerance covariance (visualized by the LE displacement but quantified by the tolerance covariance matrix  $\mathbf{S}$  (19)) will admit loss values comparable to the nominal design. Note that this envelope has much greater variability compared to the one shown in Figure 2. This is because in Figure 2, in addition to loss, the mass flow function is also constrained. We shall study how one generates samples with multiple objectives in the companion Part II paper [15].






Fig. 10. The blade envelope for loss.

### 4.3 Use or scrap?

Following the flowchart in Figure 6, we compute the ensemble mean  $\mu$  and tolerance covariance matrix  $\mathbf{S}$  associated with the different airfoils sampled from the inactive subspace. Our choice of  $H = 5000$  was dictated by the convergence of the two aforementioned statistical quantities in the  $L_2$  norm sense. Figure 11 shows the tolerance covariance matrix  $\mathbf{S} \in \mathbb{R}^{240 \times 240}$ . To correctly obtain these quantities, each airfoil profile has been carefully normalized to ensure that the horizontal coordinates across the profiles are the same, and thus only the vertical displacements, i.e., thickness / chord, are used in estimating  $\mu$  and  $\mathbf{S}$ . With these






Fig. 11. The tolerance covariance matrix  $\mathbf{S}$  generated with  $H = 5000$  inactive samples.

metrics computed, the Mahalanobis distance between any new scanned geometry and the point cloud associated with the inactive samples can be determined. We compute this distance for the 500 out of the 5000 airfoils generated from the inactive subspace and plot them in Figure 12(a). After pruning away the geome-



(a)



(b)

Fig. 12. Machine learning blade envelopes: (a) Mahalanobis distances  $\zeta$  reported for numerous blades along with their losses; (b) Conversion to a binary scrap-or-use decision where 1 indicates use and 0 indicates scrap.

tries that lie outside of the control zone, we do the same for our original 1000-point training-testing database, for which most of the airfoils admit different values of  $Y_p$ . A clear divergence in the loss values is observed for  $\zeta$  values greater than 7; all the loss invariant designs have low values of  $\zeta$ , illustrating the utility of our overall methodology. To propagate a degree of uncertainty in our assessment, we assume that designs with  $\zeta$  values between 3.5 and 7 lie within a  $\zeta$  buffer region, also demarcated in the figure.

We use this zone to tune our logistic function and arrive at the following values:  $\beta_1 = 1, \beta_2 = 5$  and  $\beta_3 = 3$ ; and illustrate the result in Figure 12(b). Here the horizontal axis plots  $\zeta$ , while the vertical axis represents the binary outcome: a value of 1 implies that the design can be used, while a value of 0 implies that the design should be scrapped. A value in-between—i.e., falling within the buffer region—indicates a geometry that requires further examination before a final scrap-or-use decision.

### 4.4 Evaluating designs from other design spaces

To put our methodology through its paces, we generate 1000 new airfoil profiles from a different design space, one comprising




Fig. 13. Sample designs from the larger design space of FFD design variables.




Fig. 14. Application of the logistic model on profiles generated from the larger design space with  $d = 30$  design variables. The values on the colorbar are CFD yielded  $Y_p$  values.

$d = 30$  FFD design variables, as shown in Figure 13. We evaluate the loss of these new profiles by running them through the  $SU^2$  suite. Then, for each profile that lies in the control zone, we compute  $\zeta$  using our previously stored values of  $\mu$  and  $S$ —obtained from the 20D design space. These distances are then fed through the aforementioned logistic model and plotted along the vertical axis in Figure 14. Here the colors of each marker correspond to the CFD computed loss value. These results attest to the notion that a blade envelope’s utility in distinguishing between tolerance-abiding airfoils is not restricted to a particular parameterization. Moreover, the fact that most of the blades with comparable loss values to the LS89 fall either below or within the  $\zeta$  buffer region gives us confidence that this workflow can be beneficial and transformative when deployed on real manufactured scans.

## 5 CONCLUSIONS

This paper introduces the concept of blade envelopes and details its underlying methodology. Blade envelopes can be thought of as a blade tolerance draft—a visual manufacturing guide—that instructs blade manufacturers and designers what tolerances can be used and whether a used component satisfies these tolerances or not.

## Acknowledgements

CYW acknowledges financial support from the Cambridge Trust, Jesus College, Cambridge, and the Data-Centric Engineering programme of The Alan Turing Institute. PS was funded through a Rolls-Royce research fellowship and AS was funded through the Digital Twins in Aeronautics grant as part of the Strategic Priorities Fund EP/T001569/1. ABD was supported by the Lloyds Register Foundation Programme on Data Centric Engineering and by The Alan Turing Institute under the EPSRC grant EP/N510129/1. Thanks are due to Bryn Noel Ubald for his help in generating some of the figures, and to Raúl Vázquez for his support of this work.

## References

- [1] Seshadri, P., Shahpar, S., Constantine, P., Parks, G., and Adams, M., 2018. “Turbomachinery Active Subspace Performance Maps”. *Journal of Turbomachinery*, **140**(4), Apr., p. 041003.
- [2] Goodhand, M. N., and Miller, R. J., 2012. “The Impact of Real Geometries on Three-Dimensional Separations in Compressors”. *Journal of Turbomachinery*, **134**(2), Mar., p. 021007.
- [3] Garzón, V. E., and Darmofal, D. L., 2003. “Impact of Geometric Variability on Axial Compressor Performance”. *Journal of Turbomachinery*, **125**(4), Oct., pp. 692–703.
- [4] Goodhand, M. N., Miller, R. J., and Lung, H. W., 2014. “The Impact of Geometric Variation on Compressor Two-Dimensional Incidence Range”. *Journal of Turbomachinery*, **137**(2), Sept., p. 021007.
- [5] Bradshaw, S., and Waitz, I., 2009. “Impact of Manufacturing Variability on Combustor Liner Durability”. *Journal of Engineering for Gas Turbines and Power*, **131**(3), May, p. 032503.
- [6] Duffner, J. D., 2008. “The effects of manufacturing variability on turbine vane performance”. Master’s thesis, Aerospace Computational Design Laboratory, Dept. of Aeronautics & Astronautics, Massachusetts Institute of Technology, Cambridge, MA, June.
- [7] Fathi, A., and Alizadeh, M., 2012. “Effects of Blade Manufacturing Deviations on Turbine Performance”. In ASME 2012 Gas Turbine India Conference, American Society of Mechanical Engineers, pp. 203–211.
- [8] Wang, X., and Zou, Z., 2019. “Uncertainty analysis of impact of geometric variations on turbine blade performance”. *Energy*, **176**, June, pp. 67–80.
- [9] Shi, W., Chen, P., Li, X., Ren, J., and Jiang, H., 2019. “Uncertainty Quantification of the Effects of Small

- Manufacturing Deviations on Film Cooling: A Fan-Shaped Hole”. *Aerospace*, **6**(4), Apr., p. 46.
- [10] Kamenik, J., Voutchkov, I., Toal, D. J. J., Keane, A. J., Högner, L., Meyer, M., and Bates, R., 2018. “Robust Turbine Blade Optimization in the Face of Real Geometric Variations”. *Journal of Propulsion and Power*, **34**(6), pp. 1479–1493.
- [11] Seshadri, P., Shahpar, S., and Parks, G. T., 2014. “Robust Compressor Blades for Desensitizing Operational Tip Clearance Variations”. In ASME Turbo Expo 2014: Turbine Technical Conference and Exposition, American Society of Mechanical Engineers Digital Collection.
- [12] Lange, A., Voigt, M., Vogeler, K., and Johann, E., 2012. “Principal component analysis on 3D scanned compressor blades for probabilistic CFD simulation”. In 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, American Institute of Aeronautics and Astronautics.
- [13] Dow, E., and Wang, Q., 2013. “Output Based Dimensionality Reduction of Geometric Variability in Compressor Blades”. In 51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, American Institute of Aeronautics and Astronautics.
- [14] Lobato, H. M. P., Gower, I. R., Powell, C. A., and Orchard, N. B., 2014. United States Patent: 8718975 - Surface profile evaluation, May.
- [15] Wong, C. Y., Seshadri, P., Scillitoe, A., Ubald, B. N., Duncan, A., and Parks, G. “Blade Envelopes Part II: Multiple Objectives and Inverse Design”. *Submitted to Journal of Turbomachinery*.
- [16] Arts, T., and de Rouvroit, M. L., 1992. “Aero-thermal performance of a two-dimensional highly loaded transonic turbine nozzle guide vane: A test case for inviscid and viscous flow computations”. *Journal of Turbomachinery*, **114**(1), pp. 147–154.
- [17] Economon, T. D., Palacios, F., Copeland, S. R., Lukaczyk, T. W., and Alonso, J. J., 2016. “SU2: An Open-Source Suite for Multiphysics Simulation and Design”. *AIAA Journal*, **54**(3), Mar., pp. 828–846.
- [18] Vitale, S., Gori, G., Pini, M., Guardone, A., Economon, T. D., Palacios, F., Alonso, J. J., and Colonna, P., 2015. “Extension of the SU2 open source CFD code to the simulation of turbulent flows of fluids modelled with complex thermophysical laws”. In 22nd AIAA Computational Fluid Dynamics Conference, p. 2760.
- [19] Samarov, A. M., 1993. “Exploring Regression Structure Using Nonparametric Functional Estimation”. *Journal of the American Statistical Association*, **88**(423), Sept., pp. 836–847.
- [20] Constantine, P. G., 2015. *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*. Society for Industrial and Applied Mathematics, Philadelphia, PA, Mar.
- [21] Seshadri, P., Narayan, A., and Mahadevan, S., 2017. “Effectively subsampled quadratures for least squares polynomial approximations”. *SIAM/ASA Journal on Uncertainty Quantification*, **5**(1), pp. 1003–1023.
- [22] Bryan, K., and Leise, T., 2013. “Making do with less: An introduction to compressed sensing”. *SIAM Review*, **55**(3), pp. 547–566.
- [23] Peng, J., Hampton, J., and Doostan, A., 2014. “A weighted  $\ell_1$ -minimization approach for sparse polynomial chaos expansions”. *Journal of Computational Physics*, **267**, pp. 92–111.
- [24] Smith, R. L., 1996. “The Hit-and-run Sampler: A Globally Reaching Markov Chain Sampler for Generating Arbitrary Multivariate Distributions”. In Proceedings of the 28th Winter Conference on Simulation, WSC '96, IEEE Computer Society, pp. 260–264.
- [25] Constantine, P., Howard, R., Glaws, A., Grey, Z., Diaz, P., and Fletcher, L., 2016. Python Active-subspaces Utility Library, Sept.
- [26] Boyd, S., and Vandenberghe, L., 2004. *Convex Optimization*. Cambridge University Press, Cambridge, UK.
- [27] Seshadri, P., and Parks, G., 2017. “Effective-quadratures (EQ): Polynomials for computational engineering studies”. *The Journal of Open Source Software*, **2**, p. 166.