

Gradient-enhanced Least-square Polynomial Chaos Expansions for Uncertainty Quantification and Robust Optimization

Tiziano Ghisu* and Diego Ignacio Lopez†
Università degli Studi di Cagliari, Cagliari, 09123, Italy

Pranay Seshadri‡
Imperial College London, London SW7 2AZ, UK

Shahrokh Shahpar§
Rolls-Royce plc, Innovation Hub, Central Technology, Derby, DE24 8BJ, UK

Regression-based Polynomial Chaos expansions offer several advantages over projection-based approaches, including their lower computation cost and greater flexibility. In the presence of expensive function evaluations, such as with computational fluid dynamics and finite element analysis, the availability of gradient information, coming from adjoint solvers, can be used to reduce the cost of least-square estimation. Particular attention needs to be paid to the accuracy of gradient information, as adjoint solvers are often more noisy than their primal counterparts. This paper compares different approaches for gradient-enhanced least-square Polynomial Chaos expansion, both for algebraic test cases, and for real-world test cases, i.e. a transonic compressor and a modern jet engine fan.

I. Introduction

Adjoint, or rapid gradient computations, form an integral component of modern computational engineering. These have, and continue to play a role in design optimization [1, 2], uncertainty quantification [3] and response surface generation [4]. Essentially, one hopes to be able to estimate more functional characteristics of a model quantity of interest (qoi) by its gradients, as oppose to using only the qoi values. For polynomial approximations, this recovery can imply the ability to approximate the coefficients of more basis terms, or the ability to have a higher degree of exactness. There is a subtle difference in the two, in that the former may not necessarily guarantee that the off-diagonal elements of the associated Gram matrix, particularly at the first few rows and columns, are zero. For expensive model evaluations, as in computational fluid dynamics (CFD) or finite element analysis (FEA), the availability of gradient information can lead to a lower cost of approximating the functional. Gradient evaluations, however, are often more noisy than function evaluation, and this needs to be considered if the objective is to attain a similar degree of exactness as with methods that employ only function evaluations.

This manuscript aims to compute the coefficients of a Polynomial Approximation via least squares; a compressed sensing exposition can be found in Peng *et al.* [5]. The appeal of least squares techniques for solving this problem, is two-fold. First, it permits more freedom in basis selection: a user need not stick to tensor or sparse grid basis, for which integration rules are well known [6, 7]. Second, there exists a vast and rich literature on least squares perturbation theory which may be used to infer error estimates (e.g. see Chapter 7 of [8] and references therein).

A. Gradient enhanced least squares Polynomial Chaos (PC) approximations

The main contribution of this paper is the use of gradient evaluations for constructing Polynomial Chaos expansions based on a least squares approximation of the coefficients. The general approach for achieving this in literature (see [9, 10]) has been to formulate the problem as a stacked least squares problem

$$\tilde{\mathbf{k}} = \underset{\mathbf{k}}{\operatorname{argmin}} \left\| \begin{pmatrix} \mathbf{A} \\ \mathbf{C} \end{pmatrix} \mathbf{k} - \begin{pmatrix} \mathbf{b} \\ \mathbf{d} \end{pmatrix} \right\|_2, \quad (1)$$

*Associate Professor, Department of Mechanical, Chemical and Materials Engineering

†PhD Student in Industrial Engineering

‡Research Fellow - Faculty of Natural Sciences, Department of Mathematics

§RR Engineering Fellow - Aerothermal Design Systems, AIAA Fellow, FRAeS

where matrices \mathbf{A} and \mathbf{C} encode the polynomial and polynomial-derivative bases, and vectors \mathbf{b} and \mathbf{d} are weighted evaluations of the model and its gradient, respectively. Entries in \mathbf{x} are the coefficients of the polynomial basis, and their estimation is one of the goals of this work.

With respect to the traditional stacked approach, in the presence of noisy gradient evaluations, which can be expected due to numerical discretization error, a constrained least squares formulation can yield more accurate coefficient estimates. This is illustrated in the numerical examples in section V. All codes to replicate our numerical results can be found in www.effective-quadratures.org/papers.

II. Notation and preliminaries

The aim is to find a polynomial approximation for a qoi $f(\mathbf{x})$ in the form:

$$f(\mathbf{x}) \approx \sum_{j=1}^n k_j \psi_j(\mathbf{x})$$

where $\mathbf{x} = (x^{(1)}, \dots, x^{(d)})$ be a d -dimensional vector of mutually independent random variables with joint probability density ρ and marginal densities ρ_i , given by $\rho(\mathbf{x}) = \prod_{i=1}^d \rho_i(x^{(i)})$. $\{\psi_j^{(i)}\}_{j=0}^{\infty}$ is a family of polynomials L^2 -orthogonal on \mathbb{R} when weighted by the density ρ_i :

$$\int_{\mathbb{R}} \psi_i^{(k)}(s) \psi_j^{(k)}(s) \rho_k(x) ds = \delta_{i,j},$$

where $\delta_{i,j}$ is the Kronecker delta function. Multivariate polynomials $\psi_j : \mathbb{R}^d \rightarrow \mathbb{R}$ can be calculated as products of univariate polynomials,

$$\psi_j(\mathbf{x}) = \prod_{k=1}^d \psi_{j_k}^{(k)}(x^{(k)}), \quad \mathbf{j} = (j_1, \dots, j_d) \in \mathbb{N}_0^d,$$

where \mathbf{j} is a multi-index that denotes the order of ψ_j and its composite univariate polynomials $\psi_{j_k}^{(k)}$. These multi-indices lie within a multi-index set denoted by \mathcal{J} , which may be either be a total order or a hyperbolic cross (see [11–13]).

A high-cardinality quadrature rule $\{(\mathbf{x}_i, \omega_i^2)\}_{i=1}^m$, with non-negative weights ω_i^2 , can be calculated for generic probability measures using the Stieltjes procedure and the Chebyshev algorithm [14, 15].

$$\sum_{i=1}^m \omega_i^2 \psi_j(\mathbf{x}_i) \psi_\ell(\mathbf{x}_i) = \int_{\mathbb{R}^d} \psi_j(\mathbf{x}) \psi_\ell(\mathbf{x}) \rho(\mathbf{x}) d\mathbf{x} = \delta_{\ell,j}, \quad \mathbf{j}, \ell \in \mathcal{J}. \quad (2)$$

III. Sampling strategies

Different approaches are available for estimating the coefficients of the expansion in (2). Differently from projection-based methods [16], a least-square regression seeks an approximation for the coefficients of the expansion in the following form:

$$\tilde{\mathbf{k}} = \underset{\mathbf{k}}{\operatorname{argmin}} \quad \|\mathbf{A}\mathbf{k} - \mathbf{b}\|_2. \quad (3)$$

A. Definition of \mathbf{A} and \mathbf{b}

Let $\mathbf{A}_{\mathcal{F}} \in \mathbb{R}^{r \times n}$ be a matrix formed by evaluating n multivariate orthonormal polynomials at a tensor grid formed of r quadrature points. Thus, individual entries of $\mathbf{A}_{\mathcal{F}}$ are given by

$$\mathbf{A}_{\mathcal{F}}(i, \mathbf{j}) = \omega_i \psi(\mathbf{x}_i), \quad \mathbf{j} \in \mathcal{J}, \quad i = 1, \dots, r,$$

where w_i are the multivariate tensor grid quadrature weights and $\psi(\mathbf{x}_i)$ is the multivariate orthonormal polynomial evaluated at the quadrature point given by \mathbf{x}_i . The index set \mathcal{J} is assumed to satisfy the condition that $|\mathcal{J}| = n \ll r$, which depends on the choice of the basis and of the quadrature rule. Matrix $\mathbf{A}_{\mathcal{F}}$ is arranged based on the weight

associated with the corresponding quadrature points, although the particular ordering chosen is not relevant for the purposes of this work.

$\mathbf{b}_{\mathcal{F}} \in \mathbb{R}^r$ is then defined to be the vector of weighted model evaluations at the r quadrature points, where individual entries of $\mathbf{b}_{\mathcal{F}}$ are given by

$$\mathbf{b}_{\mathcal{F}}(i) = \omega_i f(\mathbf{x}_i), \quad (4)$$

where $f(\cdot)$ represents the quantity of interest from the computational model.

Given that $\mathbf{A}_{\mathcal{F}}$ is a *tall* matrix, where the number of rows r determines the number of computational evaluations defined by the quadrature points, one can reduce the number of rows and still obtain accurate estimates of the n coefficients. Different approaches can be used to reduce the number of rows in $\mathbf{A}_{\mathcal{F}}$, ranging from random sampling [17, 18] to optimal subsampling by means of QR column pivoting [19, 20].

IV. Gradient enhancement

If gradient information is available, the regular least squares approach in ((3)) can be modified to include derivative information. The straightforward solution is to employ a stacked approach, i.e.

$$\tilde{\mathbf{k}} = \underset{\mathbf{k}}{\operatorname{argmin}} \left\| \begin{pmatrix} \mathbf{A} \\ \mathbf{C} \end{pmatrix} \mathbf{k} - \begin{pmatrix} \mathbf{b} \\ \mathbf{d} \end{pmatrix} \right\|_2, \quad (5)$$

A. Definition of \mathbf{C} and \mathbf{d}

Let $\mathbf{C}_{\mathcal{F}}^{(k)} \in \mathbb{R}^{r \times n}$ for $k = 1, \dots, d$ be a set of matrices formed by evaluating the derivatives of the multivariate orthonormal polynomial at a tensor grid formed of the same p quadrature points. Here d represents the dimensionality of our problem—i.e., for approximating a polynomial of two-dimensions, $d = 2$. Individual entries of $\mathbf{C}_{\mathcal{F}}^{(k)}$ are given by

$$\mathbf{C}_{\mathcal{F}}^{(k)}(i, j) = \frac{\partial \psi(\mathbf{x}_i)}{\partial x^{(k)}}, \quad j \in \mathcal{J}, \quad i = 1, \dots, r,$$

where $\frac{\partial \psi(\mathbf{x}_i)}{\partial x^{(k)}}$ is derivative of the multivariate polynomial $\psi(\mathbf{x}_i)$ with respect to the *parameter* $x^{(k)}$. Having access to f 's gradients a set of vectors $\mathbf{d}_{\mathcal{F}}^{(k)} \in \mathbb{R}^{p d}$ can be defined to be the vector of gradient evaluations at the r quadrature points, where individual entries are given by

$$\mathbf{d}_{\mathcal{F}}^{(k)}(i) = \frac{\partial f(\mathbf{x}_i)}{\partial x^{(k)}}, \quad i = 1, \dots, r, \quad (6)$$

$\mathbf{C}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{F}}$ can then be determined as follows:

$$\mathbf{C}_{\mathcal{F}} = \begin{pmatrix} \mathbf{C}_{\mathcal{F}}^{(1)} \\ \vdots \\ \mathbf{C}_{\mathcal{F}}^{(d)} \end{pmatrix} \quad \text{and} \quad \mathbf{d}_{\mathcal{F}} = \begin{pmatrix} \mathbf{d}_{\mathcal{F}}^{(1)} \\ \vdots \\ \mathbf{d}_{\mathcal{F}}^{(d)} \end{pmatrix}, \quad (7)$$

Similarly to \mathbf{A} and \mathbf{b} , \mathbf{C} and \mathbf{d} can be determined from $\mathbf{C}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{F}}$ by subsampling the quadrature points, following a given heuristic.

B. Determining the minimum number of subsamples t

The number of subsamples t to be used in the construction of matrices \mathbf{A} , \mathbf{b} , \mathbf{C} and \mathbf{d} can be determined following an iterative procedure, outlined in Algorithm 1, required for ensuring that $\operatorname{rank} \begin{pmatrix} \mathbf{A} \\ \mathbf{C} \end{pmatrix} \equiv n$.

The algorithm requires two inputs, namely the number of basis coefficients $n > 1$, and the dimensionality of the problem. It starts by generating matrices $\mathbf{A}_{\mathcal{F}}$ and $\mathbf{C}_{\mathcal{F}}^{(k)}$ for $k = 1, \dots, d$ with size $r \times n$. Quadrature points and corresponding rows in $\mathbf{A}_{\mathcal{F}}$ and $\mathbf{C}_{\mathcal{F}}^{(k)}$ are then selected iteratively only if they increase the rank of matrix \mathbf{A} and \mathbf{C} , until

Algorithm 1 Iterative algorithm for determining the optimal subsamples

```

1: procedure  $\zeta = \text{OPTIMAL\_SUBSAMPLES}(n, d)$ 
2:   Compute  $A_{\mathcal{F}}$  and  $C_{\mathcal{F}}^{(k)}$  for  $k = 1, \dots, d$ 
3:   Set rank of stacked matrix  $R = 0$ ,  $\zeta = []$ 
4:   while  $R < n$  do
5:     Select a random quadrature point  $i$ 
6:     Set  $\zeta_{new} = \zeta = [\zeta, i]$ 
7:     Select the quadrature points in  $\zeta_{new}$  to construct matrices  $A$  and  $C$ 
8:     Calculate  $R_{new} = \text{rank} \begin{pmatrix} A \\ C \end{pmatrix}$ 
9:     if  $R_{new} > R$  then
10:        $R = R_{new}$ 
11:        $\zeta = \zeta_{new}$ 
12:     end if
13:   end while
14:   return  $\zeta$ 
15: end procedure

```

the rank of the stacked matrix is equivalent to n . To reduce the effect of noise in the input data [13], some oversampling is then used by selecting a number of extra quadrature points. At the end of the procedure, the following matrices are obtained:

$$A \in \mathbb{R}^{t \times n}, \quad C \in \mathbb{R}^{(d \times t) \times n}, \quad \mathbf{b} \in \mathbb{R}^t, \quad \mathbf{d} \in \mathbb{R}^{dt}; \quad (8)$$

where t is the dimensionality of array ζ . A schematic of the methodology is given in Figure 1.

C. The noisy case

An alternative to solving the system in (5) when some level of noise in the gradient evaluations is expected is the following:

$$\begin{aligned} \tilde{\mathbf{k}} &= \underset{\mathbf{k}}{\text{argmin}} \quad \|\mathbf{C}\mathbf{k} - \mathbf{d}\|_2. \\ &\text{subject to} \quad \mathbf{A}\mathbf{k} = \mathbf{b}. \end{aligned} \quad (9)$$

The solution to (9) is unique if and only if the null spaces of A and C intersect trivially;

$$\mathcal{N}(A) \cap \mathcal{N}(C) = \{\mathbf{0}\}. \quad (10)$$

This is equivalent to requiring that $\text{rank} \begin{pmatrix} A \\ C \end{pmatrix} = n$ [21]. As discussed in Chapter 2 of Björck [22], there are two techniques for computing a solution to (9): the null space method and the direct elimination technique. The former has been selected in this work.

This method requires a QR factorization of A^T

$$\begin{pmatrix} Q_1 & Q_2 \end{pmatrix} R = A^T, \quad Q_1 \in \mathbb{R}^{n \times t}, \quad Q_2 \in \mathbb{R}^{n \times (n-t)}, \quad (11)$$

where Q_2 is an orthogonal basis for the null space of A . The solution vector can then be split into the sum of two orthonormal components

$$\mathbf{x} = Q_1 \mathbf{w}_1 + Q_2 \mathbf{w}_2 \quad (12)$$

where $\mathbf{w}_1 \in \mathbb{R}^t$ and $\mathbf{w}_2 \in \mathbb{R}^{(n-t)}$. Since A is full rank, the lower triangular matrix R^T is non-singular and the constraints

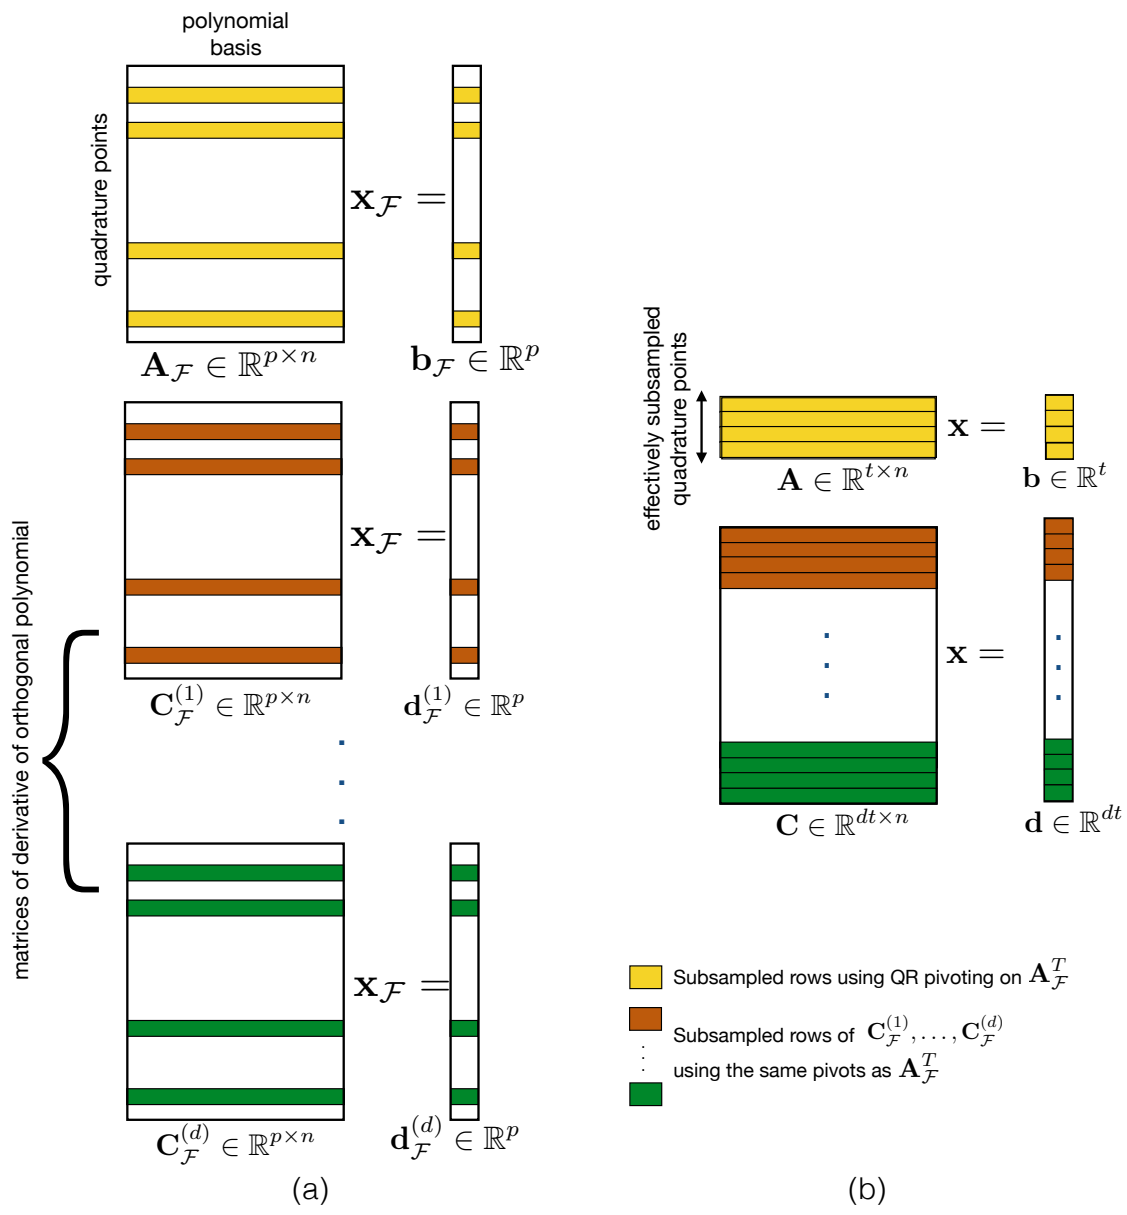


Fig. 1 Schematic outlining the relative sizes of the matrices

can be rewritten as $\mathbf{w}_1 = \mathbf{R}^{-T}\mathbf{b}$. Thus:

$$\begin{aligned}
\mathbf{C}\mathbf{x} - \mathbf{d} &= \mathbf{C}(\mathbf{Q}_1\mathbf{w}_1 + \mathbf{Q}_2\mathbf{w}_2) - \mathbf{d} \\
&= \mathbf{C}(\mathbf{Q}_1\mathbf{R}^{-T}\mathbf{b} + \mathbf{Q}_2\mathbf{w}_2) - \mathbf{d} \\
&= \mathbf{C}\mathbf{Q}_1\mathbf{R}^{-T}\mathbf{b} + \mathbf{C}\mathbf{Q}_2\mathbf{w}_2 - \mathbf{d} \\
&= \underbrace{\mathbf{C}\mathbf{Q}_2}_{\mathbf{G}}\mathbf{w}_2 - \underbrace{(\mathbf{d} - \mathbf{C}\mathbf{Q}_1\mathbf{R}^{-T}\mathbf{b})}_{\mathbf{h}}.
\end{aligned} \tag{13}$$

It follows that the original problem is equivalent to the unconstrained least squares problem

$$\underset{\mathbf{w}_2}{\text{minimize}} \quad \|\mathbf{G}\mathbf{w}_2 - \mathbf{h}\|_2, \tag{14}$$

after which \mathbf{x} can be computed using (12).

The condition number of the system is $\kappa(\mathbf{G}) = \kappa(\mathbf{C}\mathbf{Q}_2)$, which Eldén [23] states is $\leq \kappa(\mathbf{C})$, provided that \mathbf{A} is well-conditioned [22].

V. Numerical examples

The performance of the null-space method proposed is demonstrated in this section through two numerical examples, where it is compared against the standard stacked approach. The first example involves an analytical problem with the artificial addition of white noise in the gradient expression, and the second is a standard CFD turbomachinery test case involving a compressor blade which has traditionally posed challenges to simulate accurately due to its complex three-dimensional flow.

A. A 2D analytical problem

The first test case is an analytical problem, consisting in a bi-variate exponential function with uncertain variables x_1 and x_2 . Both variables have known probability density functions, uniformly distributed between [-1:1]. White noise of different magnitude is added to the gradients calculated analytically, and its impact on the evaluation of the coefficients of the expansion, as well as on mean and standard deviation, are assessed.

$$\begin{aligned}
f(x_1, x_2) &= \exp(2x_1 + x_2) \\
\nabla f &= \begin{pmatrix} 2\exp(2x_1 + x_2) + \kappa \\ \exp(2x_1 + x_2) + \kappa \end{pmatrix}.
\end{aligned} \tag{15}$$

Figure 2 presents the result of applying the stacked and linear equality approaches when the gradients calculation are exact (i.e. $\kappa = 0$). The coefficients of a 15th order polynomial expansion calculated with the two approaches are presented graphically, and both cases exhibit the decay typical for a polynomial approximation of an exponential function with uncertain variables. The global difference in the evaluation of the coefficients is of the order of 10^{-10} .

Figure 3 presents the solution of the same problem in the presence of noisy gradients. κ is now randomly selected from a normal distribution with zero mean and 0.01 standard deviation, i.e. $\kappa = \mathcal{N}(0, 0.01)$. Results in Figure 3 report the application of the two approaches to typical input data, as the values of the coefficients of the polynomial expansion depend both on the evaluation points selected using Algorithm 1, and on the actual gradient values. Figure 4 presents the distributions for the errors in the calculation of coefficients, mean and standard deviation. The advantages of the linear equality approach using the null-space method are evident, as there is a significant reduction of both the global error in the evaluation of the coefficients, and in the estimation of the standard deviation. The reduction of the error in the evaluation of the mean is more limited.

Figure 5 presents the solution of the same problem in the present of larger noise in the gradient evaluations. κ is now randomly selected from a normal distribution with zero mean and 0.1 standard deviation. As before, results in Figure 5 present the application of the two approaches to typical input data, while Figure 6 presents the distributions for the errors in the calculation of coefficients, mean and standard deviation. The advantages of the linear equality approach are evident, although as expected the errors are larger than in the previous case with $\kappa = \mathcal{N}(0, 0.01)$, due to the larger noise in the gradient evaluations.




Fig. 2 Here $\kappa = 0$ for the analytical problem in 15. Shown are the coefficients on a logarithm (base-10) scale using (a) the stacked approach; (b) the linear equality approach via the null-space method.




Fig. 3 $\kappa = \mathcal{N}(0, 0.01)$ for the analytical problem in 15. Shown are the coefficients on a logarithm (base-10) scale using (a) the stacked approach; (b) the linear equality approach via the null-space method.

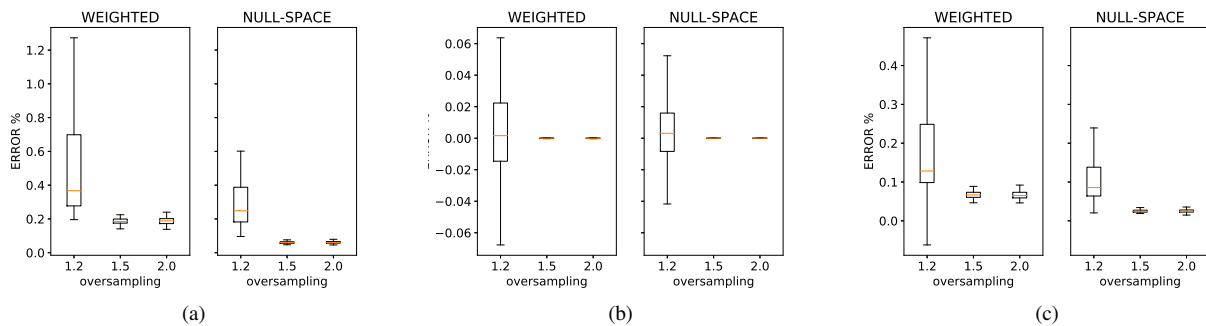


Fig. 4 $\kappa = \mathcal{N}(0, 0.01)$ for the analytical problem in 15. Shown are the errors (a) in the evaluation of the coefficients (total); (b) in the evaluation of the mean; (c) in the evaluation of the standard deviation.




Fig. 5 $\kappa = \mathcal{N}(0, 0.1)$ for the analytical problem in 15. Shown are the coefficients on a logarithm (base-10) scale using (a) the stacked approach;(b) the linear equality approach via the null-space method.

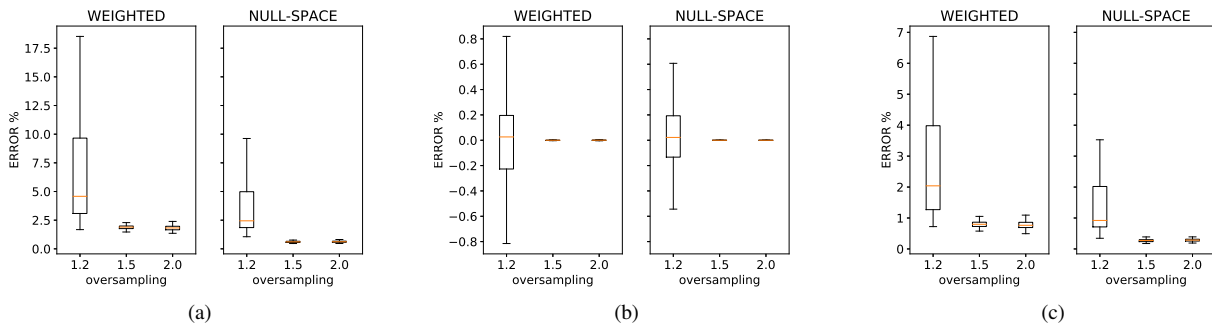


Fig. 6 $\kappa = \mathcal{N}(0, 0.1)$ for the analytical problem in 15. Shown are the errors (a) in the evaluation of the coefficients (total); (b) in the evaluation of the mean; (c) in the evaluation of the standard deviation.

B. NASA Rotor 37

As a further candidate for testing the performance of the method described, this study considered the NASA Rotor 37 [24]. It is a low aspect ratio rotor designed in the 1970s as the inlet stage of an 8 stage core compressor with an overall pressure ratio of 20 : 1. Having been tested in isolation, interaction effects on the blade can be avoided, which results in an ideal test case for validation. Further details of the stage characteristics can be found in [25].

Despite its geometrical simplicity and being subjected to a uniform inlet flow, CFD modelling of this compressor blade has historically been challenging. Airflow about the blade experiences complex and highly three-dimensional phenomena, mostly due to the stage's high pressure ratio, which results in a large tip-leakage vortex and strong shock wave-boundary layer interactions, leading to large regions of separated flow. Adequately capturing these flow features poses challenges to modern turbomachinery solvers, and whilst an accurate characterization of the compressor performance can be obtained through CFD, gradient evaluations through the adjoint of the Navier-Stokes equations are more inclined to result in a noisy response.

As such, the test case is well suited for comparing the accuracy of polynomial approximation via least squares vs more traditional approaches. The CFD set-up used in this study has been described in a previous publication [26]. The domain shown in Figure 7 is discretized with a fine multi-block structured grid. A radial distribution of total pressure and temperature is specified at the inlet, whilst at the exit, a radial equilibrium capacity value was specified. The reader is referred to [26] for further details and validation.




Fig. 7 Rotor 37 Domain [26]

Results

As an initial test case, 5 uncertainties were considered, representing a tangential movement of the blade section at different spanwise locations. The uncertainties have all been assumed to be normally distributed, while the rotor adiabatic efficiency has been considered as the qoi. Figure 8 reports the dependence of the qoi and its gradient from the 5 uncertainties. The curves represent the best quadratic fitting using 3 and 5 quadrature points, and the lines represent the values of the derivatives, calculated using finite differences (step size of 0.01) and the adjoint solver. It is clear that the derivatives are affected by some degree of noise, especially for component 1 (where the error is more random than in other cases, where it appears to be biased towards a specific direction).

Four different approaches have been compared for the calculation of mean and standard deviation of the qoi, and their derivatives with respect to the uncertain parameters. The four approaches include: a sparse grid numerical integration (SGNI) using only function values, which is taken as a reference, a sparse grid numerical integration using using gradient information (SGNI-G), and least-square approximation using the stacked approach (LSA-G) and the linear inequality approach via the null-space method (LSA-G-NS). All methods use a second-order total order bases for the polynomial approximation, apart from the SGNI-G approach that uses a first-order approximation, the rational being the similar cost of a first-order SGNI-G and second-order LSA-G and LSA-G-NS approaches. Detailed derivation of the methodologies used for calculating mean, standard deviation, and their gradients with the different approaches are summarized in [27].

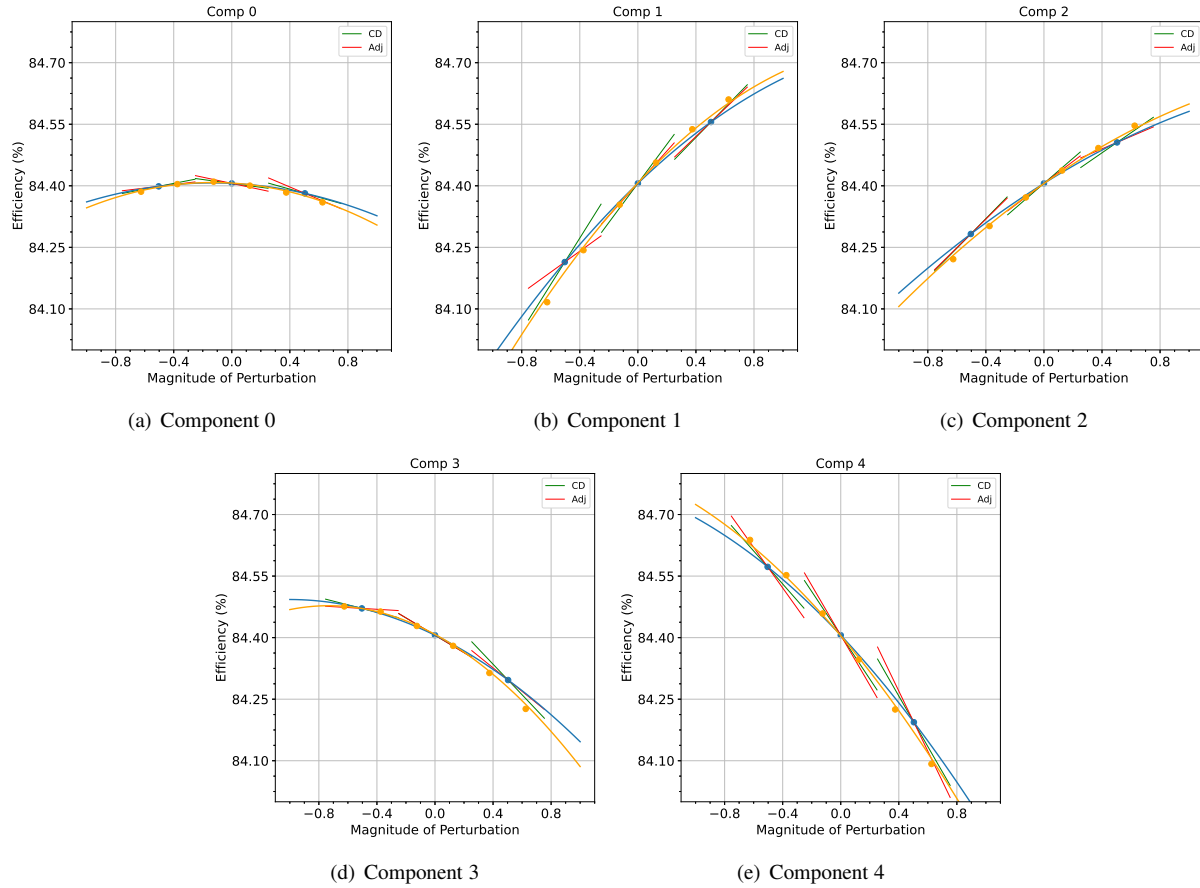



Fig. 8 Noise in the derivative information for the R37 test case

Figure 9 reports the coefficients of the PC expansion, calculated with the four approaches. The SGNI-G approach only estimates the first-order terms (up to coefficient 5), while the other methods estimate even second order terms. However, the estimate of the first order term in particular is incorrect for the stacked LSA-G approach. The inaccurate estimation of the Sobol indices, which represent the impact of each uncertainty on the variance of the qoi, is a direct consequence. The results are presented in Figure 10.




(a) 2nd-order SGNI without gradients vs 1st-order SGNI-G with gradients (b) 2nd-order SGNI without gradients vs 2nd-order LSA-G with gradients (c) 2nd-order SGNI without gradients vs 2nd-order LSA-G-NS with gradients

Fig. 9 Polynomial Chaos expansion coefficients comparison using different methodologies. LSA approaches are tested with different levels of oversampling (o-s)


Figure 11 and Figure 12 report the gradients of the first 2 moments, obtained with the different methods. SGNI-G

uses the gradient of the qoi to calculate the gradients for the moments, and produces important errors, which are due to the noise in the gradient seen in Figure 8. The noise in the gradient, however, does not affect the estimation of the moments, as these are calculated based on the qoi PC approximation. On the contrary, the LSA approaches use both qoi and qoi gradients' at the quadrature points to calculate the qoi PC approximation, and this leads to lower errors in the moments' gradients, both for the stacked and the null-space approaches, especially when employing some oversampling. Overall, considering the noise to be in the gradient evaluations and not in the QoI has a beneficial effect, as overall the predictions with the null-space LSA-NS are superior to the ones obtained by the stacked LSA approach.




(a) 2nd-order SGNI without gradients vs 1st-order SGNI with gradients (b) 2nd-order SGNI without gradients vs 2nd-order LSA with gradients (c) 2nd-order SGNI without gradients vs 2nd-order LSA-NS with gradients

Fig. 10 Sobol indices comparison using different methodologies. LSA approaches uare tested with different levels of oversampling (o-s)



(a) 2nd-order SGNI without gradients vs 1st-order SGNI with gradients (b) 2nd-order SGNI without gradients vs 2nd-order LSA with gradients (c) 2nd-order SGNI without gradients vs 2nd-order LSA-NS with gradients

Fig. 11 Gradient of the mean: comparison using different methodologies. LSA approaches uare tested with different levels of oversampling (o-s)



(a) 2nd-order SGNI without gradients vs 1st-order SGNI with gradients (b) 2nd-order SGNI without gradients vs 2nd-order LSA with gradients (c) 2nd-order SGNI without gradients vs 2nd-order LSA-NS with gradients

Fig. 12 Gradient of the standard deviation: comparison using different methodologies. LSA approaches uare tested with different levels of oversampling (o-s)

Table 1 Uncertainty quantification for the R37 efficiency with different approaches

	SGNI (2nd O.)	SGNI-G (1st O.)	LSA-G (2nd O., stacked)	LSA-G-NS (2nd O., null-space)
MEAN	84.310	84.313	84.316	84.313
STD	0.29312	0.29173	0.41204	0.29972
ERROR FOR MEAN (%)	-	0.003	0.015	0.003
ERROR FOR STD (%)	-	-0.476	4.914	1.763
	-0.013	-0.027	-0.043	-0.017
	0.329	0.386	0.293	0.339
GRADIENT OF MEAN	0.231	0.179	0.191	0.221
	-0.154	-0.030	-0.132	-0.173
	-0.373	-0.489	-0.442	-0.376
	0.007	0.007	0.003	0.012
	-0.065	0.054	-0.024	-0.068
GRADIENT OF STD	-0.026	-0.033	-0.029	-0.034
	0.031	0.014	0.004	0.028
	0.052	0.074	0.086	0.078
	-	-2.505	-5.217	-0.703
	-	10.069	-6.343	1.729
ERROR FOR GRADIENT OF MEAN	-	-9.218	-7.069	-1.688
	-	21.689	3.856	-3.466
	-	-20.309	-11.994	-0.456
GLOBAL ERROR FOR GRADIENT OF MEAN	-	63.790	34.479	8.042
	-	-0.267	-4.751	4.918
	-	128.396	43.742	-3.754
ERROR FOR GRADIENT OF STD	-	-7.659	-2.353	-8.275
	-	-18.471	-29.308	-3.337
	-	24.479	36.919	28.456
GLOBAL ERROR FOR GRADIENT OF STD	-	179.27	117.07	48.740

Table 1 summarizes the results of the different approaches for the uncertainty quantification of R37. It is evident how all methods produce a similar estimate for the mean of the efficiency, while the LSA-G with the stacked approach fails to correctly estimate the standard deviation. The sensitivities of mean and standard deviation with the null-space approach (LSA-G-NS) are also significantly more accurate, when compared to both SGNI-G and LSA-G methods

VI. Robust Optimization

Mitigating the numerical inaccuracies encountered with the stacked approach for the calculation of statistical moments through least-squares approximations provides significant advantages for industrial design applications. The linear inequality approach via the null-space method proposed enables more accurate estimations of the mean and standard deviation of performance metrics associated to a particular design, when subjected to uncertain conditions, to be obtained at a reduced expense, while maintaining the flexibility typical of least-square approaches. While traditional methods for optimizing under uncertainty have typically required designers to employ models of reduced complexity (e.g. 2D) to lessen the expense of functional evaluations [28], or consider a reduced number of uncertain parameters to limit the overall number of functional calls required [29, 30], the null-space LSA-G-NS approach provides greater flexibility in both these considerations through full exploitation of gradient information. To demonstrate this advantage, the Robust Design Optimization (RDO) of an industrially relevant problem is performed in this section. The case considered is a modern turbomachinery test case involving a low speed, high bypass ratio research fan blade, called VITAL.

A. Optimization Framework

For a function $f = f(\mathbf{x}, \omega)$ representing the response of a quantity of interest for a design parameterized by the design vector $\mathbf{x} \in \chi \subseteq \mathbb{R}^m$, and subjected to uncertain random variables $\omega \in \Omega \subseteq \mathbb{R}^k$, a general unconstrained RDO problem can be defined as a multi-objective optimization problem as per Eq. 16, where $\mu(\mathbf{x}, \omega)$ represents the mean value and $\sigma(\mathbf{x}, \omega)$, the standard deviation of f , respectively.

$$\begin{aligned}
& \underset{\mathbf{x}}{\text{minimise}} \quad \{\mu(\mathbf{x}, \omega), \sigma(\mathbf{x}, \omega)\} \\
& \mathbf{x} \in \chi \subseteq \mathbb{R}^m \\
& \omega \in \Omega \subseteq \mathbb{R}^k
\end{aligned} \tag{16}$$

Assuming the gradient $\mathbf{g}(\mathbf{x}, \omega) = \nabla f$ can be computed for all $\mathbf{x} \in \chi$ and all $\omega \in \Omega$, the value of the cost functions in (16) can be obtained by solving (14) through p quadrature sample points of f and \mathbf{g} .

Equation 16 involves two design considerations which may present conflicting optimality criteria. While the optimization of μ is aimed at identifying designs with statistically improved values of f , optimizing for σ identifies designs where the variability encountered in f due to the uncertain parameters is minimized. The trade-offs between these two objectives can be obtained by analysing the Pareto frontier of the design space, defined as the set of all solutions for which no other solution is better in all objectives.

Solving (16) typically requires numerical search algorithms that identify designs in the Pareto frontier. Multi-objective evolutionary algorithms are a popular choice for this task, since they are able to identify non-dominated individuals in the generations and evolve the entire Pareto set. In this study, the Adaptive Range Multi-Objective Genetic Algorithm (ARMOGA) optimizer, present in the optimization library SOFT [31], was employed to perform the search.

Due to the large number of function calls required by evolutionary algorithms, they are commonly employed with response surfaces, models constructed through a reduced number of evaluations that predict the behavior of the cost functions in the design space. The authors of this work have proposed coupling Artificial Neural Networks (ANNs) with Active Design Subspaces (ADS) to enable highly accurate predictions over complex high-dimensional spaces while maintaining an overall low sample requirement [32]. In [32], the approach was applied to the optimization of the nominal performance of the Vital fan, while here it is extended to a robust optimization problem, using the same strategy to regress the optimization cost functions and perform the ARMOGA evaluations. A brief description of this methodology is presented in this section, for an in-depth analysis the reader is referred to [32].

1. AI-Enabled Active Subspaces

ANNs are selected to regress the quantity of interest because, as detailed in [33, 34], provided there are sufficient neurons in the network, ANNs are able to fit any function of arbitrary shape, which enables regressing complex multi-modal functions without making any prior assumptions on their behavior.

However, fitting neural networks to high-dimensional input data typically requires a large number of samples, which can make the method infeasible in cases where function evaluations are obtained through a costly procedure. Reductions in the expense of ANNs can be obtained by reducing the dimensionality of the design vector, \mathbf{x} . ADS [35] are adopted to learn a low-dimensional representation for the function, based on non-conditioned data sets. Through the eigenvalue decomposition of the function's gradient covariance matrix, the ADS approach identifies dominant linear combinations of the components of \mathbf{x} that best describe the variability in the function. Discovering the dominant linear combinations, or active vectors, is particularly useful as they enable simplifying the behavior of the function, thus reducing the number of samples required to accurately fit ANNs.

Notwithstanding of their advantages, learning the ADS can itself be an expensive procedure, as knowledge on the gradient of the function of interest is required, which may not always be available. In this work, we consider μ and σ , which themselves are obtained with gradient evaluations from f . Therefore, an analytical derivation of $\nabla\mu$ and $\nabla\sigma$ can require significant efforts to achieve.

The strategy enables application of the ADS through fitting an initial ANN based on samples from the high-dimensional input space and differentiating these models to obtain expressions for the gradient of the quantity of interest. Upon identification of the active vectors for each function, the original samples are mapped to the active subspaces and a secondary ANN is fitted to the low-dimensional inputs. Because the number of samples is kept constant, but the function's dimensionality is reduced and its behavior simplified, the accuracy of the second ANN can be expected to increase.

B. VITAL Fan Blade

The robust aerodynamic optimization of a modern low speed, high bypass ratio fan blade was considered due to the significant influence variations in this component's efficiency have on the overall Aero-engine's performance, with increases of 1.4% in fan efficiency yielding reductions of 1% in the engine's specific fuel consumption [36]. The VITAL

fan blade, shown in Figure 13, was designed for rig studies with a span about two-thirds smaller than conventional production blades, an adjusted rotational speed, designed to reproduce the flow physics at cruise condition and an increased thickness, intended to guarantee the mechanical integrity required of a large composite fan blade.




Fig. 13 VITAL CFD Domain [32]

1. Computational Tools

The function and gradient evaluations in this section were performed with the Rolls-Royce proprietary CFD suite Hydra [37]. The primal solver employs unstructured grids with an edge-based data structure and convergence acceleration through an element collapsing multi-grid algorithm. A five-stage Runge-Kutta scheme with a block Jacobi preconditioner is employed for pseudo time-stepping when solving the steady-state Reynolds-averaged Navier-Stokes equations. The turbulence closure model employed in this work is Spalart-Allmaras. The Hydra suite also includes a discrete adjoint solver, which was employed in this work to estimate the gradients of the quantity of interest.

The computational model, shown in Figure 13, is composed of a single passage, single blade row model with the downstream splitter. The domain is meshed using the Rolls-Royce proprietary code PADRAM [38], producing a multi-block structured grid with an H-O-H topology, where the blade is enveloped in an O-mesh, H-mesh blocks are used for the upstream and downstream regions and a C-mesh is employed for the splitter. A mesh convergence study was previously undertaken to identify the optimal distribution of nodes [39], leading to a total of 5.4×10^6 cells, placing 30 mesh nodes in the tip gap. The y^+ of the mesh is below 1 on all viscous surfaces.

The whole domain is modelled on a rotating frame with the rotor, casing, splitter, inlet and exit surfaces set as stationary. A uniform, 1-dimensional Boundary condition (BC) is specified at the inlet surface, prescribing radially-varying values for total pressure, total temperature, whirl, pitch and turbulence intensity. Non-dimensional capacity values are specified at the bypass and core exit surfaces while periodic boundaries are employed on the upper and lower surfaces to simulate adjacent passages. CFD-experimental validation of this set-up have previously been reported [39, 40], achieving a good match for a number of different operating conditions.

2. Geometry Parameterization

Due to variability in the manufacturing process, random deviations from the design intent fan blade are typically encountered on the final product. The source of these deformations can be difficult to isolate and mitigating them can require significant efforts and alterations to the production process. Therefore, it is advantageous to design fan blade geometries that are robust to manufacturing deformations and maintain a high performance with low variance.

Identification of the geometrical deviations in manufactured blades can be achieved through reverse engineering processes [41], which usually decode 3D scans of the part into more meaningful data formats. This work considered such a process where the geometry of manufactured blades is encoded in a vector of uncertain parameters, ω . Each parameter, shown in Figure 14 for an aerofoil section, controls a particular Degree Of Freedom (DOF) for the geometry. The DOFs applied are: Sweep (axial movement of the section), Lean (circumferential movement of the section), Skew (rotation about the section's centroid) and Leading Edge (LE) and Trailing Edge (TE) *recambering*. The parameters are applied on five aerofoil control sections uniformly distributed through the blade span – at 0%, 25%, 50%, 75%, 100% – providing a total of 25 DOFs. The value of the deformation as a function of the blade span is achieved through

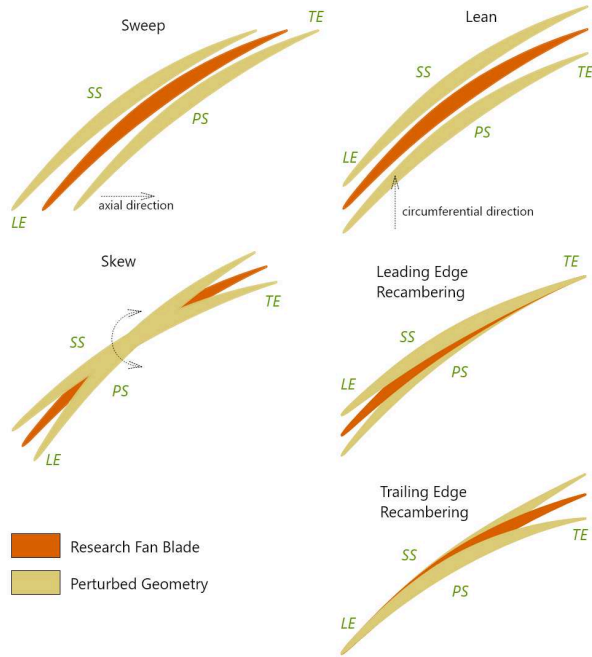


Fig. 14 Parameterization of Manufacturing Variability

smooth cubic B-spline interpolation, with multiple control points via the control sections.

The value of ω was obtained for a large number of manufactured blades (in the order of hundreds), which enables describing the parameters through their Probability Density Function (PDF). Typical distributions are shown for three parameters in Figure 15. It was found that the PDFs closely resemble normal distributions. An inference was therefore made, that all parameters vary with a normal distribution, where the shape factors were estimated from the approximation curves shown in Figure 15. This knowledge was employed to define the quadrature collocation points in the UQ of each design assessed in the optimization process.




Fig. 15 Typical PDFs for uncertain parameters

For the generation of new designs, the design vector, \mathbf{x} , was constructed using the same parametrization. In this case, the upper and lower ranges for each DOF were defined by extending the maximum and minimum values from ω by a factor of three, to allow sufficient flexibility in the design space. Two additional DOFs controlling the locality of the LE and TE *reambering* are also introduced on each span section, such that low values of these parameters cause very localised camber line alterations, and vice-versa. Sufficiently large values can propagate the perturbations

Table 2 Robust performance of Pareto optimal designs relative to datum

Design	LSA-G-NS (1st O.)		LSA-G-NS (2nd O.)	
	$\Delta\mu(\eta)$	$\Delta\sigma(\eta)$	$\Delta\mu(\eta)$	$\Delta\sigma(\eta)$
Opt-1	+0.240%	-44.353%	+0.2511	-45.625
Opt-2	+0.264%	-40.028%	-	-
Opt-3	+0.272%	-37.753%	-	-
Opt-4	+0.365%	-20.821%	+0.353%	-21.140%
DGO	+0.470%	+32.778%	-	-

through the aerofoil, thus providing complete control over the camberline. Thus, the resulting dimensionality of the design vector is 35. Such a parametrization has been previously used for design optimization studies [32].

3. Robust Design Optimization

The quantity of interest for the optimization carried out in this section is the fan bypass isentropic efficiency, defined in Eq. 17. Function and gradient evaluations for this quantity were obtained from the primal and adjoint CFD solution. A first order LSA-G-NS method with no oversampling was employed to obtain the value for the optimization cost functions defined in (16). A deterministic constraint on the Pressure Ratio (PR) was enforced as per Eq. 18 to discourage the optimizer from selecting designs with overly low PR values, which would require the low-pressure shaft to speed up to meet the engine's thrust requirement; or overly large PR values which would require operating at significantly lower shaft speeds and can lead to overloading of the LP turbine blade.

$$\eta(\mathbf{x}, \omega) = \frac{\left(\frac{p_{0_{exit}}}{p_{0_{inlet}}}\right)^{\frac{\gamma-1}{\gamma}} - 1}{\frac{T_{0_{exit}}}{T_{0_{inlet}}} - 1} \quad (17)$$

$$0.99 PR_{datum} \leq PR(\mathbf{x}) \leq 1.05 PR_{datum} \quad (18)$$

An initial data set of 250 different designs was constructed through a pseudo-random design of experiment approach based on Sobol' sequences. CFD simulations were carried out for each design, obtaining the corresponding values for μ , σ and PR . This information was employed to train ANNs through the procedure discussed in Sec. VI.A.1. The application of this methodology also enabled reducing the dimensionality of the parameterization, through the use of active vectors, from 35 parameters to 6 for μ , 9 for σ and 3 for PR . It is worthy to note that this reduction does not affect the number of DOFs allowed for the geometry, as each active vector is obtained from a linear combination of all the original design parameters.

The active vectors can be further exploited to assess the effect of each DOF on the variability of the quantities of interest. The bar graphs in Figure 16 show the components of the most dominant linear combination identified for each function. Each DOF has five bars associated with it, corresponding to the perturbation at the various spanwise locations. The first bar corresponds to the 0% span perturbation, followed by the remaining four bars for perturbations at 25, 50, 75, and 100% span, respectively. This figure shows that LE *recambering*, particularly on the upper 25% of the blade span is the most critical perturbation for both the mean and the standard deviation of the fan efficiency. The low values obtained for the parameters controlling the *locality* of this perturbation suggest that the chordwise extent is not significant and primarily the perturbation on the vicinity of the LE is driving the function. The main effects can be therefore derived to be the effective flow incidence angle and the smoothness of the expansion upstream of the shock wave, both being mainly controlled by the camber line distribution at the LE. In addition, it is noticeable that the circumferential lean at the tip section is next in significance for the mean, while it is not so relevant for the standard deviation. Tip lean primarily affects leakage flow, and this suggests that while this phenomenon is significant for the average performance of the blade, it is not sensitive to the manufacturing variability considered in this study. Thus, a close control of tip leakage flow can be a substantial design handle for robust fan blades. With respect to the PR, it was found that skew and TE *recambering* are the most critical perturbations, which is in agreement with aerodynamic principles, as they mostly control flow turning.

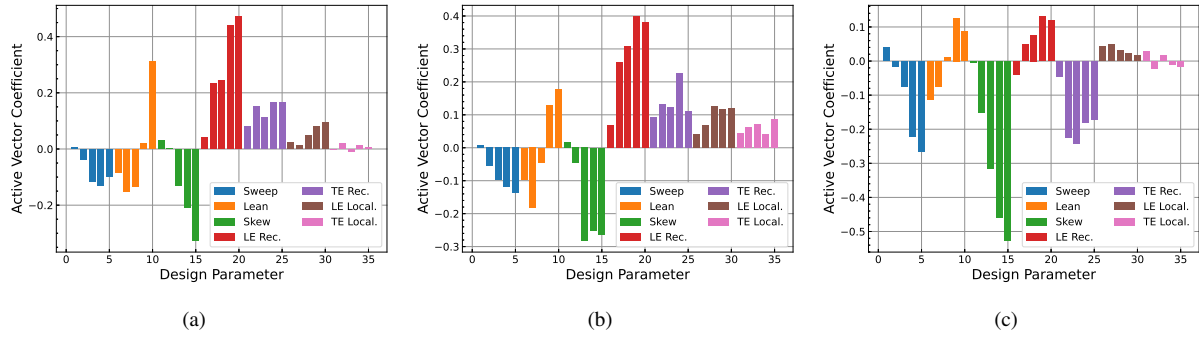


Fig. 16 Coefficients of the most dominant direction for (a): mean; (b): standard deviation and (c): pressure ratio

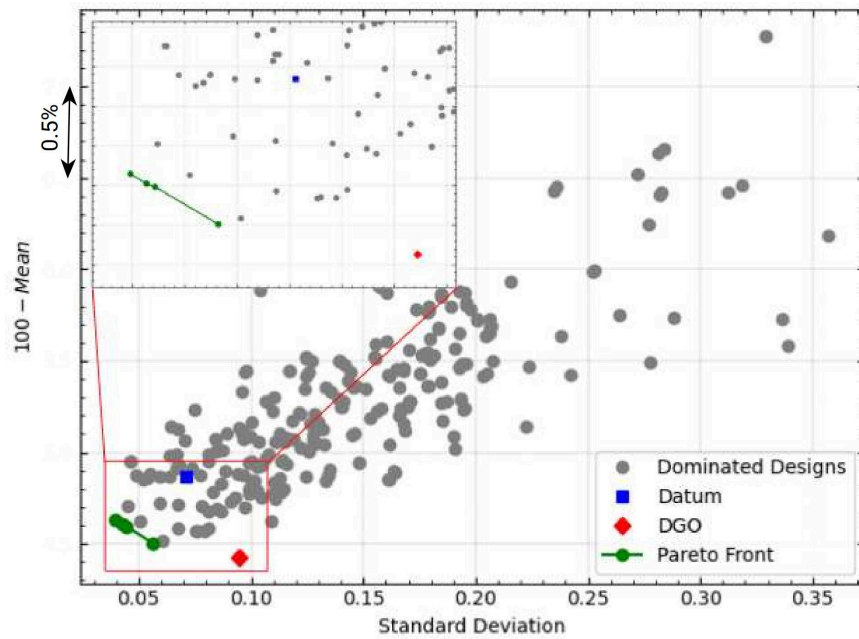


Fig. 17 Pareto frontier for the robust designs

An adaptive sampling strategy based on optimizing the system was employed to augment the number of data points and obtain more accurate predictions. As such, for every iteration, the Pareto frontier was obtained with ARMOGA, the corresponding designs evaluated with CFD and LSA-G-NS to obtain the true values of the cost functions, and the data were subsequently appended to the data set. Four iterations were required to converge the Pareto front. The results are shown in Figure 17, which also includes the values for the Datum design and the dominated designs evaluated from the DOE and the adaptive sampling iterations. It can be observed from this figure that there is not a significant trade-off between mean and variance of efficiency for the parameterization considered, and optimizing one of the functions generally improves the other. This is also in agreement with the fact that the active vectors for these quantities are quite similar, as shown in Figure 16, which implies that the search directions are aligned.

The ARMOGA process identified four designs in the Pareto frontier, namely Opt-1 through Opt-4, whose performance is gathered in Table 2. Relative to the datum design, the optimized geometries achieve improvements in mean efficiency ranging from an additional 0.24% to 0.365%, while the standard deviation has been reduced in the range of 20.82% through 44.35%. These values were obtained from first order LSA-G-NS approximations, which assume that the behavior of the cost functions in the uncertainty range is approximately linear. For the extreme designs in the Pareto front, Opt-1 and Opt-4, as well as the datum design, the value of μ and σ was also calculated with a second order LSA-G-NS employing 30% of oversampling, with the aim of producing a more accurate estimation of

their true robust performance. It was found that the trend predicted with the first order method is accurate and the performance improvement identified in the optimization process is maintained. These results demonstrate that an adequate desensitising of the fan blade to manufacturing variability has been achieved.

To provide context to these metrics, the robust performance of the Pareto optimal designs is compared against a Deterministic Global Optimum (DGO), obtained from a previous work conducted with the same parameterization and PR constraints [32]. A first order LSA-G-NS process was employed to derive the robust performance of DGO, which resulted in an increase of 0.47% in the mean efficiency but a drop of about 33% in the standard deviation. The DGO design is also Pareto-optimal, as it provides the best mean performance, however its worsened sensitivity with respect to the datum blade disfavours its potential application. Therefore, while deterministic optimization processes can lead to designs that are, on average, high-performing, large variability in the measured performance between physical manufactured blades happen. This is a particularly discouraged scenario on applications where multiple components are expected to work jointly and balanced, like turbomachinery applications.

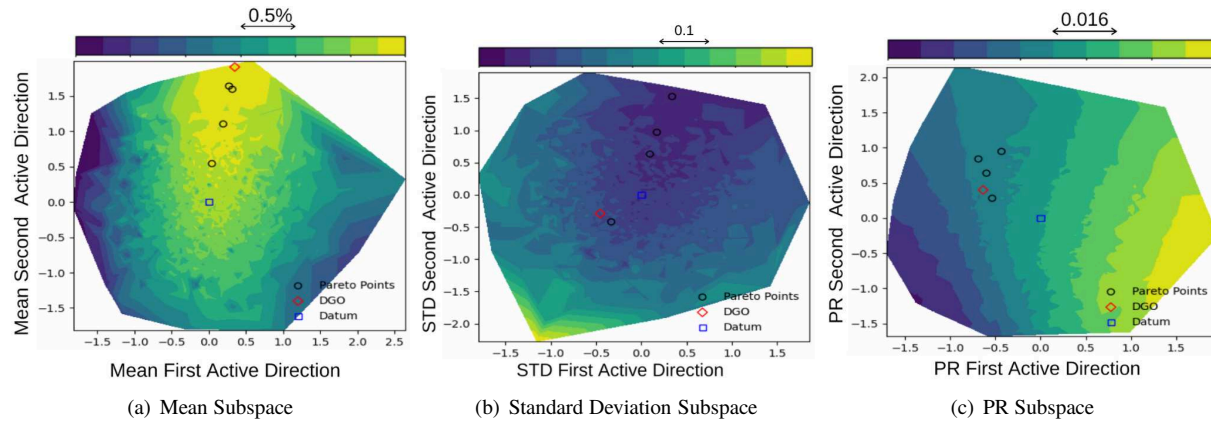


Fig. 18 2D Performance Maps for the Quantities of Interest

For visualization purposes, the performance of the optimum points can be visualised in 2-dimensional active subspace performance maps, following [42]. These maps, shown in Figure 18, demonstrate that all the designs lie in a corridor of high mean efficiency, which the optimization process has been able to identify, and that could be further exploited to device design principles for statistically high-performing blades. The standard deviation map, on the other hand, shows a less clear trend, with the implication that DGO and Opt-4 are located in a region where σ changes rapidly, making the robustness of these designs potentially highly sensitive to design alterations, as well as highly dependent on the extent of the design uncertainties.

4. Aerodynamic Analysis

To understand the leading factors contributing to the increased robustness of the optimal designs found, an aerodynamic analysis is performed for Opt-1, since it offers the lowest sensitivity to manufacturing variations, and DGO which offers the highest mean performance. The interested reader can consult [32] for a more in-depth analysis of the DGO shape.

Figure 19 shows the geometrical features of Opt-1. It can be noted that an aft sweep has been applied to the blade tip, which is typically a disfavoured design practise. This has to do with the fact that, due to centrifugal forces, the low energy flow from the boundary layer at lower span sectors is forced radially outwards and, if there is significant backwards sweep, this flow can interact with an extended portion of the blade's chord, producing a larger accumulation of delayed flow. However, in this case, the backwards sweep is enabled due to circumferential lean, which has inclined the blade pressure surface further towards the hub at the tip region, causing the blade to exert a body force on the fluid, directing it towards the hub and minimizing its outward tendency.

These design features bias the airflow to the mid-span region, causing a slight reduction in this sector's performance, as can be noted from the radial efficiency distribution in Figure 20a. Nonetheless, this performance deficit is recovered at the tip, where the sweep produces a reduction in the effective flow velocity, which strongly mitigates the shock wave and delays it further downstream, as shown in Figure 20b, thus minimizing the shock-induced separated region. This

alleviation of shock wave effects is believed to be the primary mechanism for achieving the increased robustness of the blade, as generally even slight geometrical deviations can cause the location and strength of the shock wave to vary, which translates to performance variations. By having an ameliorated shock behavior, the blade is able to cope better with such variability in its shape. With regards to the increased mean performance, the more uniform loading and better behaved airflow at the tip enables increasing the blade twist, resulting in an increased flow turning, which can be noted from the exit axial whirl angle distribution from Figure 20.




Fig. 19 Geometrical Features of Opt-1 (orange) relative to Datum (grey)

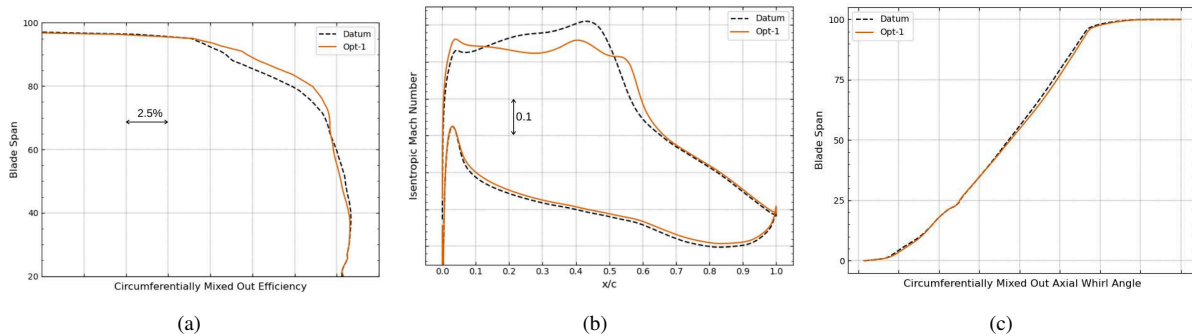


Fig. 20 Aerodynamic behavior of Opt-1 relative to Datum. (a): radial efficiency distribution; (b): lift plot at 80% span; (c): exit axial whirl angle distribution

The geometry of DGO is presented in Fig. 21. It can be appreciated that this design follows the datum shape more closely throughout the span, with the most significant design feature being a more skewed tip. This causes the tip aerofoils to be better aligned with the flow direction, thus reducing the effective incidence angle. The aerodynamic effect is a reduction in the pre-shock mach number, which weakens the shock, as can be observed from the lift plot in Fig. 22b. In turn, this reduces the shock-induced separated region of the blade tip, improving the flow attachment, and enabling the blade to apply more turning to the airflow, thus increasing the performance at around 90% span, as noticeable from the radial efficiency distribution shown in Fig. 22a.

It is worthy to highlight that the aerodynamic benefit for the DGO has been obtained from performing slight adjustments to the datum shape, thus making it more suitable to the conditions for which it was designed. However, manufacturing variability can significantly alter the operating conditions, which can render the DGO improvements ineffective and indeed counter productive, thus making this design not robust.

The effects of the optimized geometry are of a highly three-dimensional nature, and their identification has been

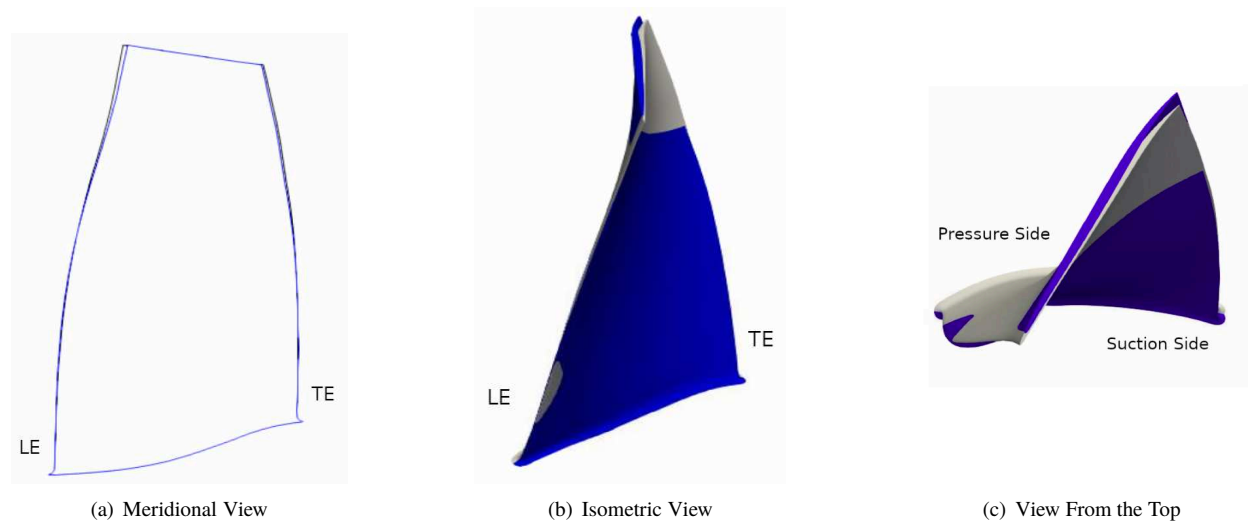


Fig. 21 Geometrical Features of DGO (blue) relative to Datum (grey)

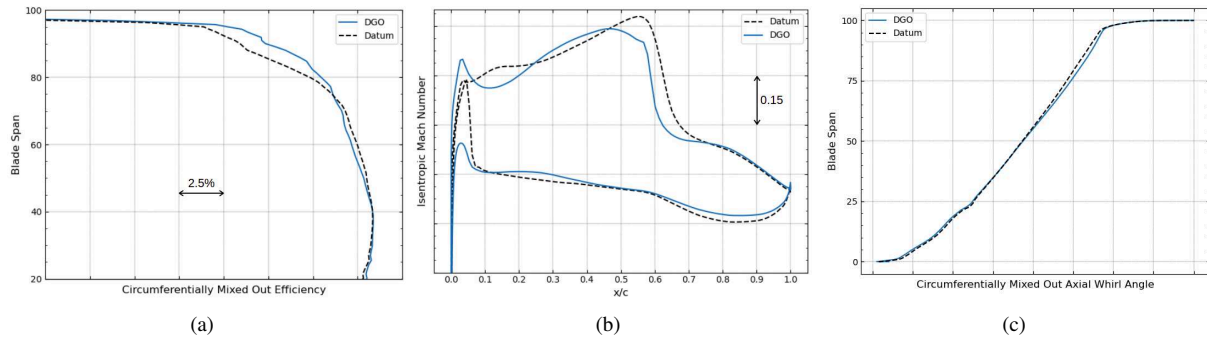


Fig. 22 Aerodynamic behavior of DGO relative to Datum. (a): radial efficiency distribution; (b): lift plot at 90% span; (c): exit axial whirl angle distribution

possible due to the application of a 3D CFD setup employed in the optimization process. In turn, such a high-fidelity model has been enabled through the reduced cost of performing UQ via the LSA-G-NS method discussed in this work.

VII. Conclusion

This paper presents a novel approach to estimate the coefficients of a Polynomial expansion of an uncertain quantity of interest. The approach is based on the solution of a linear equality approach, which provides improved results compared to the traditional stacked approach in the presence of noisy gradient information, as it is often the case when this is retrieved from running adjoint solvers.

The advantages of the proposed approach are first demonstrated on simpler test cases, including an analytical problem with the artificial addition of white noise, and a 5-variable uncertainty quantification for the well known Rotor 37 test case. Then, the method is tested on a higher-dimensionality robust optimization problem, where it is combined with dimension-reduction techniques and response surfaces to find the trade off between mean and standard deviation for the isentropic efficiency of a modern aero-engine representative fan called Vital.

VIII. Acknowledgement

The authors would like to thank Rolls-Royce for their support and permission to publish this work. The second author is funded by a Scholarship from the Italian Ministry of Education, through the University of Cagliari. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement N. 769025.

References

- [1] Chung, H.-S., and Alonso, J. J., "Using gradients to construct cokriging approximation models for high-dimensional design optimization problems," *AIAA paper*, Vol. 317, 2002, pp. 14–17.
- [2] Seshadri, P., Constantine, P., Iaccarino, G., and Parks, G., "A density-matching approach for optimization under uncertainty," *Computer Methods in Applied Mechanics and Engineering*, Vol. 305, 2016, pp. 562–578.
- [3] Lockwood, B., and Mavriplis, D., "Gradient-based methods for uncertainty quantification in hypersonic flows," *Computers & Fluids*, Vol. 85, 2013, pp. 27–38.
- [4] Simpson, T. W., Poplinski, J., Koch, P. N., and Allen, J. K., "Metamodels for computer-based engineering design: survey and recommendations," *Engineering with computers*, Vol. 17, No. 2, 2001, pp. 129–150.
- [5] Peng, J., Hampton, J., and Doostan, A., "A weighted l_1 -minimization approach for sparse polynomial chaos expansions," *Journal of Computational Physics*, Vol. 267, 2014, pp. 92–111.
- [6] Burkardt, J., "1D quadrature rules for sparse grids," *Interdisciplinary Center for Applied Mathematics and Information Technology Department, Virginia Tech*, 2010.
- [7] Burkardt, J., "The combining coefficient for anisotropic sparse grids," *Interdisciplinary Center for Applied Mathematics & Information Technology Department, Virginia Tech*, 2010.
- [8] Golub, G. H., and Van Loan, C. F., *Matrix computations*, Vol. 3, JHU Press, 2012.
- [9] Roderick, O., Anitescu, M., and Fischer, P., "Polynomial regression approaches using derivative information for uncertainty quantification," *Nuclear Science and Engineering*, Vol. 164, No. 2, 2010, pp. 122–139. <https://doi.org/10.13182/NSE08-79>.
- [10] Li, Y., Anitescu, M., Roderick, O., and Hickernell, F., "Orthogonal bases for polynomial regression with derivative information in uncertainty quantification," *International Journal for Uncertainty Quantification*, Vol. 1, No. 4, 2011. <https://doi.org/10.1615/int.j.uncertaintyquantification.2011002790>.
- [11] Temljakov, V. N., "APPROXIMATION OF PERIODIC FUNCTIONS OF SEVERAL VARIABLES WITH BOUNDED MIXED DIFFERENCE," *Mathematics of the USSR-Sbornik*, Vol. 41, No. 1, 1982, p. 53. URL <http://stacks.iop.org/0025-5734/41/i=1/a=A03>.
- [12] Bungartz, H. J., and Griebel, M., "Sparse grids," *Acta numerica*, Vol. 13, 2004, pp. 147–269.
- [13] Blatman, G., and Sudret, B., "Adaptive sparse polynomial chaos expansion based on least angle regression," *Journal of Computational Physics*, Vol. 230, No. 6, 2011, pp. 2345–2367.
- [14] Gautschi, W., "On Generating Orthogonal Polynomials," *SIAM Journal on Scientific and Statistical Computing*, Vol. 3, No. 3, 1982, pp. 289–317. <https://doi.org/10.1137/0903018>.

- [15] Ghisu, T., and Shahpar, S., "Toward affordable uncertainty quantification for industrial problems - Part I: Theory and validation," *Proceedings of the ASME Turbo Expo*, Vol. 2C-2017, 2017. <https://doi.org/10.1115/GT2017-64842>.
- [16] Le Maître, O. P., Reagan, M. T., Najm, H. N., Ghanem, R. G., and Knio, O. M., "A stochastic projection method for fluid flow: II. Random process," *Journal of computational Physics*, Vol. 181, No. 1, 2002, pp. 9–44. <https://doi.org/10.1006/jcph.2002.7104>.
- [17] Migliorati, G., Nobile, F., von Schwerin, E., and Tempone, R., "Approximation of Quantities of Interest in Stochastic PDEs by the Random Discrete L2 Projection on Polynomial Spaces," *SIAM Journal on Scientific Computing*, Vol. 35, No. 3, 2013, pp. A1440–A1460.
- [18] Cohen, A., Davenport, M. A., and Leviatan, D., "On the stability and accuracy of least squares approximations," *Foundations of computational mathematics*, Vol. 13, No. 5, 2013, pp. 819–834.
- [19] Seshadri, P., Narayan, A., and Mahadevan, S., "Effectively Subsampled Quadratures For Least Squares Polynomial Approximations," *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 5, 2016. <https://doi.org/10.1137/16M1057668>.
- [20] Ghisu, T., and Shahpar, S., "Affordable uncertainty quantification for industrial problems: Application to aero-engine fans," *Journal of Turbomachinery*, Vol. 140, No. 6, 2018. <https://doi.org/10.1115/1.4038982>.
- [21] Lawson, C. L., and Hanson, R. J., *Solving least squares problems*, Vol. 15, SIAM, 1995.
- [22] Björck, Å., *Numerical methods in matrix computations*, Springer, 2015.
- [23] Eldén, L., "A weighted pseudoinverse, generalized singular values, and constrained least squares problems," *BIT Numerical Mathematics*, Vol. 22, No. 4, 1982, pp. 487–502.
- [24] Moore, R. D., and Reid, L., "Performance of Single-Stage Axial-Flow Transonic Compressor With Rotor and Stator Aspect Ratios of 1.19 and 1.26, Respectively, and With Design Pressure Ratio of 1.82," *NASA Technical Paper 1338*, 1978.
- [25] Suder, K. L., "Experimental Investigation of the Flow Field in a Transonic, an Axial Flow Compressor With Respect to the Development of Blockage and Loss," *NASA TM 107310*, 1996.
- [26] John, A., Shahpar, S., and Qin, N., "Alleviation of Shock-Wave Effects on a Highly Loaded Axial Compressor Through Novel Blade Shaping," *Proceedings of ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition*, 2016, p. V02AT37A040. <https://doi.org/10.1115/GT2016-57550>.
- [27] Mura, R., Ghisu, T., and Shahpar, S., "Least squares approximation-based polynomial chaos expansion for uncertainty quantification and robust optimization in aeronautics," *AIAA AVIATION 2020 FORUM*, Vol. 1 PartF, 2020, pp. 1–34. <https://doi.org/10.2514/6.2020-3163>.
- [28] Ghisu, T., Parks, G. T., Jarrett, J. P., and Clarkson, P. J., "Robust Design Optimisation of Gas Turbine Compression Systems," *Journal of Propulsion and Power*, Vol. 27, No. 2, 2011, p. 282–295. <https://doi.org/10.2514/6.2008-5828>.
- [29] Seshadri, P., Shahpar, S., and Parks, G., "Robust Compressor Blades for Desensitizing Operational Tip Clearance Variations," *Proceedings of the ASME Turbo Expo*, Vol. 2, 2014. <https://doi.org/10.1115/GT2014-26624>.
- [30] Wang, X., Hirsch, C., Liu, Z., Kang, S., and Lacor, C., "Uncertainty-based robust aerodynamic optimization of rotor blades," *International Journal for Numerical Methods in Engineering*, Vol. 94, 2013, pp. 111–127. <https://doi.org/10.1002/nme.4438>.
- [31] Shahpar, S., "SOFT: A New Design And Optimisation Tool for Turbomachinery," *Evolutionary Methods for Design Optimisation And Control*, K. Giannakoglou et al., ed., 2002.
- [32] Lopez, D. I., Ghisu, T., and Shahpar, S., "Global Optimisation of a Transonic Fan Blade Through AI-Enabled Active Subspaces," *Proceedings of the ASME Turbo Expo*, (GT2021-59166) 2021.
- [33] Haykin, S., *Neural Networks: A Comprehensive Foundation*, Prentice-Hall, 1999.
- [34] Cybenko, G., "Approximation by superpositions of a sigmoidal function," *Mathematics of Control, Signals and Systems*, Vol. 2, 1989, pp. 303–314.
- [35] Constantine, P. G., *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*, SIAM Spotlights, 2015.
- [36] Zamboni, G., and Xu, L., "Fan Root Aerodynamics for Large Bypass Gas Turbine Engines: Influence on the Engine Performance and 3D Design," *Journal of Turbomachinery*, Vol. 134, No. 6, 2009, p. 061017. <https://doi.org/10.1115/GT2009-59498>.

- [37] Lapworth, L., "HYDRA-CFD: A Framework for Collaborative CFD Development," *International Conference on Scientific and Engineering Computation (IC-SEC)*, 2004.
- [38] Milli, A., and Shahpar, S., "PADRAM: Parametric Design and Rapid Meshing System for Complex Turbomachinery Configurations," *Proceedings of the ASME Turbo Expo*, Vol. 8, 2012. <https://doi.org/10.1115/GT2012-69030>.
- [39] Giugno, A., Shahpar, S., and Traverso, A., "Adjoint-Based Optimization of a Modern Jet Engine Fan Blade," *Proceedings of the ASME Turbo Expo*, 2020.
- [40] John, A., Ning, Q., and Shahpar, S., "The Influence of Parametrisation Setup on the Constrained Adjoint Optimisation of Transonic Fan Blade," *Proceedings of the ASME Turbo Expo*, 2020.
- [41] Shahpar, S., "Building digital twins to simulate manufacturing variation," *Proceedings of the ASME Turbo Expo*, Vol. 2A-2020, 2020. <https://doi.org/10.1115/GT2020-15263>.
- [42] Sheshadri, P., Shahpar, S., Constantine, P., Parks, G., and Adams, M., "Turbomachinery Active Subspace Performance Maps," *Journal of Turbomachinery*, Vol. 140, No. 041003-1, 2017.

2021-07-28

Gradient-enhanced least-square polynomial chaos expansions for uncertainty quantification and robust optimization

Ghisu, Tiziano

AIAA

Ghisu T, Lopez D, Seshadri P, Shahpar S. (2021) Gradient-enhanced least-square polynomial chaos expansions for uncertainty quantification and robust optimization. In: AIAA Aviation 2021 Forum, 2-6 August 2021, Virtual Event. Paper number AIAA 2021-3073

<https://doi.org/10.2514/6.2021-3073>

Downloaded from Cranfield Library Services E-Repository