

Robust Turbine Blade Optimization in the Face of Real Geometric Variations

Jan Kamenik*, Ivan Voutchkov†, David J. J. Toal‡, Andy J. Keane§,

University of Southampton, Southampton, SO16 7QF, UK

Lars Högner¶

TU , Dresden, 01062, Germany

Marcus Meyer||

Rolls-Royce Deutschland Ltd & Co KG, Blankenfelde-Mahlow, 15827, Germany

Ron Bates**

Rolls-Royce plc, Derby, DE24 8BJ, UK

Due to manufacturing variations, no real turbine blade exactly conforms to its nominal geometry. Even minimal deviations are known to affect aerodynamic performance, blade temperatures and blade lifespan negatively. Rather than conventional deterministic design with its costly adherence to strict control of tolerance limits, robust design optimization aims to incorporate inevitable variations into the design process itself, so that both performance mean and scatter can be optimized simultaneously. Such a workflow is presented and applied in this paper to aerodynamically optimize an industrial turbine rotor blade against realistic manufacturing variations. A set of digitized 3D laser scans from two turbofan engines forms the core of this study. On the basis of these deviations, the approach uses high-fidelity geometric models, non-intrusive uncertainty quantification and efficient robust optimization with constraints to effectively locate Pareto-optimal designs. One selected robust blade is validated and shown to be desensitized to the observed manufacturing variability. The underlying measurement data is crucial to obtain realistic results and as a consequence vital to design real robust turbine blades.

*Doctoral Student, Rolls-Royce University Technology Centre, jkamenik@soton.ac.uk

†Senior Research Fellow, Rolls-Royce University Technology Centre, i.i.voutchkov@soton.ac.uk

‡Lecturer, Rolls-Royce University Technology Centre, d.j.j.toal@soton.ac.uk

§Professor, Rolls-Royce University Technology Centre, andy.keane@soton.ac.uk

¶Research Associate, Chair of Turbomachinery and Flight Propulsion, lars.hoegner@tu-dresden.de

||CFD Methods Specialist, ET-DSE & CFD Methods, marcus.meyer@rolls-royce.com

**Corporate Specialist - Design Sciences, Engineering Capability, ron.bates@rolls-royce.com

Nomenclature

COI	Coefficient of importance
C_v	Coefficient of variation
c_p, c_v	Specific heat capacity at constant pressure, volume
\mathbf{d}	Design variable vector
E	Expected value
h	Optimal kernel bandwidth
\dot{m}	Mass flow rate
\mathbf{P}	Delta parameter sample matrix
p	Pressure
PCE	Percentage error
R^2	Coefficient of determination
R_s	Specific gas constant
r_s	Spearman's rank correlation coefficient
\hat{s}	Sample standard deviation
T	Temperature
Var	Variance
w	Specific work
\bar{x}	Sample mean
$\Delta\bar{\mathbf{p}}$	Vector with radially averaged delta parameters
η	Isentropic efficiency
θ, q, λ	Kriging hyperparameters
κ	Heat capacity ratio
ν	Kinematic viscosity
$\tilde{\nu}$	Spalart variable
ρ	Density
ϕ	Inlet capacity
Ψ	Kriging correlation matrix

I. Introduction

No turbine blade ever exactly conforms to its nominal design intent geometry and almost never operates at the nominal conditions it was initially designed for. For example, deviations can stem from manufacturing variations, corrosion, foreign object damage, creep, fatigue and operating condition abnormalities such as hot streaks. In the literature, it is well-known that even minimal perturbations in certain areas, whatever the cause, can have substantial detrimental effects on blade performance [1, 2, 3, 4]. For turbines in particular, this concerns the capacity, pressure losses, blade temperatures and as a consequence thereof blade life. This motivates the need for uncertainty quantification (UQ) and robust design optimization (RDO) with the goal to design turbine blades, which operate well despite unavoidable variations. In the past, the problem of manufacturing variations to some degree was addressed by safety factors and strict control of tolerances, which, however, can quickly lead to over-engineering and increased costs [5]. In contrast, the fundamental idea here is to establish a non-deterministic methodology to create robust turbine blade *by design*, which is vital in moving away from suboptimal safety factors and the expensive practice of tightening quality control tolerance limits.

The contribution of this paper relative to prior work is the incorporation of a previously developed parameterization method [2] and UQ workflow [6] into a coherent RDO workflow for turbine blades, which, to the best of our knowledge, is the first such heuristic to incorporate real variations of the entire blade shape. It is shown that artificially generated blades for UQ can recreate the stochastic performance of real scans and that the optimization effectively produces robust turbine blade designs. The optimization strategy also improves existing techniques, can handle constraints and has been integrated into Rolls-Royce's optimization tool OPTIMAT v2 ^a. The difference to similar surrogate-based robust optimization methods is that UQ is carried out on a single surrogate and the statistics are passed straight to the optimizer. The construction of another layer of surrogates for mean and standard deviation is not required. For details, see Sec. III. Moreover, as an improvement to related work on compressor blades, where deviations were approximated with a truncated multivariate normal distribution [8], probability density functions (PDFs) of random pa-

^aOPTIMAT v2 is a surrogate-based optimization tool developed at the University of Southampton's Rolls-Royce UTC [7].

parameters are estimated nonparametrically through kernel density estimation (KDE). This is more flexible and allows the incorporation of random variables with any distribution.

The methodology in this paper is significant because it permits robust turbine blade design with realistic geometric variations so that blades can be optimized and desensitized towards real measured variations and as a consequence perform better. This is preferable to robust optimization with made-up and possibly unrealistic variations based on assumptions. The industrial relevance stems from the fact that every engine manufacturer faces adverse effects due to manufacturing variations in one way or another, and the workflow presents a way to tackle this. It is efficient due to its use of fully-featured high-fidelity geometric models, 3D Reynolds-averaged Navier-Stokes (RANS) equation-based computational fluid dynamics (CFD), easy-to-use non-intrusive UQ and its use of parallel computing. It can also handle large numbers of parameters, which can be useful for potential large-scale industrial applications.

II. Background

A. Related Work

Several relevant UQ studies describe the effects of turbine blade deviations based on actual manufacturing process-related data [1, 2, 3, 4, 9]. Notably, a previous version of the correlation-based parameterization method also used in this paper was used as part of a UQ study [2]. The statistical basis consisted of 500 digitized new and used high-pressure (HP) turbine blades. With a 3D RANS-based CFD model of an HP turbine stage, a sensitivity study based on 50 MC samples was carried out and the stagger angle was found to be the single most important parameter for stage capacity and reaction. The effects of manufacturing variability on the aerodynamic performance of turbine vanes have also been investigated using 2D MISES CFD and principal component analysis (PCA) by Duffner [3]. It was found that the regions most sensitive to perturbations were the throat and trailing edge (TE). Upstream of the throat, performance was found to be nearly completely insensitive to perturbations. Changes in the minimum throat area and shock waves were identified as the main underlying mechanisms. Another study of a Rolls-Royce HP nozzle guide vane (NGV) explored the influence of geometric variations on capacity settings using parametric

NGV TE slots [4]. A one-factor-at-a-time analysis and MC simulations were used and TE position and the suction side (SS) profile offset were found to impact capacity significantly. In another UQ study, manufacturing data was distilled into tolerances for four parameters: ± 0.3 mm for the chord length, $\pm 20'$ (arcminutes) for the stagger angle, -0.1 mm for the TE radius and ± 0.15 mm for the blade height [9]. This study identified the stagger angle and TE variations to impact the performance most significantly.

In general, RDO is not yet widely adopted in industry and subject to a lot of current research [8, 10, 11, 12, 13]. Past studies often lack real manufacturing data, and probability distributions for uncertain parameters are often made-up. For example, uniform probability distributions were assumed for a centrifugal impeller optimization [11]. Similarly, a β -distribution was assumed for compressor tip clearance variations [14]. The problem is that in reality probability distributions could be entirely different, e.g., in terms of range, shape and correlation, but the variability model is crucial for the final design and its performance. There also have been studies with underlying variability models based on measurements. In one such study of compressor blades, probability distributions and correlations were approximated with a truncated multivariate normal PDF and each blade was analyzed with three 2D MISES solver runs [8]. The study hence relied on the assumption that each parameter's marginal distribution is also normally distributed. The restriction to 2D CFD was chosen due to the long CFD runtimes, but three-dimensional effects in the flow were neglected. The method presented in this paper improves upon both of these aspects. In another study, a 3D RANS-based robust optimization of a fully-featured turbine blade with PDFs based on over 100 measurements was carried out [10]. But besides three operational uncertainties, this study only considered the leading edge (LE) film cooling row and the tip gap as the sole two geometric noise parameters as the main focus was the inclusion of conjugate heat transfer to compute the actual metal temperatures.

Various approaches to robust optimization have been compared by Keane [12, 13]. In particular, it was shown that co-Kriging can be very efficient, but also that there is no universal method of choice. Co-Kriging or other multi-fidelity approaches are typically only options when high- and low-fidelity model outputs are well correlated and otherwise cannot be used. Surrogate model-

assisted optimization is typically a suitable choice, whereas direct NSGA-II (Non-dominated Sorting Genetic Algorithm II [15]) searches involving a vast number of function evaluations in conjunction with long CFD runtimes are often prohibitively expensive in practice. Other often-used UQ techniques such as polynomial chaos-based surrogate models are typically restricted to problems with only a few variables [16]. Sparse grid methods, although highly efficient for smooth functions, cannot cope with potentially noisy, discontinuous or missing data points, e.g., due to diverged CFD or parameterization failures. It is also worth noting that robust optimization can be construed differently, that there is no universal definition for robustness and that different statistics can be used as objectives. For example, minimizing expectation and dispersion, expectation and nominal value or nominal value and dispersion are three possible formulations [17]. By far the most frequent form is to minimize expected value and standard deviation, which is also used here. Hence, the stochastic objectives are to simultaneously minimize the sample mean and the unbiased sample standard deviation in Eq. (1):

$$\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i \quad \text{and} \quad \hat{s} = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2} \quad (1)$$

B. Optical Scans & Blade Variations

Fifty-eight HP turbine blade scans in the form of STL files (stereolithography file format) were provided by Rolls-Royce Germany. Although the blades have been in service briefly with two different airlines and were taken from *used, low-cycle* engines, they essentially represent manufacturing variations. In a previous work, all blades were digitized using the scanning system GOM ATOS SO 4M, a high-resolution, structured-light 3D scanner [6]. Due to their significance for this study, some facts from this reference are repeated here: The measurement accuracy of the scan system according to the manufacturer is 5 μm , but previous internal investigations with a test specimen revealed that the real observed measurement accuracy is around 8 μm for planar surfaces and about 27 μm near strongly curved areas. The blades were aligned to the nominal part using the upper firtree surfaces due to their contact with rotating disk surfaces during operation. This ensures the blades are realistically aligned in accordance with their actual operating state. Any measured

deviation is thus a superposition of a blade's positional variation, shape variation and variation of the selected firtree surfaces. Also, gaps from scanned cooling holes are removed from each scan using purpose-built in-house software from TU Dresden [18]. Doing this ensures that complete airfoil sections are available, which are required for the parametrization to work. With smoothed patches approximating the local blade curvature, the application of this tool introduces negligible effects for the results.




Figure 1 High-pressure turbine blade CAD model, scan and reconstruction.

As an illustration of one of the scanned HP turbine blades, Fig. 1b depicts the STL file of scan #1 with about 2 million triangles. This scan can be compared with the nominal CAD model, the design intent geometry, depicted in Fig. 1a. Figure 1c shows a reverse engineered NX model of the blade geometry for CFD analysis.

To illustrate some observed deviations from the nominal CAD model, Fig. 2 depicts the geometric variations of all 58 scans' leading and trailing edges at 50% span. In Fig. 2b near the TE, axially shifted and twisted airfoil shapes are striking. One outlier with a particularly short TE stands out. The manufactured mean is clearly different from the nominal CAD design intent.




Figure 2 Variations of all 58 blade scans (mean - -) vs. the nominal CAD design intent (—) at half-span.

III. Computational Methodology

The four basic robust optimization steps can be broken down as follows:

1. Construction and evaluation of a design of experiments (DOE) comprising the design *and* noise parameter space.
2. Tuning of an ordinary Kriging surrogate model.
3. UQ through sampling the surrogate (MC or quasi-Monte Carlo (QMC) sampling) with Eq. (1).
4. Multi-objective optimization of mean and variance using NSGA-II [15].

One part of this robust optimization methodology is novel and differs from existing approaches, e.g., those due to Lee and Park [19] or Dellino et al. [20]. Here, one single Kriging model for the combined design and noise parameters is constructed and then sampled. In both references, however, an intermediate Kriging model is built first, which is then sampled to construct another layer of Kriging models for mean and variance. Here, mean and standard deviation are passed directly to NSGA-II without these additional surrogates, which is more efficient (since only one surrogate has to be optimized) and more accurate (due to the lack of redundant surrogates). The key advantage of this process is the cheap computation of statistics through sampling the surrogate model. In step 3, either pseudorandom numbers, so Monte Carlo (MC) samples, or low-discrepancy sequences such

as Halton, Sobol or Faure sequences, referred to as quasi-Monte Carlo (QMC) methods, can be used. QMC samples are more evenly distributed and QMC can be superior to MC due to its faster rate of convergence close to $\mathcal{O}(1/N)$, versus $\mathcal{O}(1/\sqrt{N})$ for MC [21]. One of the advantages of the proposed workflow is that the surrogate is independent of the number of UQ samples. To ensure fast UQ, the surrogate, as opposed to the expensive black box function, can be sampled efficiently thousands of times to calculate the statistics.

A. Robust Optimization with an Analytical Function

As an analytical toy problem, a stochastic version of the modified Branin function [22] in Eq. (2),

$$\begin{aligned}
 f(\mathbf{x}) = & \overbrace{\left(x_2 - \frac{5.1}{4\pi^2} x_1^2 + \frac{5}{\pi} x_1 - 6 \right)^2 + 10 \left(1 - \frac{1}{8\pi} \right) \cos x_1 + 10 + \frac{1}{3} (x_1 + 5)}^{\text{Modified Branin function [22]}} \\
 & + 50000 \cdot \underbrace{\left(\frac{\frac{x_3}{8} (0.05 + \cos(\frac{x_1}{8} + 1))^4 + \cos(\frac{x_2}{8} + 7)^4 + \frac{x_4}{8} \cos(\frac{x_1}{8} + 1)^2 \cos(\frac{x_2}{8} + 7)^2}{(\frac{x_1}{8} + 1)^2 + 2(\frac{x_2}{8} + 7)^2} \right)}_{\text{Extra uncertain / noise term (cf. Ref. [5])}} \quad (2)
 \end{aligned}$$

is used, where in addition to the usual design parameters, $x_1 \in [-5, 10]$ and $x_2 \in [0, 15]$, two additional noise parameters, $x_3, x_4 \in [-8, 8]$, are part of the function. Without noise, the modified Branin function is a multimodal function and has a single global minimum in the upper left corner at $f(-\pi, 12.275) \approx 1.0174$.

A bound-constrained multi-objective optimization problem formulation is given by Eq. (3):

$$\begin{aligned}
 & \underset{\mathbf{x}}{\text{minimize}} && \text{E}[f(\mathbf{x})] \text{ and } \text{Var}[f(\mathbf{x})], \\
 & \text{subject to} && x_1 \in [-5, 10], \\
 & \text{and} && x_2 \in [0, 15], \\
 & \text{and} && x_3, x_4 \in [-8, 8].
 \end{aligned} \quad (3)$$

To solve this, a 20-point Latin Hypercube DOE for design and noise parameters is generated and evaluated first. The second step is to construct the Kriging surrogate given the parameters and their responses. Third, the statistics are computed using 100,000 QMC Sobol samples. Fourth,

a multi-objective search is run with NSGA-II. Both population and generation size are set to 100 and a space-filling DOE is used inside NSGA-II at the start. The result of this search is a set of non-dominated, Pareto-optimal update points, which can be added to the initial DOE. The update point selection method first selects points in the middle of the Pareto front, followed by endpoints and points in between. Every new point is compared to existing points to ensure that only truly new points are added. The optimization loop can be repeated until convergence is detected or until a fixed time limit has been reached. Convergence is monitored after every iteration. New Pareto-optimal infill points are compared to all previous points to make sure they are truly new and not too close to existing points. The search is stopped if not a single new Pareto-optimal point can be found for four consecutive iterations. Besides Pareto-optimal update points (exploitation), space-filling update points can also help improve surrogate model accuracy through better space exploration. They are computed with the Euclidean distance by selecting points furthest away from existing points (DOE and update points). In this example, five Pareto-optimal and five space-filling update points are requested in every iteration. Additionally, whenever update points are clustered too closely, again based on a minimum allowed Euclidean distance, a filtering method using a d -dimensional hypersphere replaces such points with space-filling update points.

The results from this methodology applied to Eq. (3) are depicted in Fig. 3. The resulting estimates of both mean and standard deviation in the design parameter space are shown in Figs. 3a-3b respectively. Figure 3c shows the final Pareto front in the objective space after 100 iterations through the loop. For validation, the metamodel-based search was compared to a direct NSGA-II search using a 20-point Sobol sequence DOE and 100,000 Halton sequence UQ samples per design point. The agreement of both Pareto fronts in Fig. 3c confirms the previous results.

To gauge surrogate model accuracy, metrics such as the coefficient of determination (R^2), the

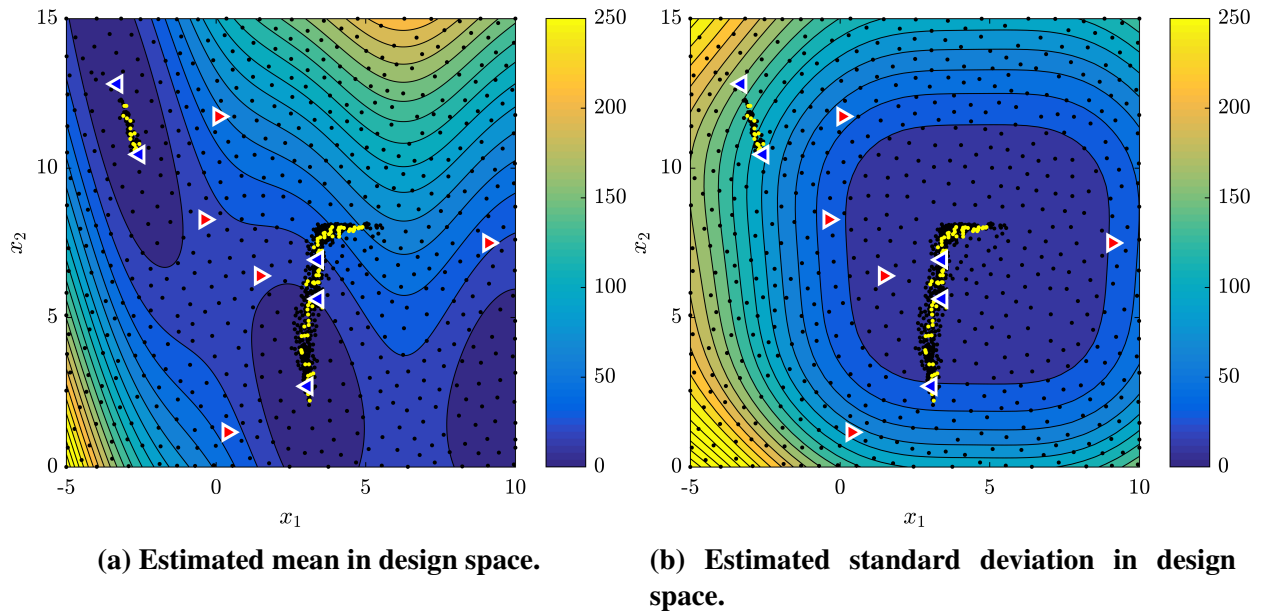


Figure 3 Robust optimization using Eq. (2) with 5 Pareto-optimal update points (\blacktriangleleft), 5 space-filling update points (\blacktriangleright), all Pareto-optimal points (\bullet) and all evaluated points (\bullet).

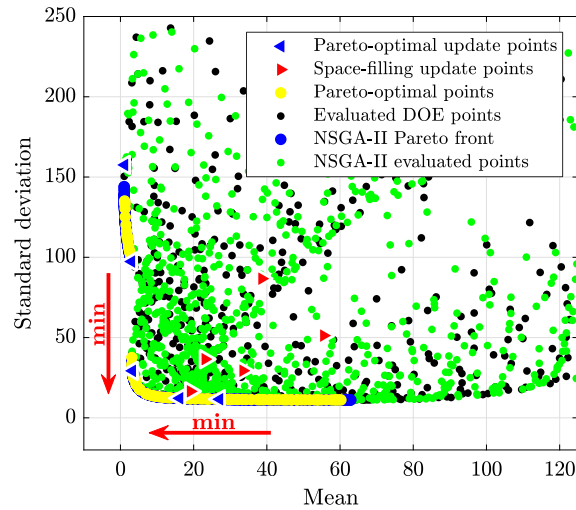


Figure 3 (cont.) In addition to markers from Figs. 3a-3b, direct NSGA-II Pareto front points (\bullet) and corresponding evaluated points (\bullet) are depicted for validation.

root-mean-square error (RMSE) and the maximum absolute error (MAE) as defined in Eq. (4),

$$R^2 = 1 - \frac{\sum_{i=1}^n (f_i - \hat{f}_i)^2}{\sum_{i=1}^n (f_i - \bar{f})^2}, \quad (4a)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (f_i - \hat{f}_i)^2}, \quad (4b)$$

$$\text{MAE} = \max(|f - \hat{f}|), \quad (4c)$$

can be computed using k -fold cross-validation, i.e., using k subsets of points. RMSE and MAE are reported in percent with regard to the full range of values, e.g., in Table 1. Expectedly, given the huge number of points in only four dimensions, the final surrogate model accuracy is excellent.

Table 1 Cross-validated accuracy estimates for surrogate models of $f(\mathbf{x})$.

Error metric	Beginning (20 points)	End (1020 points)
R^2 [%]	91.2214	99.9958
RMSE [%]	6.4896	0.0593
MAE [%]	15.2618	1.0988

Ordinary Kriging surrogate models approximate black box function responses $f(\mathbf{x})$ at any point \mathbf{x} with the stochastic process model $y = \mu + \epsilon(\mathbf{x})$, where μ is a constant mean and ϵ a Gaussian error with zero mean and variance σ^2 . ϵ is assumed to depend on the distance between observed data points. Based on the choice of basis function, the correlation matrix Ψ is then constructed. In engineering applications, this is typically the correlation function in Eq. (5) [22]:

$$\Psi(\mathbf{x}^{(i)}, \mathbf{x}^{(l)}) = \exp\left(-\sum_{j=1}^d \theta_j |x_j^{(i)} - x_j^{(l)}|^{q_j}\right) + \lambda \delta_{ij} \quad (5)$$

Surrogate model hyperparameters θ_j , q_j and λ are tuned using maximum likelihood estimation (MLE), where the concentrated log-likelihood function (CLLF) is optimized with a hybrid optimizer, viz., particle swarm optimization (PSO) followed by a greedy downhill search. λ is a regression constant added to the diagonal of Ψ to deal with potentially noisy CFD outputs by either interpolating ($\lambda = 0$) or regressing ($\lambda > 0$) observed data points. When $q_j = 2$, Eq. (5) is called the Gaussian correlation function. A new ordinary Kriging prediction \hat{f} at update a point \mathbf{x}^N is given by Eq. (6) with $\psi(\mathbf{x}^N) = \Psi(\mathbf{x}^N, \mathbf{x}^{(i)})$ as the correlation vector between the point and the DOE:

$$\hat{f}(\mathbf{x}^N) = \hat{\mu}_r + \psi^T (\Psi + \lambda \mathbf{I})^{-1} (\mathbf{f} - \mathbf{1} \hat{\mu}_r), \quad (6)$$

$$\text{where } \hat{\mu}_r = \frac{\mathbf{1}^T (\Psi + \lambda \mathbf{I})^{-1} \mathbf{f}}{\mathbf{1}^T (\Psi + \lambda \mathbf{I})^{-1} \mathbf{1}}.$$

B. Robust Optimization Heuristic for Turbine Blades

The robust turbine blade optimization workflow is an extension of the principles described in Sec. III. A for the analytical example. Figure 4 depicts the entire workflow, which can be divided into three main subgroups: In steps 1-3, the initial DOE is generated and evaluated. This involves the geometry construction process as well as the black box CFD objective and constraint functions described later in Sec. III. E. The unnumbered steps below steps 1-3 constitute the statistical basis. As described in Sec. III. 2, based on the 58 optical blade scans, real parameter samples are generated from fitted probability distributions. Also, a target correlation matrix is extracted to generate realistically correlated pseudorandom samples with Iman & Conover's method (see Sec. III. D. 2 and Sec. III. D. 3). Steps 4-9 constitute the main robust optimization loop. Steps 4-5 correspond roughly to what has been described in the analytical example in Sec. III. A. The real workflow also allows the definition of constraint functions for any quantity of interest. Surrogate models are then built for each objective and each constraint function. Step 5, NSGA-II, yields a set of Pareto-optimal points. Step 6 involves a simple twofold check: On one hand, whether or not to stop because the optimization budget is exhausted, and on the other to leave the loop in case the Pareto front has not changed meaningfully in between four successive iterations. In step 7 the workflow uses incremental increases of the number of UQ samples to speed up UQ search times. Steps 8 and 9 are detailed in subsequent sections, Sec. III. D and Sec. III. E respectively.

C. Geometry and Analysis of the Real Scans' Performance

First, to serve as a reference for validation of the parametrized blades for UQ in the next section, the original scans' performance is analyzed by accurately reverse engineering the scanned blade geometries. Having removed the film cooling holes, the blade geometry is extracted directly from STL files by slicing the triangle edges at n radial positions to extract intersection points. These slices are extracted between 10% and 90% span to ignore the fillet regions below and above. The only reconstruction errors are due to the partitioning of scans into discrete airfoil sections and due to the cooling hole removal process. Both are negligibly small. Results are presented in Sec. IV. A.

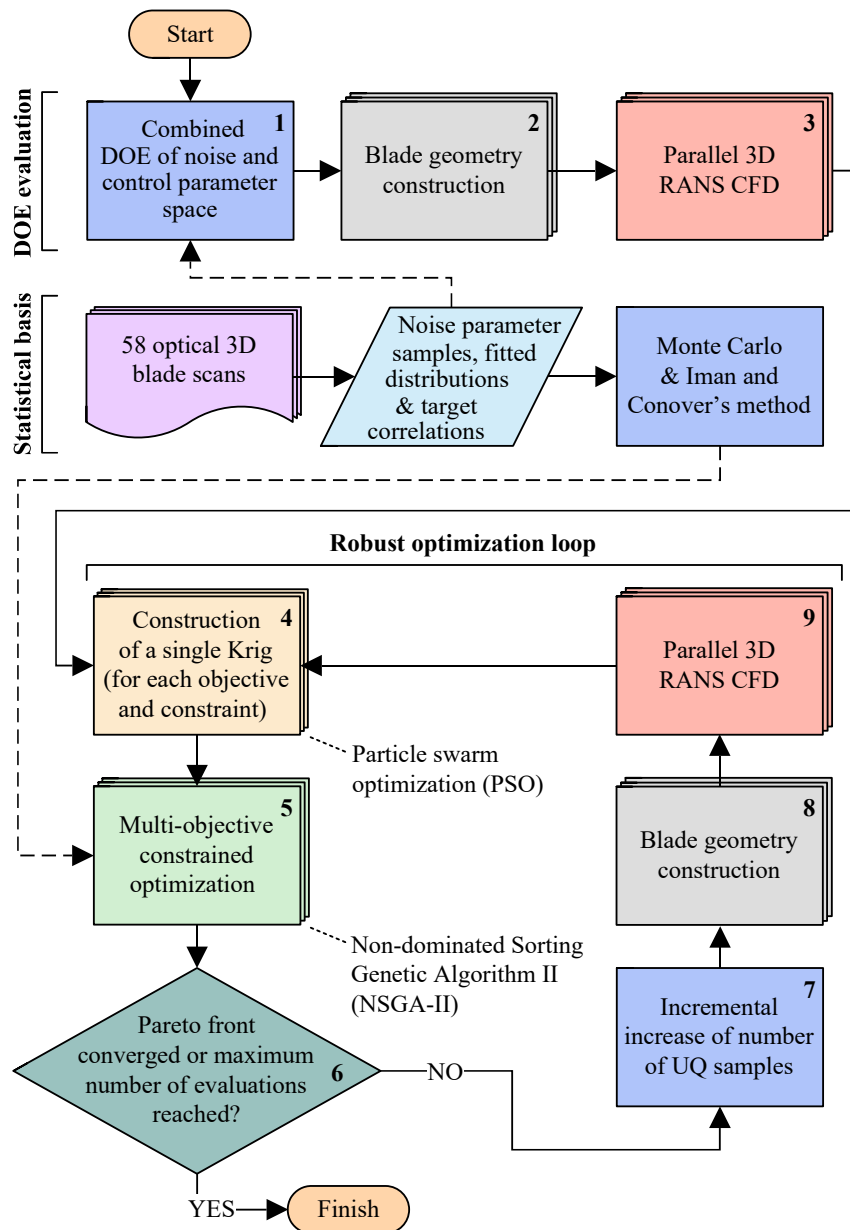


Figure 4 Measurement-based robust turbine blade optimization flowchart.

D. Parametric Construction of Realistic Blade Models

Parametrized blades, as opposed to the original scan files, are necessary to be able to construct as many artificial, but realistic blades as required for UQ. In principle, various methods such as free-form deformation or principal component analysis (PCA) can be used for this, e.g., see Ref. [23] for PCA. Here, a reverse engineering process based on a set of 14 classical airfoil parameters, implemented in a tool called Blade2Parameter (B2P) developed at TU Dresden, is used [6]. Figure 5 depicts all 14 parameters applied to a turbine blade section. As opposed to parameterizing the en-

tire airfoil, based on the assumption that differences are small, deviations are modeled in relation to a baseline blade shape such as the CAD model. This parameterization has two distinct advantages: First, unlike modes from PCA, variations of standard airfoil parameters are straightforward to interpret, whereas PCA modes are more abstract than familiar airfoil parameters. Second, only 14 noise parameters (or even fewer, as shown in the sensitivity analysis in Sec. IV. B) are required to accurately model geometric deviations. In contrast, the number of PCA modes can be far greater than 14 and also vary significantly depending on the problem, as has been shown in previous UQ studies, e.g., by Refs. [23, 24]. The number of parameters should be as low as possible in any optimization due to the curse of dimensionality and the associated increased complexity.

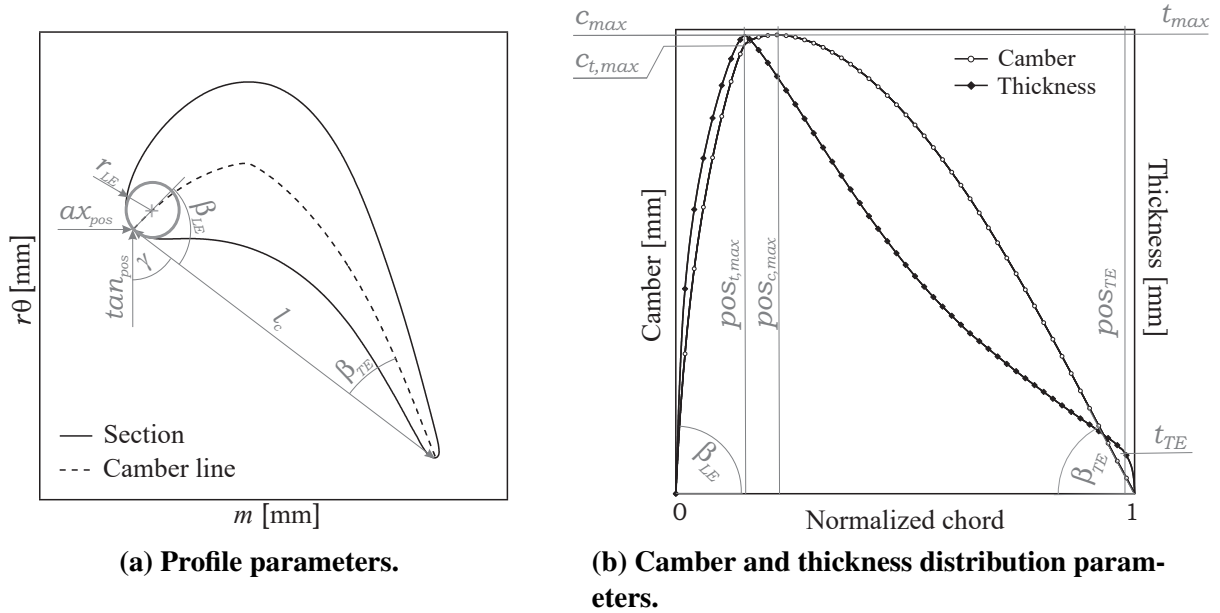


Figure 5 Illustration of Blade2Parameter's noise parameters.

Through Eq. (7),

$$\mathbf{p}^i = \mathbf{p}^{\text{baseline}} + \Delta\mathbf{p} \quad (7)$$

the parameterization creates new blade shapes \mathbf{p}^i by adding delta parameters $\Delta\mathbf{p}$ to a baseline blade shape $\mathbf{p}^{\text{baseline}}$ with camber, chord and thickness distributions. Samples for $\Delta\mathbf{p}$ are obtained by subtracting the baseline model's parameters from those of the scans. The 14 delta parameters

for a single section are listed in Eq. (8),

$$\Delta \mathbf{p} = [ax_{pos} \ tan_{pos} \ l_c \ \gamma \ \beta_{LE} \ \beta_{TE} \ t_{max} \ pos_{t,max} \ c_{max} \ pos_{c,max} \ c_{t,max} \ pos_{TE} \ t_{TE} \ r_{LE}]^T \in \mathbb{R}^{14}. \quad (8)$$

$\Delta \mathbf{p}$ consists of axial and tangential distances ax_{pos} and tan_{pos} , chord length l_c , stagger angle γ , LE and TE camber angles β_{LE} and β_{TE} , maximum thickness t_{max} and its position $pos_{t,max}$, maximum camber c_{max} and its position $pos_{c,max}$, camber at the position of maximum thickness $c_{t,max}$, TE position pos_{TE} , TE thickness t_{TE} and LE radius r_{LE} .

First, B2P is used to extract parameter samples from all 58 blade scans to create the statistical basis for UQ and optimization. Delta parameter samples can then be created by subtracting parameters of a real scan from those of the nominal CAD model. Second, B2P is used again to construct artificial turbine blades given a pseudorandom sample of $\Delta \mathbf{p}$. Given that the parameterization is for 2D airfoil profile sections, n sections are stacked spanwise to create the final 3D blade models. For each 2D section in length-preserving $r\theta$ - m space, ax_{pos} and tan_{pos} are used to position the blade. Given this point, stagger angle γ and chord length l_c are added. Next, camber and thickness distributions are both morphed towards their target distribution using several morphing vectors. Then, the morphed camber line is added to the chord line and the morphed thickness distribution is applied to the camber line. This process yields a complete airfoil profile and is repeated for all specified span positions.

Applying Eq. (8) to a single blade section technically results in 14 parameters *per section*. For example, when a blade is modeled with 21 2D airfoil sections, there would be $21 \cdot 14 = 294$ delta parameters in total. As this would be far too many parameters for many purposes, all delta parameters can be radially averaged to yield $\Delta \bar{\mathbf{p}}$, which reduces the number of parameters to just 14. This radial averaging, termed the one mean area approach, can be justified due to the high radial correlations between sections. The reason this works is that the delta parameters $\Delta \bar{\mathbf{p}}$ are added to the baseline blade to add variability. As an example, the high correlations for Δt_{max} are shown in Fig. 6, but other delta parameters exhibit similar correlations.

Also, the regions below 10% span and above 90% span require special attention. This is




Figure 6 Spanwise correlations for Δt_{max} noise parameter.

because the parameterization does not work for fillets, i.e., sections there are not airfoil profiles. To bridge the gap, models are extended through extrapolation using a special in-house code `parapy`. Local displacement vectors are computed from two corresponding points on the outermost profiles and smoothed using a Savitzky-Golay filter before new profiles are created. All required sections from 0 %-100 % span are then available to construct the final 3D blade geometry.

Another in-house tool, which fits all airfoil section points with four cubic B-spline curves, is used to construct the blades from sections. Each surface segment is C^2 continuous to ensure smooth transitions. The fitting essentially acts as a data point smoothing. From all radially stacked splines, the final 3D blade model is then created, as shown in Fig. 1c. As a final step, each Siemens NX blade model is sectioned again to export an arbitrary number of sections into text input files for the meshing tool.

It is also worth noting that the reverse engineering workflow also supports constructing and varying other geometric features such as the shroud geometry or film cooling hole positions and dimensions. Hub and shroud fillets could be added too. Here, every blade's shroud is shifted axially. This shift is computed by comparing the axial chord midpoints of the topmost aerofoil sections of the artificial blade to that of the nominal CAD geometry and the nominal shroud geometry is moved accordingly.

During the optimization, PADRAM, one of Rolls-Royce's in-house codes [25], is used to generate new turbine blade shapes given various engineering *design* parameters (cf. Sec. III. E. 4).


Fine, unstructured, triangular surface meshes of each blade are then constructed with ANSYS ICEM CFD and exported as STL files. Then, baseline parameters and distributions are reverse engineered and the *noise* parameters are applied.

1. Comparison of Geometric Reconstruction Accuracy

The excellent reconstruction accuracy of the described approach is depicted in Fig. 7. In Fig. 7a on the left, the magnitude of one of the scanned blade's deviations are shown. Figure. 7b in the middle depicts the deviations of a one-to-one, nonparametric reconstruction as described in Sec. III. C with 19 sections extracted directly from the scan's triangles. On the right in Fig. 7c is the fully parameterized blade. The colors and histograms demonstrate that reconstruction errors from the parameterized blade are very small compared to the actually measured deviations. Note that the baseline model used is a so-called mean value model (MVM), which is essentially an average of all blade scans' parameters, thickness, camber and chord distributions. Using the MVM in lieu of the nominal CAD model can lead to better reconstruction accuracy, but differences are very small, i.e., on the order of a few tens of micrometers, and hence negligible.

2. Sample Generation

Pseudorandom samples have to be generated to construct new artificial blade geometries with realistic geometric variations. Equation (8) is key to doing this: delta parameter samples are generated by first applying the blade parameterization to each blade scan. Subtracting each scan's parameters from the baseline model's parameters yields a 58×14 delta parameter sample matrix \mathbf{P} , which constitutes the statistical basis for the entire workflow. For example, the samples in the first row of \mathbf{P} create the parametrized reconstruction of scan #1 as shown in Fig. 7. New delta parameter samples $\Delta\bar{\mathbf{p}}$ have to be generated based on the samples in \mathbf{P} . Since visual inspection and statistical tests showed that parametric PDFs do not fit the data well, each PDF was estimated using KDE. KDE is nonparametric and more flexible, since random variables with any distribution can be handled. Here, a kde function using Gaussian kernels and an automatic optimal bandwidth selection process is used, see Ref. [26] for details. Any KDE estimate is a sum of many "bumps", i.e., kernels,



(a) Nominal CAD model vs. scan #1. (b) Nonparametric reconstruction vs. scan #1. (c) Parametric reconstruction vs. scan #1.

Figure 7 Surface comparisons to illustrate measured blade deviations and reconstruction accuracy (units are normalized).

given the observed samples. Here, $n = 58$ realizations for each parameter are extracted from the scans. Each PDF is fitted separately. The kernel density estimator is defined in Eq. (9) as

$$\hat{f}_h(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right), \quad \text{where } K(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} \quad (9)$$

is the Gaussian kernel [26]. h is the bandwidth, which specifies the kernel width. As an example for the maximum thickness Δt_{max} delta parameter, the empirical cumulative distribution function (ECDF) in Fig. 8 illustrates that KDE is able to approximate the given sample data adequately. Similar fits were generated for all other delta parameters. For each delta parameter, minimum and maximum parameter bounds are extracted from the original data in \mathbf{P} to effectively truncate the PDFs below and above these values. Random samples with values larger or smaller than these bounds are discarded as in reality blades with extreme variations in violation of the allowed tolerance limits are sorted out. New samples are generated by adding values drawn uniformly at random from each sample X_i with replacement to values drawn from the Gaussian $\mathcal{N}(X_i, h)$, i.e.,

centered at X_i with variance h .

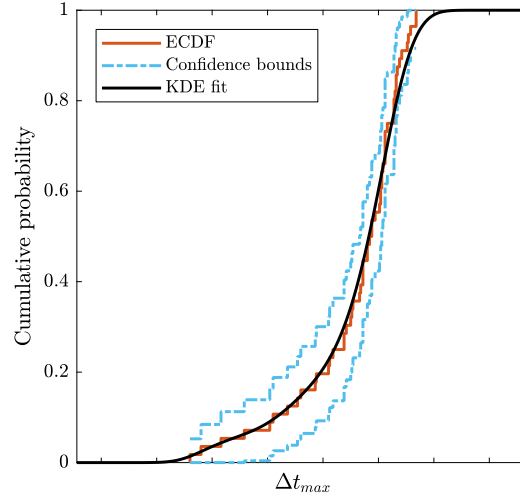


Figure 8 Kernel density-estimated cumulative distribution function fit of Δt_{max} parameter.

3. Sample Correlation Control

To create an artificial turbine blade geometry, the delta parameters are sampled independently from each KDE-fitted PDF. A matrix with N rows of samples for $\Delta \bar{\mathbf{p}}$ could be created and used to construct blades. However, these new blades would be unrealistic without also considering the strong correlations of the original scans' parameters. This is illustrated in Fig. 9 using Spearman's rank correlation coefficient r_s to reveal monotonically dependent delta parameters. r_s is defined in Eq. 10 with the operator $\text{rg}(X_i)$ for the ranked scores and $\bar{\text{rg}}_X$ for their average [27]:

$$r_s = \frac{\sum_i (\text{rg}(X_i) - \bar{\text{rg}}_X)(\text{rg}(Y_i) - \bar{\text{rg}}_Y)}{\sqrt{\sum_i (\text{rg}(X_i) - \bar{\text{rg}}_X)^2} \sqrt{\sum_i (\text{rg}(Y_i) - \bar{\text{rg}}_Y)^2}} \quad (10)$$

It is crucial to ensure the original correlation structure is preserved for newly generated samples and Iman and Conover's method is used to do this here. Iman and Conover's method is distribution-free and allows the reintroduction of rank correlations *after* independently generating N rows of Monte Carlo samples. The target correlation matrix stems from the original blade sample's parameters. For details of the method, see Ref. [27].




Figure 9 Target delta parameter correlation matrix.

E. Computational Fluid Dynamics Model & Setup

The CFD model consists of the turbine rotor blade only to keep the computational cost as low as possible. Rotationally periodic boundary conditions are used, so that only a single blade passage needs to be modeled. Although this is a frequently used simplification, one has to keep in mind that this introduces a perfect correlation between all blades around the annulus and entails overestimating the variability, e.g., see Ref. [4]. However, this is of no concern for the actual trade-off between performance mean and variability, because only their absolute magnitude is affected. Boundary conditions (BCs) are specified for all inlets and outlets and all rotating and stationary surfaces. All walls are adiabatic. A meridional view of the complete turbine stage with all BCs can be found in Ref. [2]. The main inlet is a normal subsonic inflow and the outlet a non-reflecting subsonic outflow. Mass flow inlets are used for both hub cavity and rim seal. Rolls-Royce's standard in-house solver HYDRA is used. A multigrid method with a 5-grid-level W-cycle is used to accelerate convergence. The general CFD approach is to solve the steady-state 3D compressible RANS equations defined in Eqs. (11a)-(11b) using Reynolds and Favre averaging [28]:

$$\frac{\partial \bar{\rho}}{\partial t} + \frac{\partial}{\partial x_j} (\bar{\rho} \tilde{u}_j) = 0 \quad (11a)$$

$$\frac{\partial}{\partial t} (\bar{\rho} \tilde{u}_i) + \frac{\partial}{\partial x_j} (\tilde{u}_i \bar{\rho} \tilde{u}_j) = -\frac{\partial \bar{p}}{\partial x_i} + \frac{\partial \bar{\sigma}_{ij}}{\partial x_j} - \frac{\partial R_{ij}}{\partial x_j} \quad (11b)$$

$$\frac{\partial}{\partial t}(\bar{\rho}\tilde{E}) + \frac{\partial}{\partial x_j}(\tilde{u}_j\bar{\rho}\tilde{H}) = \frac{\partial}{\partial x_j}(\overline{u_i\sigma_{ij}}) - \frac{\partial\bar{q}_j}{\partial x_j} - \frac{\partial\bar{Q}_j}{\partial x_j} \quad (11c)$$

In Eqs. (11a)-(11c), $\bar{\sigma}_{ij} \approx 2\tilde{\mu}_t(\tilde{S}_{ij} - 1/3\tilde{S}_{kk}\delta_{ij})$ is the viscous stress tensor, $R_{ij} = \bar{\rho}\tau_{ij} = \overline{\rho u_i'' u_j''} = \bar{\rho} \overline{u_i'' u_j''}$ is the turbulent stress tensor (Reynolds stress), $S_{ij} = 1/2(\partial u_i/\partial x_j + \partial u_j/\partial x_i)$ is the strain rate tensor, $\bar{\rho}\tilde{E} = \bar{\rho}c_v\tilde{T} + 1/2\bar{\rho}\tilde{u}_i\tilde{u}_i + 1/2\bar{R}_{ii}$ is the averaged total energy field, $\bar{\rho}\tilde{H} = \bar{\rho}c_p\tilde{T} + 1/2\bar{\rho}\tilde{u}_i\tilde{u}_i + 1/2\bar{R}_{ii} = \bar{\rho}c_p\tilde{T}_0$ is the averaged total enthalpy, $\bar{q}_j = -\overline{k_T \partial T/\partial x_j} \approx -c_p\tilde{\mu}/Pr \partial\tilde{T}/\partial x_j$ is the mean heat flux vector (Fourier's law) and $\bar{Q}_j = \bar{\rho}c_p\overline{u_j'' T''} + 1/2\bar{\rho}\overline{u_i'' u_i'' u_j''} + \bar{\rho}\tilde{u}_i\tau_{ij}$ is the turbulent scalar flux. The turbulence closure problem stems from the turbulent stress $\bar{\rho}\tau_{ij}$ and all terms in $\bar{Q}_j = \bar{\rho}c_p\overline{u_j'' T''} + \bar{\rho}\overline{u_i'' u_i'' u_j''}/2 + \bar{\rho}\tilde{u}_i\tau_{ij}$. Here, the one-equation Spalart-Allmaras turbulence model defined in Eq. (12) is solved for the Spalart variable $\tilde{\nu}$:

$$\begin{aligned} \frac{\partial\tilde{\nu}}{\partial t} + u_j \frac{\partial\tilde{\nu}}{\partial x_j} = c_{b1}(1 - f_{t2})\tilde{S}\tilde{\nu} - \left[c_{w1}f_w - \frac{c_{b1}}{\kappa^2}f_{t2} \right] \left(\frac{\tilde{\nu}}{d} \right)^2 \\ + \frac{1}{\sigma} \left[\frac{\partial}{\partial x_j} \left((\nu + \tilde{\nu}) \frac{\partial\tilde{\nu}}{\partial x_j} \right) + c_{b2} \frac{\partial\tilde{\nu}}{\partial x_i} \frac{\partial\tilde{\nu}}{\partial x_i} \right] \end{aligned} \quad (12)$$

In compressible flows, the turbulent dynamic eddy viscosity μ_t is computed with the relation $\mu_t = \rho\tilde{\nu}f_{v1}$. For the sake of brevity, exact model constants used in HYDRA are not provided here and all additional model equations, e.g., $f_{v1} = \dots$, can be found in Ref. [29]. The working fluid air is modeled as a calorically perfect gas with heat capacity ratio $\kappa = 1.3$. Thus, it obeys the ideal gas law equations defined in Eq. (13),

$$p = \bar{\rho}R_s\tilde{T}, \quad \kappa = \frac{c_p}{c_v}, \quad c_p - c_v = R \quad \text{and} \quad e = c_v T, \quad (13)$$

where κ , R_s , c_p and c_v are all constant. In reality, the fluid is a two-phase air-fuel mixture, but the presence of fuel is neglected here.

Also note that for each blade, the CFD control volume position is fixed in space, so that each blade is installed with whatever deviations it was made with, i.e., as if its firtree root was attached to the rotating turbine disk. For the CFD model, there are two geometrical simplifications vis-à-vis reality, i.e., the very small shroud and hub gaps are not modeled and instead are continuous in the annulus, as is commonly done. Also, the case here is run ‘‘cold’’ and a cold-to-hot transformation

would be necessary for real-engine results. Moreover, since film cooling is not part of the scope of this study, the inlet total pressure was increased to offset the missing coolant mass flow rate from all film cooling holes, i.e., to ensure the same operating point as if cooling was enabled. This is necessary since for this blade film cooling accounts for roughly 5% of the total mass flow rate through the annulus, which is significant. Besides the fully-featured case, a defeatured model only consisting of the rotor blade from 10% to 90% span is used briefly in Sec. IV. A, because it does not involve extrapolated sections.

1. Mesh & Quantities of Interest

A structured $y^+ \approx 1$ mesh was created with PADRAM, one of Rolls-Royce's in-house meshing tools [25]. A mesh independence study was carried out to determine the final mesh size. Variations from the mesh are quantified in detail in Sec. III. 3. The final mesh with a size of roughly 5.5 million cells is depicted in Fig. 10.




Figure 10 Fully-featured CFD domain and mesh.

Quantities of interest monitored and extracted from each CFD solution include the isentropic efficiency defined in Eq. (14),

$$\eta = \frac{\sum_{i=inlets} (\dot{m}h_0)_i - \sum_{i=exits} \dot{m}h_0}{\sum_{i=inlets} (\dot{m}h_0)_i - \sum_{i=sec_exits} (\dot{m}h_0)_i - (\dot{m}h_{0,ideal})_{main_exit}}, \quad (14)$$

as the ratio of actual to ideal, isentropic work for turbines, the inlet capacity in Eq. (15),

$$\phi = \frac{\dot{m}_4 \sqrt{T_{04}}}{p_{04}}, \quad (15)$$

and the specific work produced.

$$w = \frac{\sum_{i=inlets} (\dot{m}H_0)_i - \sum_{i=exits} (\dot{m}H_0)_i}{(\dot{m}T_0)_{main,inlet}}. \quad (16)$$

The work is computed from the rotor torque $\tau = \int \mathbf{r} \times \mathbf{F} dA$ with all forces from pressure and viscous shear stress terms acting on all rotating walls. Torque and specific work are both directly related to the generated power: $P = \tau\Omega = \dot{m}w$. During post-processing, values from the last 100 CFD iterations are averaged for each quantity of interest due to small oscillations from the steady-state CFD approach. All flow quantities are mass-weighted.

2. CFD Validation

For the scanned turbine rotor no experimental data has been recorded, so CFD validation had to be done by inference and a two-step approach was pursued: first, it was ensured that the present case matches Rolls-Royce's internal CFD solution. Second, the CFD can be validated by using the same settings as in Sec. III. E in a different but similar case, for which experimental data is available. One such turbine with time-averaged measurements is the MT1 research rig. The experimental data can be found in Ref. [30] and geometry and boundary conditions in Ref. [31]. The CFD was set up with a $y^+ \approx 1$ grid, similar to the HP rotor. Figure 11 shows that there is very good agreement with both experimental and numerical results, all using the Spalart-Allmaras turbulence model. Due to the experiments' time-averaging, it comes as no surprise that unsteady phase-lagged BCs perform best. For both steady-state results from HYDRA and Fluent, reproduced from Ref. [31], there is some disagreement near the SS shock, which is attributable to unsteady flow. For the rest, the agreement is very good.

3. Epistemic vs. Aleatoric Uncertainties

In UQ studies, it is important to ensure that epistemic (modeling) uncertainties, i.e., volatility stemming from the computational model itself, cause smaller variations than the effects of real aleatoric (inherent) aberrations, i.e., the variations to be quantified. Here, epistemic uncertainties stem from sectioning the

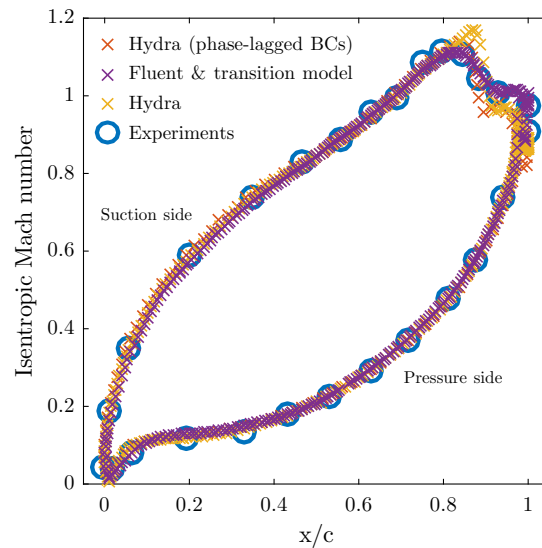


Figure 11 Isentropic Mach number distribution of MT1 turbine stator at half-span (○ from Ref. [30] and × from Ref. [31]).

blade, meshing, mass averaging over postprocessing planes and averaging of 100 samples for each quantity of interest. Geometric differences as a result of sectioning the blade are negligibly small ($\ll 10 \mu\text{m}$) and hence are of no concern.

As will be shown in Sec. IV. A, aleatoric geometric blade variations lead to output variations of roughly 0.2%, 0.5% and 0.7% for efficiency, capacity and specific work respectively. This stands in contrast to maximal meshing variations of 0.02% for efficiency, -0.003% for capacity and -0.015% for specific work, which were obtained by comparing the outputs from the selected mesh-independent 5.5 million cell mesh with those from finer 8.2 and 15.3 million cell meshes. Hence, meshing uncertainties are at least an order of magnitude smaller than aleatory variations. Also, an iteration reduction study was carried out because it helps to significantly speed up each black box function's evaluation time in HYDRA. With 40 preceding W-cycle multigrid iterations, errors from fewer iterations are very small, viz., around hundredths of percent. Thus, the number of iterations could be reduced from 1000 to 300.

The convergence of the sample mean and the sample standard deviation is shown in Fig. 12. It is evident that fluctuations of the mean are small in comparison to the much slower convergence of the standard deviation.

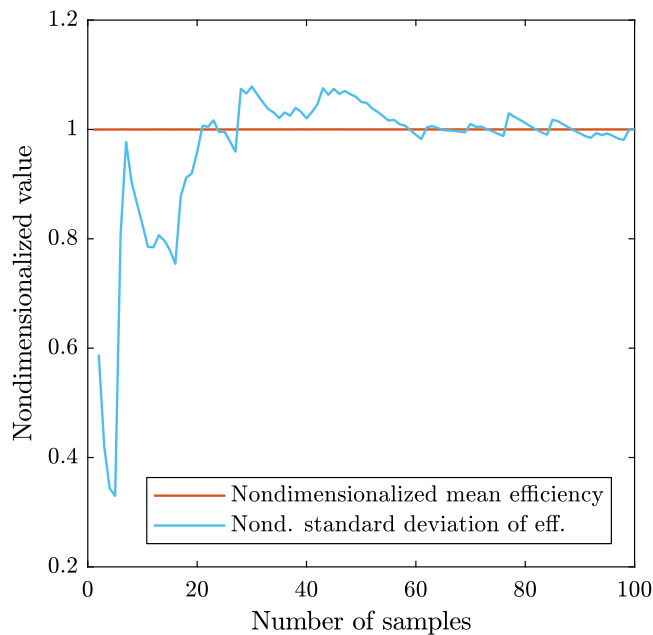



Figure 12 Convergence of objectives for 100 artificial blades.

4. Design Parameters

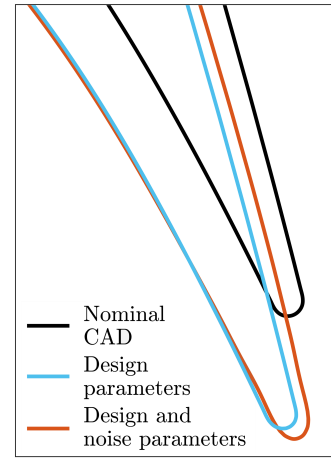
A set of five 3D engineering design parameters are used as controllable design parameters: skew, sweep, lean and LE & TE recambering. As an example, Fig. 13a depicts the effects of all five parameters applied at three different radial positions, 0 %, 50 % and 100 % span, so 15 parameters in total. This is a solid choice as, for example, this parameterization has been used previously in a deterministic aerodynamic turbine optimization to alleviate negative effects from high swirling flows in modern lean-burn combustors [32]. The same reference also details the geometrical changes each parameter brings about. Essentially though, each parameter in some way modifies the 2D airfoil sections: e.g., skew rotates them, lean shifts them circumferentially, sweep shifts them axially and both recambering parameters shears them either at the LE or TE. Radially, a B-spline ensures a smooth transition across all airfoil sections. Figure 13b illustrates the effects of both the design and the noise parameters.

F. Optimization Problem Definition

For a basic aerodynamic optimization of a turbine rotor, the end result ought to be a blade that at some capacity delivers a certain amount of power with a smaller total pressure drop, i.e., higher efficiency. Hence, isentropic efficiency defined in Eq. (14) is selected as the objective function. For the purposes of optimization, efficiency is multiplied by negative one to minimize it. Equality constraints for capacity defined in



(a) Design parameters at three span positions.



(b) Design *and* noise parameters.

Figure 13 Illustration of design and noise parameter effects.

Eq. (15) and extracted work defined in Eq. (16) are implemented as two inequality constraints each. The capacity constraint for a single HP turbine rotor is chosen to be $\pm 0.1\%$ because often the stator chokes and the rotor is somewhat less critical. The same constraint limits are chosen for specific work. As will be shown in Fig. 14b, the nominal CAD geometry does not satisfy these requirements and hence is no valid solution. The complete problem formulation is defined mathematically in Eq. (17), where the *controllable* PADRAM design parameters are contained in vector \mathbf{d} . Since the main intention in this study is to showcase the implementation and design improvements, a 15 design parameter limit is chosen for \mathbf{d} . In addition, the vector $\Delta\bar{\mathbf{p}}$ contains the *uncontrollable*, pseudorandom delta noise parameters from Eq. (8). The optimization hence involves 29 parameters in total. Upper (\mathbf{d}_U) and lower (\mathbf{d}_L) control parameter bounds in Eq. (17) are ± 0.2 mm for sweep, $\pm 0.1^\circ$ for skew, $\pm 0.1^\circ$ for lean and $\pm 0.05^\circ$ for LE recambering and $\pm 0.25^\circ$ for TE recambering. Although these bounds appear to be relatively small in magnitude, the parameters change the blade shape much more than the manufacturing variations. The lower ($\Delta\bar{\mathbf{p}}_L$) and upper ($\Delta\bar{\mathbf{p}}_U$) bounds for the delta noise parameters ensure that no pseudorandom samples with extreme variations are generated, as

described in Sec. III. D. 2. The problem tackled is this:

$$\begin{aligned}
& \underset{x}{\text{minimize}} && \text{E}[-\eta_{is}(\mathbf{d}, \Delta\bar{\mathbf{p}})] \text{ and } \text{Var}[\eta_{is}(\mathbf{d}, \Delta\bar{\mathbf{p}})], \\
& \text{subject to} && 0.999 \cdot w_{\text{CAD}} \leq \text{E}[w(\mathbf{d}, \Delta\bar{\mathbf{p}})] \leq 1.001 \cdot w_{\text{CAD}}, \\
& \text{and} && 0.999 \cdot \phi_{\text{CAD}} \leq \text{E}[\phi(\mathbf{d}, \Delta\bar{\mathbf{p}})] \leq 1.001 \cdot \phi_{\text{CAD}}, \\
& \text{and} && \mathbf{d}_L \leq \mathbf{d} \leq \mathbf{d}_U, \\
& \text{and} && \Delta\bar{\mathbf{p}}_L \leq \Delta\bar{\mathbf{p}} \leq \Delta\bar{\mathbf{p}}_U,
\end{aligned} \tag{17}$$

where subscript CAD refers to the nominal CAD model’s reference performance. There is a key point to be made regarding the definition of constraints in Eq. (17), namely that only constraints for *mean values* are enforced here. Given that all stochastic outputs entail some inherent variation, this means certain realizations can violate these mean constraints as optimizers often tend to drive designs close the specified constraint limits.

G. Computational Implementation & Resources

Due to the large computational demands, the implementation allows for parallel computations for surrogate model training, UQ and searches through MATLAB’s parallel pool and CUDA GPU support. The computational setup here consists of an Intel Core i7-4790 and an Nvidia GeForce GTX 750 Ti. CFD simulations run remotely on 14 nodes of a Linux high-performance computing (HPC) cluster with eight-core Intel Xeon E5-2670 processors. On each cluster node, one instance of the CFD solver can use 16 message passing interface (MPI) processes. The multigrid method and the reduced number of solver iterations also entail a considerable speed-up. Despite all that, the optimization presented in the next section took about two months in total (DOE \approx 20 days; optimization \approx 30 days), largely due to the long evaluation time of about eight hours for the CFD. However, using more nodes on a larger cluster and parallelizing more tasks would bring this time down further.

IV. Results

A. Validation of Artificial Blade Performance Statistics

The first step is to ensure that randomly built blades exhibit the same statistical performance outputs as the real, original sample of blades. All 58 blade scan models are sliced and built straight from the original

STL files, as described in Sec. III. C. 100 artificial blades are constructed for comparison. Due to the good reconstruction accuracy, the expectation is to obtain output performance statistics similar to the real blades. Besides the mean, the coefficient of variation (CV) as defined in Eq. (18)

$$C_v = \frac{\hat{s}}{\bar{x}} \quad (18)$$

is used to quantify the amount of variability relative to the mean. CFD results of the defeatured model ($0.1 \leq s \leq 0.9$) are illustrated in Fig. 14a. In the figure, the overlapping histograms with outputs from both cases (blue and yellow) show that artificial blades can imitate actual blade variations very well. Indeed, the results show that the artificial blades' mean efficiency, capacity and specific work is underpredicted by only 0.01 %, 0.02 % and 0.03 % respectively. Similarly, CVs were found to be overpredicted by 0.40 %, 1.77 % and 1.44 % only. Good agreement is also observed in Fig. 14b for the fully-featured CFD model. The effects the variations have on the fully-featured model are similar to those in Fig. 14a for the mean capacity (only -0.02 % difference), but somewhat worse in terms of mean efficiency (-0.07 % difference) and mean specific work (-0.28 % difference). The mean output quantities are all underpredicted by only 0.01 %. Variation is again overpredicted, likely due to the extrapolated sections below 0.1 and above 0.9 span: CVs are 16.19 %, 20.70 % and 2.80 % larger than those of the original scans. In short, results in Fig. 14 show that artificially constructed blades are able to faithfully model the performance of real blade scans. The defeatured CFD model can do so a little bit better, probably because it lacks the extrapolated sections used in the fully-featured case.

Figure 14 also reveals the impact of the manufacturing variations on blade performance. For the fully-featured blade, mean efficiency drops by about 0.06 %, mean capacity is 0.21 % higher and work produced is 0.08 % lower. Hence, mean capacity is very similar in both models, whereas there are larger differences for efficiency and work. The shift in capacity is attributable to the relatively large deviations from the nominal CAD geometry near the TE, as previously depicted in Fig. 2. Due to the manufacturing variations, most blades are skewed with widened throat areas and this has led to the increase in mean capacity.

B. Sensitivity Analysis Results

Before optimizing, it is common practice to conduct a sensitivity analysis to reduce the required number of parameters in the model. This can bring about significant computational savings as each omitted variable reduces the size of the sampling space exponentially. Important and potentially irrelevant parameters

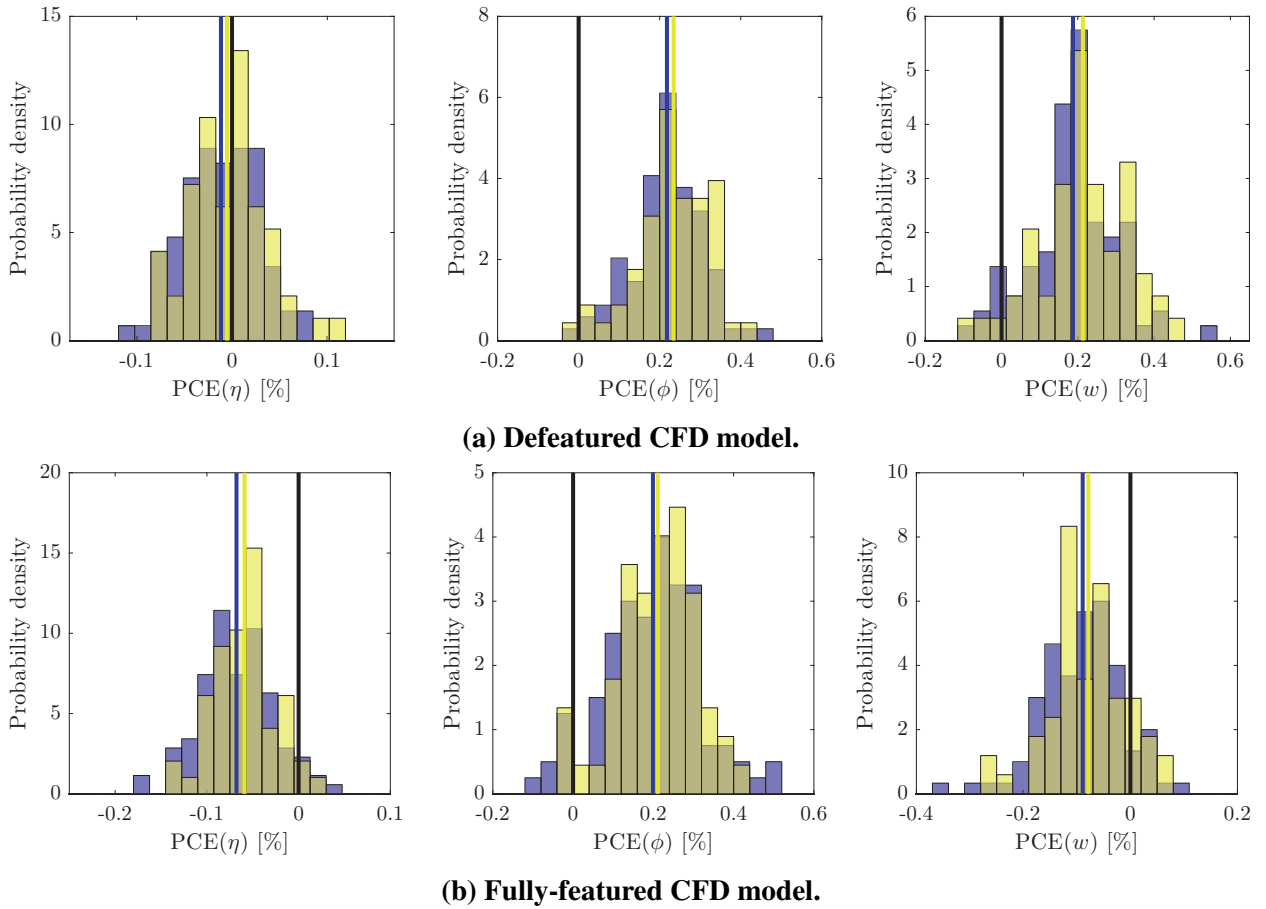



Figure 14 Comparison of CFD results of 100 artificial blades (samples ■ and mean —) and all 58 original blade scans (samples ■ and mean —) relative to the nominal CAD design intent (—).


can be revealed, but the analysis also helps to verify and to understand the model better. Here, for brevity, this is only presented for the noise parameters. Both correlation coefficient-based and surrogate model-based approaches are used, since already evaluated points can be reused and no additional effort is required. Although usage of Spearman's rank correlation coefficient is ill-advised due to the inherent parameter correlations, it is still often used. Other coefficients such as partial correlation coefficients (PCC) or coefficients of importance (COI) overcome this problem somewhat by canceling out effects of other parameters, e.g., see Ref. [33]. Here, the COI is used since it is far less sensitive towards correlated input parameters than the rank correlation coefficient. The COI is defined in Eq. (19) as

$$\text{COI} = R^2 - R_o^2, \quad (19)$$

where subscript o indicates the model built with one omitted parameter [2]. Figure 15 presents both the rank correlation coefficient and COI computed with 100 samples and the defeatured CFD model. Many rank correlation coefficients in Fig. 15a are not statistically significant, so no conclusions can be drawn from this data. However, TE position and maximum thickness parameters seem to be important. COI-based sensitivities in Fig. 15b place much more weight on few individual parameters and many relations found through rank correlations are not corroborated by the COI. Stagger angle and TE position seem to have the largest effect on efficiency. Similarly, the COI singles out the stagger angle as the most important parameter for capacity and work, but $R_o^2 < 0.8$ is not a reliable result. This corroborates previous results based on a smaller 50-point sample from an earlier UQ-only study based on an older parameterization and a different blade sample [2].



(a) Rank correlation coefficient (SS if $p < 0.05$ black, else white).



(b) Normalized COI (reliable if $R_o^2 > 0.8$ black, else white).

Figure 15 Sampling-based sensitivity analysis results.

Surrogate model-based sensitivities were computed to verify the coefficient-based results. Kriging's θ hyperparameters when suitably non-dimensionalized can serve as a statistical measure of parameter activity and can be used to produce parameter rankings [22]. Here, maximum likelihood estimation as described in Sec. III. A is used to tune Kriging θ hyperparameters within the log-search bounds $\log_{10}(\theta) \in [-10, 3]$. Logarithmic results are presented in Table 2 to help identify irrelevant parameters. The previous coefficient-based results are corroborated and the stagger angle is singled out again as the most important parameter for capacity and work, whereas the TE position again is most important for efficiency. Many other parameters

are also found to have a meaningful impact. Most importantly though, Kriging identifies the influence of Δtan_{pos} , the tangential position of the blade, as insignificant for all three outputs. Hence, this is the prime candidate for removal from the model. To test whether or not Δtan_{pos} indeed has no significant influence, 100 simulations were repeated without this parameter. Since it still required to build turbine blades though, the value of Δtan_{pos} was fixed to an average taken from the original scan sample matrix \mathbf{P} . The results, omitted here for brevity, show that Δtan_{pos} can indeed be held fixed.

Table 2 Kriging surrogate-based sensitivity analysis ($\log_{10}(\theta) \in [-10, 3]$).

Hyperparameters	Δax_{pos}	Δtan_{pos}	Δl_c	$\Delta \gamma$	$\Delta \beta_{LE}$	$\Delta \beta_{TE}$	Δt_{max}	Δpos_{tmax}	Δc_{tmax}	Δc_{max}	$\Delta pos_{c,max}$	Δr_{LE}	Δpos_{TE}	Δt_{TE}
$\log_{10}(\theta_\eta)$	-2.38	-10.00	-2.38	-1.88	-4.41	-4.62	-2.20	-2.51	-3.39	-2.34	-2.73	-3.92	-2.68	-1.14
$\log_{10}(\theta_\phi)$	-2.60	-10.00	-1.47	-0.96	-10.00	-3.86	-1.52	-1.81	-3.52	-1.25	-3.05	-5.82	-2.69	-1.50
$\log_{10}(\theta_w)$	-3.65	-10.00	-2.78	-2.07	-10.00	-4.36	-2.46	-3.04	-4.32	-2.40	-4.44	-5.42	-3.94	-2.27

C. Optimization Results

Following Sec. IV. A’s findings that artificial blades match the real scans performance and the sensitivity analysis and parameter reduction study in Sec. IV. B, the main optimization was carried out. First, a 400-point space-filling Latin hypercube was generated and evaluated. Iman and Conover’s method is applied to all noise variables and this does not interfere with the space-filling structure of the Latin hypercube. 14 update points were requested in every iteration. A threshold of 800 points in total was chosen. For efficiency and accuracy, statistics are computed with 100 UQ samples at the start and up to 1000, with incremental increases by 50 samples in between. Design parameter update points were selected from predicted Pareto front and combined with space-filling noise parameter update points to construct proper blades. For NSGA-II, the number of generations was set to 300 and the population size to 100. Also, it should be noted that the optimization was carried out with all 14 noise variables. The final result of the workflow is the Pareto front depicted in Fig. 16. Frequent constraint violations are particularly evident. The axis grid line spacing reveals that mean efficiency is two orders of magnitude larger than the standard deviation. The Krig-predicted performance of the selected Pareto-optimal point is predicted to improve by 0.11 % for the mean and 127.25 % for the standard deviation (compared to the nominal design’s performance) and the constraints are predicted to be met. The mean performance improvement is also much larger than the known epistemic uncertainties. A separate validation study is required to verify these surrogate-based predictions, however.

As a first step towards validation, it is straightforward to assess all surrogate model accuracies by comparing the predictions to the baseline design’s known performance. Based on 1000 samples, the surrogate

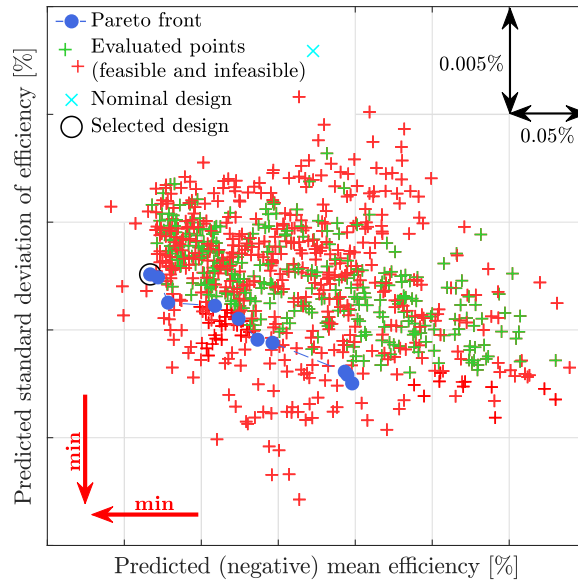


Figure 16 Final predicted Pareto front in objective space.

can be used to predict mean efficiency and standard deviation of the nominal blade. Respective percentage errors are 0.03 % and -25.67 %. The prediction of the mean efficiency is hence very accurate, whereas the standard deviation is underpredicted. The surrogate model error metric estimates in Table 3 also suggest that discrepancies exist. The efficiency surrogate is the least accurate of the three with an MAE of 14.88 %. Update points were added only to the Pareto front and no space-filling update points were added to improve the global accuracy, so it comes as no surprise that the final surrogates are found to be less accurate.

Table 3 Cross-validated surrogate accuracy metrics.

Error metric	η_{is}	ϕ	w
Beginning (400 points)			
R ² [%]	95.7897	99.9620	97.2490
RMSE [%]	3.5985	0.2768	3.0048
MAE [%]	15.2438	1.7487	10.2997
End (800 points)			
R ² [%]	97.5879	99.933	97.5152
RMSE [%]	3.1168	0.0311	2.3619
MAE [%]	14.8887	2.3565	10.6251

D. Validation of Results

Any surrogate model's predictions must be validated, especially given the estimated error in the standard deviation's prediction of 25.67%. Using the rotor's black box function and one of the designs from the Pareto front, it has to be proven *without* the surrogates that the probabilistic performance really is better than that of the baseline design and that it does indeed satisfy the constraints. The encircled design in Fig. 16 is selected for this purpose. Figure 17b depicts the selected blade's shape. The comparison with the nominal model's shape in Fig. 17c highlights the new design.

100 correlated samples for this selected design were created and evaluated with full CFD to test the probabilistic response, which is depicted in Fig. 18. The comparison to the baseline design's response shows that the mean efficiency has indeed been improved by 0.11% but the standard deviation in actual fact by 32.79%. Then, the surrogates predictions can also be taken into account. The improvement of the standard deviation was overpredicted by the surrogates with 127.25%, which is far too optimistic, whereas the mean performance prediction is excellent in comparison. The difference between the predicted and the actual improvement is only 2.12% and -16.88% for mean and standard deviation respectively. In addition, it is evident in Fig. 18 that all constraints are satisfied and that the optimizer has driven them very close to the upper and lower constraint limits. One likely cause for the differences is that the improvement of the standard deviation is more affected by UQ sampling noise than the mean (cf. Fig. 12). Also, the real noise model is more complex to approximate for the surrogates than the simple shape in Sec. III. A's analytical example. Compared to 1000 points in the analytical example with only four dimensions, i.e., 250 points per dimension, now 800 points are used to approximate a much larger space. Consequently, it is no surprise that the surrogates are somewhat inaccurate. Nevertheless, if so desired, prediction errors could be reduced by adding more points, e.g., space-filling update points. It should also be stressed that the practical purpose of surrogate models is to guide the optimizer to regions with better designs and that the surrogate's accuracy itself is not a priority as long as a superior design can be found, as has been the case here. All in all, the entire approach has been shown to be a viable option for robust design optimization. While standard deviation predictions are unmistakably less accurate than with the simple analytical noise model, a robust design with respectable performance gains for both objectives has been found. The mean constraints were accurately modeled too. Validation results have shown that the probabilistic performance prediction must be verified in any case.




Figure 17 Turbine blade shape comparison.




Figure 18 Robustness assessment of baseline (samples ■ and mean —) and selected robust blade (samples ■ and mean —) relative to the nominal CAD design intent (—).

V. Conclusions

A complete workflow for robust design optimization of turbine blades given real manufacturing variations was presented. The performance impact of these deviations was analyzed and assessed. It was shown that artificial blades with the same performance statistics as the original scans can be constructed and used for UQ. Through a sensitivity analysis, it was shown that prior to optimization parameters could be removed from the model. Finally, a robust aerodynamic optimization of a turbine rotor was presented. Validation of one selected Pareto-optimal design showed that both probabilistic objectives were improved and that tight constraints for average capacity and specific work were met. In the future, the workflow could be extended to include other geometric features such as film cooling holes and their positional variations, heat transfer computations and a cold-to-hot transformation.

Acknowledgements

The authors gratefully acknowledge the financial support of Rolls-Royce plc and the UK's Engineering and Physical Sciences Research Council under grant EP/L015382/1. The authors also thank TU Dresden's Chair of Turbomachinery and Jet Propulsion and Rolls-Royce Germany for the provision of their blade scans and in-house software. Special thanks are also due to Rolls-Royce for the permission for publication.

References

- [1] Montomoli, F., Carnevale, M., D'Ammaro, A., Massini, M., and Salvadori, S., *Uncertainty Quantification in Computational Fluid Dynamics and Aircraft Engines*, 1st ed., Springer, Cham, 2015.
doi:[10.1007/978-3-319-14681-2](https://doi.org/10.1007/978-3-319-14681-2)
- [2] Scharfenstein, J., Heinze, K., Voigt, M., Vogeler, K., and Meyer, M., "Probabilistic CFD Analysis of High-pressure Turbine Blades Considering Real Geometric Effects," *ASME Turbo Expo 2013: Turbine Technical Conference and Exposition*, American Society of Mechanical Engineers, 2013, pp. V06BT43A005–V06BT43A005.
doi:[10.1115/GT2013-94161](https://doi.org/10.1115/GT2013-94161)
- [3] Duffner, J. D., "*The Effects of Manufacturing Variability on Turbine Vane Performance*", Ph.D. Thesis, Massachusetts Institute of Technology, 2008.
doi:[1721.1/57599](https://doi.org/1721.1/57599)
- [4] Högner, L., Voigt, P., Nasuf, A., Voigt, M., Meyer, M., Goenaga, F., Berridge, C., and Vogeler, K., "Analysis of High Pressure Turbine Nozzle Guide Vanes Considering Geometric Variations," *ASME Turbo Expo 2016: Turbine Technical Conference and Exposition*, American Society of Mechanical Engineers, 2016, p. V02CT45A025.
doi:[10.1115/GT2016-57502](https://doi.org/10.1115/GT2016-57502)
- [5] Keane, A. and Nair, P., *Computational Approaches for Aerospace Design: The Pursuit of Excellence*, 1st ed., John Wiley & Sons, Chichester, 2005, pp. 327-329.
doi:[10.1002/0470855487](https://doi.org/10.1002/0470855487)
- [6] Heinze, K., Meyer, M., Scharfenstein, J., Voigt, M., and Vogeler, K., "A Parametric Model for Probabilistic Analysis of Turbine Blades Considering Real Geometric Effects," *CEAS Aeronautical Journal*,

Vol. 5, No. 1, 2014, pp. 41–51.

doi:[10.1007/s13272-013-0088-6](https://doi.org/10.1007/s13272-013-0088-6)

- [7] Voutchkov, I., and Keane, A., “Multi-objective Optimization Using Surrogates,” *Computational Intelligence in Optimization*, edited by Y. Tenne and C.-K. Goh, Springer, Berlin Heidelberg, 2015, pp. 155–175.
doi:[10.1007/978-3-642-12775-5](https://doi.org/10.1007/978-3-642-12775-5)
- [8] Bestle, D., Flassig, P. M., and Dutta, A. K., “Robust Design of Compressor Blades in the Presence of Manufacturing Noise,” *9th European Conference on Turbomachinery, Fluid Dynamics and Thermodynamics*, Vol. 1, edited by Sen, M., Bois, G., Manna, M. and Arts, T., European Turbomachinery Society - Euroturbo, Florence, Italy 2011, pp. 1303–1314, <https://opus4.kobv.de/opus4-UBICO/frontdoor/index/index/docId/7198>.
- [9] Kolmakova, D., Baturin, O., and Popov, G., “Effect of Manufacturing Tolerances on the Turbine Blades,” *ASME 2014 Gas Turbine India Conference*, American Society of Mechanical Engineers, 2014, pp. V001T02A004–V001T02A004.
doi:[10.1115/GTINDIA2014-8253](https://doi.org/10.1115/GTINDIA2014-8253)
- [10] Vinogradov, K., Kretinin, G., Otryahina, K., Didenko, R., Karelin, D., and Shmotin, Y., “Robust Optimization of the HPT Blade Cooling and Aerodynamic Efficiency,” *ASME Turbo Expo 2016: Turbine Technical Conference and Exposition*, American Society of Mechanical Engineers, 2016, p. V02CT45A003.
doi:[10.1115/GT2016-56195](https://doi.org/10.1115/GT2016-56195)
- [11] Javed, A., Pecnik, R., and van Buijtenen, J., “Optimization of a Centrifugal Compressor Impeller for Robustness to Manufacturing Uncertainties,” *Journal of Engineering for Gas Turbines and Power*, Vol. 138, No. 11, 2016, pp. 112101.
doi:[10.1115/1.4033185](https://doi.org/10.1115/1.4033185)
- [12] Keane, A. J., “Comparison of Several Optimization Strategies for Robust Turbine Blade Design,” *Journal of Propulsion and Power*, Vol. 25, No. 5, 2009, pp. 1092–1099.
doi:[10.2514/1.38673](https://doi.org/10.2514/1.38673)

- [13] Keane, A. J., “Co-Kriging for Robust Design Optimization,” *AIAA Journal*, Vol. 50, No. 12, 2012, pp. 2351–2364.
doi:[10.2514/1.J051391](https://doi.org/10.2514/1.J051391)
- [14] Seshadri, P., Shahpar, S. and Parks, G. T., “Robust Compressor Blades for Desensitizing Operational Tip Clearance Variations” *ASME Turbo Expo 2014: Turbine Technical Conference and Exposition*, American Society of Mechanical Engineers, 2017, pp. V02AT37A043–V02AT37A043.
doi:[10.1115/GT2014-26624](https://doi.org/10.1115/GT2014-26624)
- [15] Deb, K., Pratap, A., Agarwal, S. and Meyarivan, T., “A Fast and Elitist Multiobjective Genetic Algorithm NSGA-II,” *IEEE Transactions on Evolutionary Computation*, Vol. 6, No. 2, 2002, pp. 182–197.
doi:[10.1109/4235.996017](https://doi.org/10.1109/4235.996017)
- [16] Xiu, D. and Karniadakis, G. E., “Modeling Uncertainty in Flow Simulations via Generalized Polynomial Chaos,” *Journal of Computational Physics*, Vol. 187, No. 1, 2003, pp. 137–167.
doi:[10.1016/S0021-9991\(03\)00092-5](https://doi.org/10.1016/S0021-9991(03)00092-5)
- [17] Das, I., “Robustness Optimization for Constrained Nonlinear Programming Problems,” *Engineering Optimization+ A35*, Vol. 32, No. 5, 2000, pp. 585–618.
doi:[10.1080/03052150008941314](https://doi.org/10.1080/03052150008941314)
- [18] Högner, L., Voigt, M., Vogeler, K., Meyer, M., and Berridge, C., “A Curvature Based Algorithm for Treatment of Cooling Holes in Polygon Meshes of Turbine Blades,” *ASME Turbo Expo 2015: Turbine Technical Conference and Exposition*, American Society of Mechanical Engineers, 2015, p. V02BT39A025.
doi:[10.1115/GT2015-42841](https://doi.org/10.1115/GT2015-42841)
- [19] Lee, K.-H. and Park, G.-J., “A Global Robust Optimization using Kriging Based Approximation Model,” *JSME International Journal Series C Mechanical Systems, Machine Elements and Manufacturing*, Vol. 49, No. 3, 2006, pp. 779–788.
doi:[10.1299/jsmec.49.779](https://doi.org/10.1299/jsmec.49.779)
- [20] Dellino, G., Kleijnen, J. P. C., Meloni, C., and Meloni, C., “Metamodel-Based Robust Simulation-Optimization: An Overview,” *Uncertainty Management in Simulation-Optimization of Complex Systems: Algorithms and Applications*, 1st ed., edited by G. Dellino and C. Meloni, Springer US, New

York, 2015, pp. 27–54.

doi:[10.1007/978-1-4899-7547-8](https://doi.org/10.1007/978-1-4899-7547-8)

- [21] Asmussen, S. and Glynn, P. W., “Stochastic Simulation: Algorithms and Analysis”, Springer Science & Business Media, New York, 2007, p. 265.

doi:[10.1007/978-0-387-69033-9](https://doi.org/10.1007/978-0-387-69033-9)

- [22] Forrester, A., Sobester, A., and Keane, A., “Engineering Design via Surrogate Modelling: A Practical Guide”, John Wiley & Sons, Chichester, 2008, pp. 51-56, 196.

doi:[10.1002/9780470770801](https://doi.org/10.1002/9780470770801)

- [23] Lange, A., Voigt, M., Vogeler, K., and Johann, E., “Principal Component Analysis on 3D Scanned Compressor Blades for Probabilistic CFD Simulation,” *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, American Institute of Aeronautics and Astronautics, 2012, pp. 1–16.

doi:[10.2514/6.2012-1762](https://doi.org/10.2514/6.2012-1762)

- [24] Garzon, V. E. and Darmofal, D. L., “Impact of Geometric Variability on Axial Compressor Performance,” *Journal of Turbomachinery*, Vol. 125, No. 4, 2003, pp. 692–703.

doi:[10.1115/1.1622715](https://doi.org/10.1115/1.1622715)

- [25] Milli, A. and Shahpar, S., “PADRAM: Parametric Design and Rapid Meshing System for Complex Turbomachinery Configurations,” *ASME Turbo Expo 2012: Turbine Technical Conference and Exposition*, American Society of Mechanical Engineers, 2012, pp. 2135–2148.

doi:[10.1115/GT2012-69030](https://doi.org/10.1115/GT2012-69030)

- [26] Botev, Z. I., Grotowski, J. F., Kroese, D. P., “Kernel Density Estimation via Diffusion,” *The Annals of Statistics*, Vol. 38, No. 5, 2010, pp. 2916–2957.

doi:[10.1214/10-AOS799](https://doi.org/10.1214/10-AOS799)

- [27] Iman, R. L. and Conover, W.-J., “A Distribution-free Approach to Inducing Rank Correlation Among Input Variables,” *Communications in Statistics-Simulation and Computation*, Vol. 11, No. 3, 1982, pp. 311–334.

doi:[10.1080/03610918208812265](https://doi.org/10.1080/03610918208812265)

- [28] Gatski, T. B. and Bonnet, J.-P., *Compressibility, Turbulence and High Speed Flow*, 2nd ed., Elsevier Science, Oxford, UK, 2013, pp. 21-59.
doi:[10.1016/B978-0-12-397027-5.00009-5](https://doi.org/10.1016/B978-0-12-397027-5.00009-5)
- [29] Spalart, P. R. and Allmaras, S. R., “A One-equation Turbulence Model for Aerodynamic Flows,” *30th Aerospace Sciences Meeting and Exhibit, Aerospace Sciences Meetings*.
doi:[10.2514/6.1992-439](https://doi.org/10.2514/6.1992-439)
- [30] Beard, P. F., Smith, A. D., and Povey, T., “Effect of Combustor Swirl on Transonic High Pressure Turbine Efficiency,” *Journal of Turbomachinery*, Vol. 136, No. 1, 2014, pp. 011002.
doi:[10.1115/1.4024841](https://doi.org/10.1115/1.4024841)
- [31] Maffulli, R. and He, L., “Impact of Wall Temperature on Heat Transfer Coefficient and Aerodynamics for Three-Dimensional Turbine Blade Passage,” *Journal of Thermal Science and Engineering Applications*, Vol. 9, No. 4, 2017, pp. 041002.
doi:[10.1115/1.4036012](https://doi.org/10.1115/1.4036012)
- [32] Shahpar, S. and Caloni, S., “Aerodynamic Optimization of High-pressure Turbines for Lean-burn Combustion System,” *Journal of Engineering for Gas Turbines and Power*, Vol. 135, No. 5, 2013, pp. 055001.
doi:[10.1115/1.4007977](https://doi.org/10.1115/1.4007977)
- [33] Iooss, B. and Lemaître, P., “A Review on Global Sensitivity Analysis Methods,” *Uncertainty Management in Simulation-Optimization of Complex Systems: Algorithms and Applications*, edited by G. Dellino and C. Meloni, 1st ed., Springer, Cham, 2015, pp. 101–122.
doi:[10.1007/978-1-4899-7547-8](https://doi.org/10.1007/978-1-4899-7547-8)