



UNICA

UNIVERSITÀ
DEGLI STUDI
DI CAGLIARI

Ph.D. DEGREE IN
Industrial Engineering
Cycle XXXV

TITLE OF THE Ph.D. THESIS

Design Optimisation and Flow Characterisation for Future Aeroengine Axial Fan
Blades

Scientific Disciplinary Sector(s)
ING-IND/08

Ph.D. Student: Diego Ignacio Lopez
Supervisor: Dr Tiziano Ghisu

Final exam. Academic Year 2021/2022
Thesis defence: April 2023 Session

Questa tesi può essere utilizzata, nei limiti stabiliti dalla normativa vigente sul Diritto d'Autore (Legge 22 aprile 1941 n. 633 e succ. modificazioni e articoli da 2575 a 2583 del Codice civile) ed esclusivamente per scopi didattici e di ricerca; è vietato qualsiasi utilizzo per fini commerciali. In ogni caso tutti gli utilizzi devono riportare la corretta citazione delle fonti. La traduzione, l'adattamento totale e parziale, sono riservati per tutti i Paesi. I documenti depositati sono sottoposti alla legislazione italiana in vigore nel rispetto del Diritto di Autore, da qualunque luogo essi siano fruiti.

This Thesis can be used, within the limits established by current legislation on Copyright (Law 22 April 1941 n. 633 and subsequent amendments and articles from 2575 to 2583 of the Civil Code) and exclusively for educational and research purposes; any use for commercial purposes is prohibited. In any case, all the uses must report the correct citation of the sources. The translation, total and partial adaptation, are reserved for all countries. The documents filed are subject to the Italian legislation in force in compliance with copyright, from wherever are used.

Abstract

The increased need to design higher performing aerodynamic shapes has led to design optimisation cycles requiring high-fidelity CFD models and high-dimensional parametrisation schemes. The computational cost of employing global search algorithms on such scenarios has typically been prohibitive for most academic and industrial environments. In this work, a novel methodology is presented, called AInADS. This strategy leverages the capabilities of Artificial Neural Networks for regressing complex unstructured data, while coupling them with dimensionality reduction algorithms. This approach enables employing global-based optimisation methods on high-dimensional applications through a reduced computational cost. The capabilities of AInADS are demonstrated on three turbomachinery applications of significant industrial relevance.

The first case performs the efficiency optimisation of a modern jet engine fan blade with constrained pressure ratio. This outcome is compared against a state-of-the-art adjoint-based approach. Results indicate the strategy proposed achieves comparable improvements to its adjoint counterpart with a reduced computational cost and run time, and can scale better to multi-objective optimisation applications.

The second application concerns the stability range of axial fan blades. Historically, the tip clearance size has been considered to be the main factor driving its behaviour. This work reveals that the stall characteristics are defined by the axial momentum flux of the tip leakage flow and that tip clearance is primarily a strong driver for this metric. AInADS is employed for carefully tailoring the axial momentum via three-dimensional design, which enables a higher degree of control over the stability range for cases where the tip clearance responds to other considerations and cannot be defined for this purpose. The effect of the axial momentum on efficiency is also addressed and the trade-off between operability range and design point performance derived. The results show that the conditions for optimal stability differ from those for optimal efficiency and that control over the axial momentum enables tuning the design for a desired exchange. Numerical simulations have been

employed to drive the analysis through a high-fidelity computational model whose behavior is supported by rich set of experimental data. Contrary to current belief, results further indicate that an accurate characterisation of stall, including onset mechanism, can be achieved through steady-state simulations, minimising the need for expensive time-accurate computations during the design phase.

The final application introduces uncertainties in the design process of axial fan blades. As they are manufactured, blades deviate from the design intent shape and such geometrical variability translates to performance drifts. This work makes use of AInADS, coupled with Uncertainty Quantification methods, to estimate the statistical behaviour of a fan blade when subjected to typical manufacturing deviations. This knowledge is employed to optimise the shape of the fan blade and improve its robustness to uncertain shape deviations.

Preface

The author would like to thank Rolls-Royce for their support and permission to publish this work. The research conducted for this thesis forms part of some international collaborations. In particular, the work detailed in Chapter 3 and Chapter 5 received funding from the European Union’s Horizon 2020 research and innovation program under Grant Agreement No. 769025. The research conducted for Chapter 4 received funding from the UK Aerospace Technology Institute and InnovateUK under Grant Agreement No. 113286.

The author has been funded by a Scholarship from the Italian Ministry of Education, through the University of Cagliari for the duration of their doctoral studies. This PhD thesis would not exist were it not for this support.

Much of the work collected in this text has been previously published and/or presented at scientific conferences by the author. Following is a list of such publications.

- [1] Diego I. Lopez, Tiziano Ghisu, and Shahrokh Shahpar. “Global Optimization of a Transonic Fan Blade Through AI-Enabled Active Subspaces”. *Journal of Turbomachinery*, 144(1), 09 2021. 011013.
- [2] Diego I. Lopez, Tiziano Ghisu, Timoleon Kipouros, Shahrokh Shahpar, and Mark Wilson. “Extending Highly Loaded Axial Fan Operability Range Through Novel Blade Design”. *Journal of Turbomachinery*, 144(12), 09 2022. 121009.
- [3] Tiziano Ghisu, Diego I. Lopez, Pranay Seshadri and Shahrokh Shahpar. “Gradient-enhanced Least-square Polynomial Chaos Expansions for Uncertainty Quantification and Robust Optimization”. AIAA 2021-3073. *AIAA AVIATION 2021 FORUM*. August 2021.
- [4] Diego I. Lopez, Tiziano Ghisu, Timoleon Kipouros and Shahrokh Shahpar. “An AI-Assisted Method for Simulation-based Design Analysis and Optimisation in Turbomachinery.” *NRC22 UK*. June 2022.

Acknowledgements

I could not have undertaken this journey without the support and guidance of two incredible researchers. My deepest gratitude to Tiziano Ghisu and Shahrokh Shahpar, you are the best supervisors one can wish for.

I am greatly indebted to Timoleon Kipouros. Thank you for welcoming me in Cranfield and your invaluable advice. I can only regret not having more time to work together.

Special thanks to Mark Wilson and Luis Llano, from Rolls-Royce for your patience and for being so generous with your knowledge. To Indi Tristante, who was always available for technical discussions, and to Matt Harris, for the all the help through the years. I had the pleasure of working with Domenico Sgrò. Thank you for being an amazing source of new ideas. I'm really proud of the outcome of our work together.

I'm very grateful to have shared this journey with my friends from Cagliari, Mario and Irene, and all the friends from Spinelli's office at Cranfield. Thank you for making this a great experience.

Thank you to my parents, Adriana and Jorge, and to my better half, Pilar, who were with me every step of the way. I could not have done it without you.

Contents

1	Introduction	18
1.1	Background	18
1.2	Motivation	22
2	AI Enabled Active Subspaces	24
2.1	Engineering Design Optimisation	24
2.2	Fitting Artificial Neural Networks	26
2.3	Coupling Active Design Subspaces	30
2.3.1	Active Subspaces	30
2.3.2	Enabling ADS through AI	32
2.4	Rewriting the Optimisation Problem	33
3	Benchmarking AInADS performance	36
3.1	Comparison to an adjoint gradient-based approach	36
3.2	Application to a jet engine fan blade	38
3.2.1	Computational Tools	38
3.2.2	Optimisation Problem	39
3.2.3	Global Optimisation Approach	41
3.2.4	Adjoint-based Approach	45
3.2.5	Remarks on Computational Efficiency	48
3.3	Conclusion	50
4	Understanding Fan Stall Range through AInAS	52
4.1	Introduction	53
4.2	Computational Model	57
4.2.1	Blade Geometry	57
4.2.2	CFD Domain	58
4.2.3	CFD Solver	61
4.2.4	Identification of stall point	62
4.2.5	Model Validation	63
4.3	Design Space Exploration	66

4.3.1	AI-Enabled Active Subspaces	66
4.4	Exploratory Design Analysis	68
4.5	Effect of Tip Leakage Flow on Stability Range	76
4.5.1	Results at Top of Climb	82
4.6	Conclusions	84
5	Robust Design Optimisation	86
5.1	Introduction	86
5.2	Uncertainty Quantification	88
5.3	Optimisation Framework	89
5.4	Computational Tools	91
5.4.1	Geometry Parametrisation	91
5.5	Robust Design Optimisation	94
5.6	Aerodynamic Analysis	99
5.7	Conclusions	103

List of Figures

1.1	Historic evolution of aviation efficiency and global traffic [1]	19
2.1	General structure of perceptron neuron for three inputs	26
2.2	Activation functions for perceptrons	27
2.3	Schematic of the feed-forward neural network employed	28
2.4	Iterative process proposed to learn the ADS and fit an ANN to CFD samples	33
3.1	Primal convergence history for mean-flow residual and non-dimensionalised cost function and constraint values	37
3.2	Research blade CFD domain	38
3.3	Geometry parametrisation through EDP. The magnitude of the perturbations has been enlarged for clarity	40
3.4	Eigenvector convergence	42
3.5	Eigenvalue decay and cumulative energy plots for efficiency and pressure ratio gradient covariance matrices	43
3.6	2D active subspace maps for the optimisation functions	44
3.7	Radial efficiency profile for datum and optimised designs in (a); lift plot at 90% span in (b)	45
3.8	Aerofoil geometry at 90% span. (a): camberline distribution; (b): curvature vs arc length. The datum curve has been shifted for clarity	46
3.9	Static pressure contours at 90% span. (a): datum blade; (b): optimised design	46
3.10	Adjoint-based optimisation history. The optimum found is highlighted in green	47
3.11	Evolution of efficiency gradient through optimisation. The gradient has been normalised by l_2 norm.	48
3.12	Camberline distribution and lift plots of datum and optimum designs at (a), (b): 90% span; and (c), (d): 50% span	49

3.13	Blade metal angle distribution. 3.13a: Inlet angle, 3.13b: Exit angle	50
3.14	Fan characteristic curves for optimised designs	51
4.1	Useful markers for analysing flow mechanisms present at stall.	55
4.2	Illustration leakage jet impingement at the adjacent blade's pressure surface.	55
4.3	Running up process of cold static blade shape.	58
4.4	Convergence of the state to state process at a representative speed.	59
4.5	Illustration of the CFD domain for the research blade. Image distorted and not to scale.	60
4.6	Rotor mesh details	60
4.7	Radial distribution of nodes at mesh zone interfaces	61
4.8	Schematic of the function implemented for stall identification via Fibonacci search method.	62
4.9	Validation of performance characteristics for model against experimental data.	64
4.10	Radial profiles at the BOGV LE for 95% speed at the ADP and near stall conditions.	65
4.11	Degrees of freedom employed to manipulate the geometry of the fan blade [2].	67
4.12	Sketch of the sub-region approach followed for exploration of the Efficiency Design Space.	70
4.13	Cumulative energy plots for the covariance matrix of Efficiency, PR and Tip Leakage Axial Momentum Flux. Only the first 10 eigenmodes are shown.	71
4.14	Performance maps showing the response of the QOIs in the design space.	72
4.15	Coefficient activation for the first active direction of the QOIs. Higher magnitude coefficients highlight parameters with increased effect.	73
4.16	Stall margin and ADP efficiency versus tip leakage axial momentum flux, normalised by datum value. The QOIs present a quadratic response to changes in axial momentum.	78
4.17	Contours of radial vorticity at stall for the three representative designs at 95% span. The enhancement of the leakage jet initially suppresses the corner separation leading to increased stall range. Further increases in its strength cause impingement on the adjacent blade's pressure surface and early spillage.	79
4.18	Isentropic Mach number distributions at 95% span.	80

4.19	Top view of the datum and efficiency-optimised designs.	81
4.20	Secondary design drivers for stall margin. For a given axial momentum, off-loading the tip and applying lean towards the pressure side produce increased stability.	82
4.21	Variation of fan performance versus axial momentum flux as the speed is increased.	83
5.1	Illustration of probability distribution for the response of the quantities of interest in the event space Ω	88
5.2	Parametrisation of Manufacturing Variability	92
5.3	Typical PDFs for uncertain parameters	93
5.4	Coefficients of the most dominant direction for (a): mean of the efficiency; (b): standard deviation and (c): pressure ratio .	96
5.5	Pareto frontier for the robust designs	97
5.6	2D Performance Maps for the Quantities of Interest	98
5.7	Geometrical Features of Opt-1 (orange) relative to Datum (grey)	100
5.8	Aerodynamic behaviour of Opt-1 relative to Datum. (a): radial efficiency distribution; (b): lift plot at 80% span; (c): exit axial whirl angle distribution	101
5.9	Geometrical Features of DGO (blue) relative to Datum (grey)	102
5.10	Aerodynamic behaviour of DGO relative to Datum. (a): radial efficiency distribution; (b): lift plot at 90% span; (c): exit axial whirl angle distribution	103

List of Tables

2.1	Hyper-parameters tuned in grid search	29
3.1	Network prediction accuracy trained with high-dimensional inputs and active vectors	44
3.2	Summary of computational efficiency for the optimisation approaches	50
4.1	Performance metrics for the response surfaces constructed through AI-Enabled ADS.	74
5.1	Robust performance of Pareto optimal designs relative to datum	95

Chapter 1

Introduction

1.1 Background

The civil aviation industry has experienced an accelerated growth in the last 50 years, with the number of Available Seat Kilometers (ASK) increasing over 7% in 2018 alone [1, 3, 4], see Fig. 1.1. With a more acute environmental awareness, governments have increased funding and support to develop greener products in aviation, as well as introduced legislation to penalise the commercialisation of polluting solutions. As a result, modern day aircraft engines are required to meet increasingly tighter regulations in terms of their performance and by-product emissions, which increases the need for every component in the system to be carefully designed so as to achieve its operational requirements while employing the least amount of resources.

For decades, aviation engineers worldwide have striven to increase the efficiency of their designs. This goal has been largely achieved, with modern turbofan engines reaching improvements in efficiency of over 25% with respect to the performance of engines at the turn of the century [5]. However, as it can be observed on Fig. 1.1, the rate of improvement has slowed down, and the trend suggests that a limit in the maximum improvement is nigh.

Indeed, in recent years, the need to extract the last ounce of performance from every design has led to an increased need for higher fidelity models, which are able to represent more faithfully the complex physical characteristics of the system, thus making the design process more likely to lead to novel, high-performing shapes. From an aerodynamic standpoint, design optimisation based on Computational Fluid Dynamics (CFD) is one of the primary tools engineers have used to improve their aerodynamic designs [6–8]. In these workflows, the system under study is *parametrised* with a set of features that control its behaviour. In some cases, these are defined by the

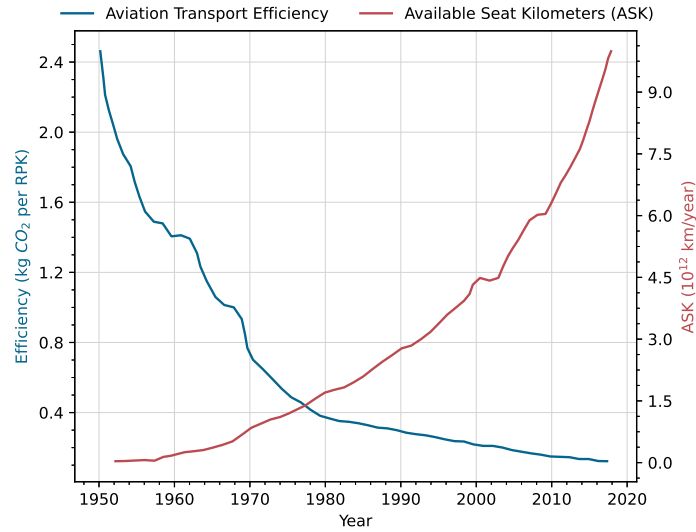


Figure 1.1: Historic evolution of aviation efficiency and global traffic [1]

designer based on prior knowledge. Most commonly, however, the available information about the system is not sufficient to expertly express the features driving its behaviour. As a result, generic and flexible parametrisations are sought. In fact, capitalising on the increased model fidelity usually requires a very fine parametrisation, enabling the system sufficient degrees of freedom to respond to even the smallest features, which increases the complexity of the process. Moreover, in such cases, in addition to *improving* the system, a desired outcome of optimisation processes is *characterising* it and identifying, from this fine set of features, the leading drivers of its behaviour. Both of these goals require the definition of objective functions, metrics on which to base the performance of the system; and constraint functions, which set the bounds of the problem and limit the scope of possibilities. Typical objective functions in aviation design include efficiency, drag or weight, while constraints might be lift, thrust or capacity.

While undoubtedly powerful, high dimensional parametrisation schemes pose many challenges to traditional optimisation workflows, which are unprepared to handle the large amount of design features. Perhaps the most strenuous of them all is an effect commonly referred to as *the curse of dimensionality* [9], which will be a recurring topic in this manuscript as methods are developed to overcome it. This provoking name describes a phenomena whereby, as the number of design features increases, the volume of the hyperspace spanned by them, the design space, expands exponentially. In these cases, adequate sampling of the design space requires an exponentially

growing number of data points, otherwise samples become sparse and pose risks of misleading further analyses [10]. As a result, and especially when considering high-fidelity CFD models, which tend to be computationally expensive, increasing the degrees of freedom of the system can quickly make the design problem intractable.

Such is the magnitude of the dimensionality challenge that many research efforts have been devoted to it [11–14]. In CFD applications, gradient-based computational optimisation methods have gathered much attention, due to the increased availability of *adjoint* flow solvers. The adjoint method [15–18], allows computation of the gradient of an objective function with respect to design parameters while maintaining the cost of the operation almost independent of the number of parameters employed. First order information provides a wealth of knowledge about the behaviour of the underlying function in the vicinity of the sampled point. Therefore, adjoint solvers are able to provide deep insights about the system at an approximately constant cost. First order search algorithms employ this information to travel in the direction of steepest descent and efficiently optimise the objective functions [19–21].

This approach has been successfully applied to optimise turbomachinery components with high-dimensional parametrisation schemes [22–29]. However, the method is not exempt of limitations. As described in [17, 30], the residual of the Navier-Stokes equations is assumed to be zero when computing the adjoint, which requires the primal CFD to achieve exceptional convergence and a near-zero residual. This implies longer solve times than the ones that would be employed for traditional CFD analysis and a corresponding increase in the computational expense of the flow solver. Moreover, as documented in [23], typical running times for discrete adjoint solvers are about three times higher than the ones of their primal counterpart, making the adjoint calculation a costly endeavour. In optimisation scenarios involving several objective function or constraints that depend on the flow behaviour, the adjoint calculation must be repeated for each objective and constraint, reducing the scalability of this approach. In addition, much of the outcome of gradient-based approaches is dictated by the initial design and, if starting from already high-performing shapes close to a local optimum, the optimisation process might not yield significantly improving designs, regardless of the parametrisation scheme [31, 32].

An alternative to gradient-based solvers are global search algorithms. These only require zeroth order information from the objective and constraint functions, which can usually be obtained with one CFD computation for any number of them. Moreover, these algorithms are stochastic in nature and this, in principle, makes them capable of identifying the global optimum of the objective function [33–35]. The application of such algorithms

to turbomachinery design optimisation has been an active area of research, with approaches falling into one of two categories: *direct* application to the CFD solution [36–41] and *surrogate based* optimisation, where a number of computations are first employed to regress the behaviour of the cost function in the design space and the search algorithm is subsequently applied on this analytical model [42–46]. Global search algorithms typically require function evaluations on the order of hundreds to thousands, depending on the dimensionality of the problem, which causes direct approaches on high-fidelity CFD to be prohibitive for most academic and industrial environments.

Indeed, with the advent of machine learning and data-driven methods, novel surrogate models are being researched which enable highly accurate predictions over complex high-dimensional datasets, making them ideal for global optimisations [47–50]. In particular, *Artificial Neural Networks* (ANNs) have become a primary area of study due to their capability to interpolate unstructured data [51, 52]. Unlike other machine learning methods, ANNs do not assume a predefined shape for the cost function, thus providing a higher level of flexibility and predictive accuracy in cases where the system cannot be well expressed through explicit functions. In the field of turbomachinery, they have been applied to develop physics-based models for unsteady compressor behaviour [53], to predict the performance of compressor cascades from design considerations [54, 55], and on global, surrogate-based optimisation environments [56–59]. However, the fact that they are unstructured also implies that the number of data points required to adequately fit them to high-dimensional, multi-modal systems –such as the parametrised flow behaviour of transonic blades– is typically higher than other methods. Thus, the cost-effectiveness of ANNs can be enhanced by coupling them with dimensionality reduction (DR) algorithms.

Several of such DR algorithms exist and have been extensively documented, such as Factor Analysis [60], Principal Component Analysis (PCA) [61], Discriminant Analysis [62], Sobol’ sensitivity indices [63], etc. [64]. In recent years, a novel set of ideas that facilitate subspace-based dimension reduction has emerged, called *Active Design Subspaces* (ADS) [65]. Instead of determining a subset of the input parameters which are most important, ADS identifies dominant linear combinations of all the parameters that best describe the variability in the output. While PCA is also able to identify such subspaces from the parameter space, with that approach the data must be correlated or conditioned by some process, like the pareto-optimal design points from a multi-disciplinary optimisation cycle [37]. ADS, on the other hand, can learn the dominant subspaces from a non conditioned dataset, which enables its application to uniformly-spread sampled points, such as those coming from a Design Of Experiment approach, which are typically

used in regression processes. In the field of turbomachinery, active subspaces have been linked with fundamental aerodynamic principles to infer pedigree design rules to improve a compressor's performance [66], to visualise the robustness of a design point with respect to uncertain parameters [26] and to identify dominant subspaces for stagnation temperature probes [67]. The ADS process requires estimating the gradient of the cost function at various random points in the design space, for which [26] employed adjoint computations and [66] regressed a second order polynomial to CFD data, and evaluated the gradients from that model. Both these approaches suffer from previously described issues: the former does not scale well for multi-objective scenarios and the latter infers the function behaviour can be well described by a second order polynomial.

1.2 Motivation

This work develops a novel strategy to approach modern engineering design, particularly in the field of gas turbines for aviation, called AI Enabled Active Subspaces, or *AInAS*. This strategy exploits the unstructured nature of ANNs by closely coupling them with ADS to reduce the dimensionality of the input space, thus providing an ameliorated performance with minimal sample requirements. The framework also supports enhanced visualisation of the system under study, which leads to increased understanding of its behaviour, in addition to parameter sensitivity ranking, which proves instrumental in identifying the leading drivers of high dimensional systems. The development of this methodology was driven by the needs for such a framework in the design and analysis of fan blades for modern turbofan engines. As a result, sections of this text are devoted to concepts regarding the aerodynamic performance of the resulting novel designs.

Three distinct applications are presented which employ AInAS to approach challenging engineering design problems. The first of these, intended mostly as a proof-of-concept, performs the aerodynamic efficiency optimisation of a fan blade under pressure ratio constraints. The capabilities of the AInAS method are compared against a state-of-the-art gradient-based approach using adjoint. It will be shown that the two methods achieve comparable improvements in performance, but the novel strategy is both faster and more economical. The second application makes use of AInAS to characterise a highly complex system, namely, the aerodynamic stability of a next generation fan blade. The effect of three dimensional geometry on the tip leakage flow is assessed and its consequences for stability range are derived, leading to an increased understanding of the phenomena in place and the

opportunity to design higher performing shapes which are also more stable (further away from stall). The final application performs a task of particular engineering relevance and at the limit of current gradient-based capabilities. The role of manufacturing variability in the final shape of the fan blade will be addressed and optimisation under uncertainty is performed, aiming to improve the mean performance of the fan blade and its sensitivity to typical geometry deviations introduced in manufacturing.

Chapter 2

AI Enabled Active Subspaces

2.1 Engineering Design Optimisation

The behaviour of engineering systems is usually described through a set of functions of interest, $\{F_k(\mathbf{x})\}$, where $k = 1, \dots, K$, that model the response observed to a set of inputs \mathbf{x} , where $\mathbf{x} \in \chi \subset \mathbb{R}^m$. Here, χ denotes the design space, where all the allowable values for the inputs to F_k live. The degrees of freedom for the system are encoded in the dimensionality of the independent variable, m , which, as described previously, can be arbitrarily large.

The exact form of these functions is typically unknown or very complex, and usually models are employed through which only functional calls are available. In other words, the information is typically limited to the value of the functions at certain discrete points contained in the design space. For engineering relevance, it is commonly the case that understanding the way F_k varies with changes in \mathbf{x} provides substantial value to designers. In particular, answers to the following questions are desired:

- Which design features are associated with the output? Are all m degrees of freedom equally significant, or can the system be well described using only a subset of them?
- What effect does a change in x_p cause to F_k ? Does increasing x_p increase the value of F_k or is the opposite true?
- Can the relationship between F_k and \mathbf{x} be described through an analytical model or is the relationship more complex?

The answer to these questions requires making use of available samples $(F_k^{(r)}, \mathbf{x}^{(r)})$ to build a statistical learning model as per (2.1), where the error term is assumed to be normally distributed with zero mean.

$$F_k(\mathbf{x}) = \hat{f}_k(\mathbf{x}) + \epsilon \quad (2.1)$$

The analysis of \hat{f}_k leads to statistical inference driven by the search for answers to the questions about the system. This constitutes one of the main objectives of an optimisation workflow, which is commonly overlooked in lieu of the other objective, which gives the process its name.

Indeed, the second objective is tasked with optimising the set of inputs \mathbf{x} to achieve a desired performance improvement. In this case, the $\{F_k(\mathbf{x})\}$ are partitioned in two groups, objective functions and constraints. A general optimisation problem can be expressed as in Eq. 2.2 for n_o objective function and n_c constraints, where $n_o + n_c = K$. Note that, conventionally, the objectives are *minimised*. For functions that are to be maximised, it is sufficient to negate their response.

This problem is typically solved employing numerical search algorithms that require repeated evaluations of the objective and constraint functions and, in some cases, higher order information like the functions' gradient or Hessian matrices. The number of iterations of the optimisation process is usually dictated by the dimensionality and range of the design vector, \mathbf{x} and whether local or global optima are desired.

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimise}} && \{f_i(\mathbf{x})\} ; i = 1, \dots, n_o \\ & \text{subject to} && g_j(\mathbf{x}) \leq h_j, j = 1, \dots, n_c. \\ & && \mathbf{x} \in \chi \subset \mathbb{R}^m. \end{aligned} \quad (2.2)$$

For single-objective problems, the output of (2.2) is a single design point \mathbf{x}^* where the system achieves its peak performance, that is $f(\mathbf{x}^*) \leq f(\mathbf{x})$ for $\mathbf{x} \in \chi$. In cases where $n_o > 1$, due to typically conflicting objective functions, it is appropriate to generate a Pareto set of solutions $\{\mathbf{x}_o\}$ where no solution *dominates* another. In other words, for every \mathbf{x}_o it is not possible to further minimise a certain objective function without worsening the performance of at least another objective [68, 69]. Equation 2.3 expresses this relationship, where \mathbf{x} dominates \mathbf{x}' if and only if:

$$\begin{aligned} & f_i(\mathbf{x}) \leq f_i(\mathbf{x}') \text{ for all } i \text{ in } \{1, \dots, n_o\} \\ & \text{and } f_i(\mathbf{x}) < f_i(\mathbf{x}') \text{ for some } i \end{aligned} \quad (2.3)$$

Due to the typically very large number of functional calls required to find optimal designs, it is standard practise to refer to models such as (2.1). However, in this case, both for the single- and multi-objective frameworks, the $\{\hat{f}_k\}$ can be treated as black boxes, since the greatest concern is not




Figure 2.1: General structure of perceptron neuron for three inputs

quite understanding their behaviour, but rather the accuracy with which comparisons can be made to detect dominated designs.

Achieving the engineering optimisation objectives therefore requires a model which is both interpretable to allow inference and highly accurate to enable identification of non-dominated designs. The strategy developed in this study is designed to meet these objectives and it is particularly tailored for cases where a global optimum is sought, the inputs' dimension, m , is arbitrarily large and functional calls are expensive such that the available sampled data is scarce. Moreover, no inferences are made on the complexity of the objective or constraint functions.

2.2 Fitting Artificial Neural Networks

Artificial Neural Networks (ANNs) are mathematical constructions where information from a set of inputs, \mathbf{v} , is linked, in a non-linear manner, to a set of outputs $Y = \hat{f}(\mathbf{v})$ by having the information flow across a structure of neurons, whose behaviour can be trained to achieve a desired outcome. ANNs are selected to model the behaviour of the cost function S and constraints because, as detailed in [70, 71], provided there are sufficient neurons in the network, ANNs are able to fit any function of arbitrary shape, which makes them ideal for interpolating complex multi-modal systems, typically encountered in turbomachinery design optimisation problems. Moreover, in the presence of high input dimensionality, the number of training samples required to fit ANNs grows more slowly than other regression methods [72, 73], making them more robust to the curse of dimensionality.

The fundamental element in ANNs is the perceptron neuron, shown in Fig. 2.1 for an example input vector of three dimensions. The parameters influencing the neuron's behaviour are the weights, \mathbf{w} , the bias, b , and the activation function F , where the response follows (2.4). In this way, the




Figure 2.2: Activation functions for perceptrons

value of each input is ranked with respect to their effect on the output by the weights. The bias encodes a minimum energy level associated with the neuron and defines the average activation value. Finally, the activation function provides an opportunity to introduce non-linearity to the output. Notice in (2.4) that if F takes the form of the identity function, the system collapses into a traditional linear model of the inputs.

$$y = F(\mathbf{w}^T \mathbf{V} + b) \quad (2.4)$$

This work considers two types of activation function to construct perceptrons, sigmoid and Rectifier Linear Unit (ReLU), shown in Fig. 2.2. The existence of the linear part in ReLU enables implementation of continuous activation functions on the basis of such elements, while the sigmoid function possesses invertibility properties and continuous differentiability properties [74].

A traditional perceptron feed-forward network structure is constructed from the combination of multiple individual neurons producing highly flexible mapping machines. The collection of neurons placed at the same level in the structure form a *layer*. The first layer is designed to match the dimensionality of the inputs, and the last layer is formed of a single neuron, which outputs the network's prediction, as shown schematically in Fig. 2.3. The weights of every neuron connection and the biases of individual neurons are defined through an optimisation process where the cost function is designed




Figure 2.3: Schematic of the feed-forward neural network employed

to maximise the network's prediction accuracy over a training dataset, for which the true values of the function are known. This work makes use of the Mean Squared Error (MSE) loss function, computed between the network's predictions and the real function value for the training samples. The back-propagation algorithm [75] is used to compute the gradient of the network's parameters with respect to the inputs and facilitate the training process. Because gradient calculation through backpropagation is expensive, this process is performed in mini-batches of samples, selected randomly from the whole set. Not only does this approach ease the computational resources required to train the network, but mini-batching also provides a stochastic approximation of the expected value of the gradient of the loss function over the training set, which has been shown to improve generalisation performance and optimisation convergence [76, 77]. A variation of stochastic gradient descent, RMSProp [78] with a learning rate of 0.001, is employed to make use of these gradients and tune the network.

One pass through all the training data is called a training epoch. In this work, the number of training epochs is defined via an Early Stopping callback [79], monitoring the convergence of the MSE evaluated on a validation set comprising 20% of the samples, which are spared and not used for training.

Table 2.1: Hyper-parameters tuned in grid search

Parameter	Lower Limit	Upper Limit
No. of HL	1	4
No. of neurons in each HL	5	500
Regularisation Coefficient	0.0	0.02
Drop-out Factor	0.0	0.5
Activation Function	sigmoid	relu

A maximum of 500 epochs without change in the MSE of the validation set is specified to prevent over-fitting.

In fact, due to the potentially very large number of neurons being employed, over-fitting the neural networks is a possible but unwelcome outcome of the training process. To prevent this scenario, in addition to early stopping, a Tikhonov regularisation term [80] is included in the MSE loss function, as well as neuron drop-out [81, 82].

The Tensorflow framework [83] was employed to construct and train the networks, through its python wrapper Keras [84]. It was discussed that the network design at the input and output layers follows the dimensionality of the problem. The layers in between are called *hidden* and are a major consideration when deciding the network’s architecture. How many Hidden Layers (HL) to use and the number of neurons per layer define the networks capacity, its ability to map complex functions. In this work, both these considerations are defined through a hyper-parameter tuning process.

Since this work deals mainly with small datasets, in the order of hundreds of samples, the training of the networks is not a particularly costly procedure. Therefore, the hyper-parameter tuning can be achieved through a grid-search, evaluating a large number of different parameter combinations. For the activation function, as described, only sigmoid and ReLU were employed. Table 2.1 shows the hyper-parameters tuned and their corresponding high and low levels. The objective function for the tuning process was set to maximise the coefficient of determination, R^2 , using 10-fold Cross Validation (CV) [85]. Some network parameters are general for this application and thus are fixed for all network configurations evaluated in the grid search:

- Number of neurons in the input layer: set to match the dimensionality of the design vector.
- Output layer definition: a single neuron without any activation function.

Thus, the network fitting process proposed requires two datasets at every CV step: one for training the weights and biases and a second for tuning the network architecture. The former is additionally split to allow for a validation set and further prevent over-fitting. The proportions of the total number of samples being placed in each dataset is as follows: 64% for the weights and biases training, 20% for hyper-parameter tuning and 16% for early stopping. The splitting of the dataset is performed using seeded random number generators, such that each network evaluated in the grid search is trained and tested using the same samples. It is worthy to note that the samples employed in this study correspond to a same model. While current AI research indicates NN capabilities that could enable the generalisation of these networks to other models, these features have not been considered in this work.

The performance of the network, as determined by the MSE and R^2 , was discovered to increase significantly when trained and tested on inputs and outputs that were normalised to have zero mean and a standard deviation of one, obtained through the transformation defined in Eq. 2.5 for a given vector \mathbf{v} .

$$\mathbf{v}_n = \frac{\mathbf{v} - \mathbb{E}[\mathbf{v}]}{\sigma(\mathbf{v})} \quad (2.5)$$

2.3 Coupling Active Design Subspaces

Adequately fitting a neural network to high-dimensional input data might require a large number of samples, which can make the method infeasible in cases where function evaluations are obtained through a costly procedure. To enable the application of neural networks to high-fidelity CFD data, this work proposes coupling them with Active Design Subspaces (ADS) [65].

2.3.1 Active Subspaces

Let F_k be a function of interest to the engineering system under study, where F_k is differentiable and Lipschitz continuous. Let \mathbf{x} be an m -dimensional input vector to F_k living in a design space χ as per (2.2), such that $F_k : \chi \rightarrow \mathbb{R}$, where $\chi \subset \mathbb{R}^m$. A condensed vector \mathbf{y} is sought, where $\mathbf{y} \in \Phi \subset \mathbb{R}^k$, with $\Phi \subseteq \chi$ and $k \leq m$. If \mathbf{y} is such that the variability in the function is expressed more efficiently, $F_k(\mathbf{y})$ is simpler to optimise.

Two approaches can be followed to identify \mathbf{y} . One is to identify a subset of $\{x_i\}$ that, when varied, do not significantly affect the response of F_k . Following, \mathbf{y} is formed of all $\{x_j\}$, where $j \neq i$. This is referred to as screening

and can lead to powerful dimensionality reduction. However, the expense of confidently screening all the input parameters can well exceed the available resources in high dimensions. Another approach is to identify an optimal basis in \mathbb{R}^m that is tuned to efficiently express F_k . Given that \mathbf{x} is full-rank, the new basis can be identified by a rotation of the original coordinate system, or, likewise, a linear combination of the original dimensions. In such a system, it's straightforward to realise directions along which F_k varies little and perform screening. This is called subspace-based dimensionality reduction. Moreover, thresholds can be implemented in the linear combination that produces the new basis, such that low coefficient values indicate parameters that can be ignored without significant loss of information. It can be noted, therefore, that subsets are a generalisation of screening tests that enable more flexibility in the implementation of the dimensionality reduction.

The ADS approach identifies the linear combination of the input parameters that best describes the variability in an objective function, through the eigenvalue decomposition of the function's gradient covariance matrix, C , defined in Eq. 2.6, from [86].

$$C = \mathbb{E}[\nabla_x F_k \nabla_x F_k^T] \quad (2.6)$$

The eigenvalue decomposition of C , computed as $C = W\Lambda W^T$ yields the dominant directions in the columns of W . Based on the decay of the eigenvalues, W and Λ can be partitioned as per (2.7), such that the active subspace is captured in the matrix W_1 , which maps the m -dimensional inputs, \mathbf{x} , to a k -dimensional *active* vector, \mathbf{y} , through the transformation expressed in (2.8).

$$\Lambda = \begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix}, \quad W = [W_1 \ W_2], \quad W_1 \in \mathbb{R}^{m \times k} \quad (2.7)$$

In practice, the matrix C is approximated through a Monte Carlo method, by drawing M independent samples $\{\mathbf{x}_i\}$ according to the sampling density $\rho = \rho(\mathbf{x})$ in χ , and computing the gradient for each sample, as defined in Eq. 2.9. Constantine [65] suggests adopting $M = \alpha k \log(m)$, where α is an oversampling factor, and k is the maximum number of dimensions in the ADS that can be accurately resolved.

$$\mathbf{y} = W_1^T \mathbf{x} \quad (2.8)$$

$$C \approx \frac{1}{M} \sum_{i=1}^M \nabla_x F_{K_i} \nabla_x F_{K_i}^T \quad (2.9)$$

The application of ADS to the optimisation problem stated in Sec. 2.1 requires building such a covariance matrix for all the F_k functions. If the number of functions of interest, K , is sufficiently low and an adjoint code is available, then computing the gradients with adjoint can be a viable alternative, requiring M primal CFD computations and $M \times (K + 1)$ adjoint computations. However, for a general scenario, K might be overly large or adjoint codes might not be available. In such scenarios, an analytical model such as (2.1) can be trained on sampled data and the gradients estimated from said model. The computational expense of building C in this case is dependent on the model employed. Based on the discussion from Sec. 2.2, this work employs ANNs for this purpose.

2.3.2 Enabling ADS through AI

An ADS is sought to reduce the number of samples needed to adequately fit a neural network to high-dimensional input data, and a neural network is required to learn this ADS. This impasse is solved through the iterative process, shown schematically in Fig. 2.4. This process is defined based on the knowledge that, regardless of the number of input parameters, as an ANN is trained with increasing number of samples, provided there is no overfitting (for which the fitting procedure in Sec. 2.2 accounts), the prediction of said network will converge to the true function from which the samples were taken. Thereby, the estimation of C via gradients evaluated with the ANN also approximates the true matrix and its eigenvector decomposition converges to the same set of dominant directions. This work proposes iterating on the number of samples employed to fit the ANN while monitoring the convergence of the first k eigenvectors of C , through Eq. 2.10, which measures the angular variation, θ , between the estimated dominant directions from successive iterations.

$$\theta_i^{(j)} = \cos^{-1} \left(e_i^{(j)} \cdot e_i^{(j-1)} \right); \quad i = 1, \dots, k. \quad (2.10)$$

Where, $e_i^{(j)}$ is the i th eigenvector of C , estimated at the j th iteration.

Once θ is small enough, the ANN prediction approximates the function behaviour in the dominant directions sufficiently well to reveal the ADS. This can therefore be exploited to reduce the dimensionality of the inputs. Mapping the samples to the active directions and fitting a final ANN to this low-dimensional dataset generally increases the accuracy of the network predictions since the data is now structured and there are more data points per dimension. The initial number of samples, n , and the increment, p , in the

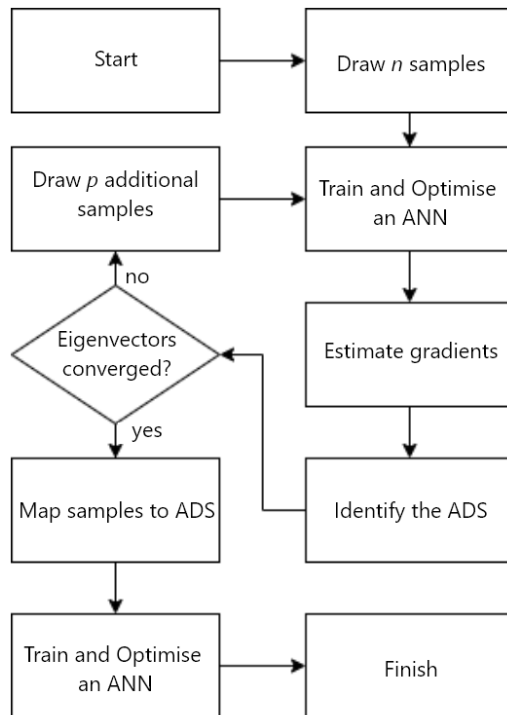


Figure 2.4: Iterative process proposed to learn the ADS and fit an ANN to CFD samples

iterative process can be selected to exploit parallel computing capabilities, if available, to reduce the time requirements of the loop.

2.4 Rewriting the Optimisation Problem

The transformation defined in Eq. 2.8 is called a *forward map* and, as described, transforms the input vector $\mathbf{x} \in \mathbb{R}^m$ to an active vector, $\mathbf{y} \in \mathbb{R}^k$. There is no reason why the active vectors of different functions should be the same, hence, after the fitting procedure, K ANNs are obtained, all responding to different active vectors in different subspaces. This section details a reformulation to the optimisation problem described in Eq. 2.2 that allows to *navigate* through the active subspaces to find an optimum in the m -dimensional space, while exploiting the low-dimensional structure discovered through the ADS.

The forward map provides a unique vector \mathbf{y} for each vector \mathbf{x} . However, the converse is not true. There are infinitely many \mathbf{x} that satisfy the inverse

map for a given \mathbf{y} . Let \hat{f} be the neural network prediction of the cost function for its active vector:

$$f(\mathbf{x}) \approx \hat{f}(W_{1_0}^T \mathbf{x}) \quad (2.11)$$

Where W_{1_0} contains the first k eigenvectors of the cost function's C matrix. Additionally, let the constraints be approximated by a neural network built on their active vector:

$$g_i(\mathbf{x}) \approx \hat{g}_i(W_{1_i}^T \mathbf{x}); \quad i = 1, \dots, n_c. \quad (2.12)$$

A new function F is defined as:

$$F(\mathbf{y}_0) = \hat{f}(W_{1_0}^T \mathbf{x}^*) \quad (2.13)$$

Where,

$$\begin{aligned} \mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \quad & \frac{1}{2} \|W_{1_0}^T \mathbf{x} - \mathbf{y}_0\|_2^2 \\ \text{subject to} \quad & \hat{g}_i(W_{1_i}^T \mathbf{x}) \leq h_i, \quad i = 1, \dots, n_c. \\ & \mathbf{x} \in \chi \subset \mathbb{R}^m. \end{aligned} \quad (2.14)$$

Each evaluation of $F(\mathbf{y}_0)$ requires solving a constrained least-squares minimisation problem designed to select from the infinitely many \mathbf{x} that solve the inverse map, one that is feasible in terms of the n_c constraints. In case there are no feasible points that satisfy this, then the feasible \mathbf{x} that maps to the *closest* active vector to \mathbf{y}_0 is selected. This is a convex minimisation problem that can be easily solved by a gradient-based search algorithm, employing finite differences for the gradients of the constraints, $\nabla \hat{g}_i$, and Eq. 2.15 for the cost function gradient. It is worthy to note that there might still be infinitely many \mathbf{x} that solve Eq. 2.14. In such cases, all design points are considered equal since they map to the same active vector and only differ in their *inactive* directions. Hence, the cost function should experience little variation between these points. The introduction of a regularisation term (e.g. Tikhonov) to the cost function in Eq. 2.14 is possible without loss of generality, in the case where there is preference for particular types of \mathbf{x} . In Ch. 4, this approach is followed.

$$\nabla_{CF} = W_{1_0} (W_{1_0}^T \mathbf{x} - \mathbf{y}_0) \quad (2.15)$$

The constrained optimisation problem defined in Eq. 2.2 can thus be rewritten as an unconstrained optimisation through function F, as per Eq. 2.16.

$$\min_{y_0} F(\mathbf{y}_0) \quad (2.16)$$

This formulation enables employing the ANN built on the cost function's ADS for global optimisation. Exploiting the dominant directions usually accelerates the convergence of the optimisation procedure since the function is very respondent to changes in its active variable. Moreover, selecting the feasible point through Eq. 2.14 allows employing the ANNs built for the constraints, which can be obtained cheaply, through the algorithm described in Fig. 2.4.

Chapter 3

Benchmarking AInADS performance

Nomenclature

\hat{f}	Neural Network approximation of function f
C	Gradient Covariance Matrix
W	Eigenvectors of the covariance matrix
Λ	Eigenvalues of the covariance matrix
m	Dimensionality of design vector
k	Dimensionality of active vector
n_c	Number of constraints in optimisation problem
\mathbf{x}	Vector of design parameters
\mathbf{y}	Vector of design parameters in the active directions
$\mathbb{E}[\cdot]$	Expected value of the argument in $[\cdot]$
$\sigma(\cdot)$	Standard deviation of the argument in (\cdot)

3.1 Comparison to an adjoint gradient-based approach

To compare the performance of the global strategy proposed, the optimisation problem in Eq. 2.2 is also solved employing an adjoint-based technique. The search algorithm chosen is the Sequential Least Squares Programming (SLSQP) from scipy [87]. The method wraps the SLSQP routine developed by [88]. It employs a Han-Powell quasi-Newton method with a BFGS update of the B-matrix for defining the search direction. This update requires gradi-

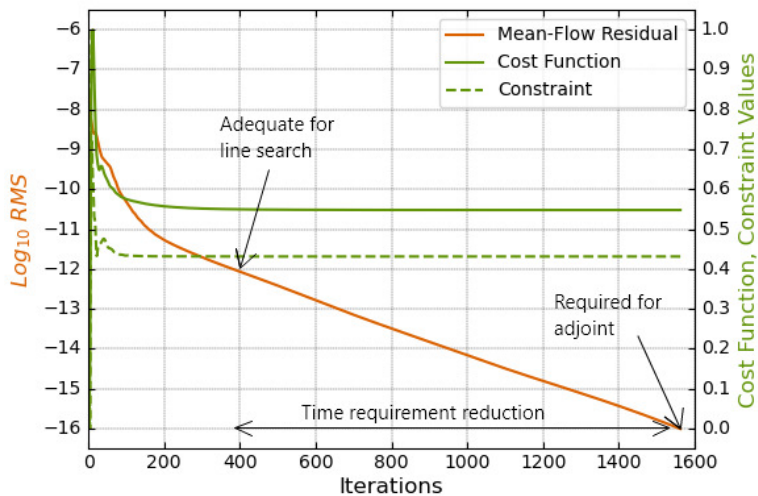


Figure 3.1: Primal convergence history for mean-flow residual and non-dimensionalised cost function and constraint values

ent information of the cost function and the constraints, hereby referred to as the *objectives*, which in this section is obtained through the adjoint approach. Next, SLSQP selects the optimum step size via a line-search optimisation, which only requires solving a primal to obtain the the objectives' value at each iteration.

As discussed in Sec. 1.1, the adjoint computation requires a near zero residual from the primal solver which increases the computational expense of that stage. Additionally typical adjoint running times are about three times longer than the flow solver, making the B-matrix update a costly endeavour. This work aims to lessen the overall time expense by employing parallel computing capability and solving the adjoint computation for the objectives simultaneously.

For the step-length algorithm, there is no need for the primal to achieve such exemplary convergence and adequate results can be obtained much earlier. Fig. 3.1 shows a typical CFD primal convergence history for an axial fan. By iteration 400, the objective convergence curves have stabilised and only minor variations in their value is noticed thereafter. From this point, in traditional CFD-based SLSQP implementations, the primal solver would be continued for an additional 1200 iterations to comply with the adjoint convergence requirements. This implies a fourfold increase in the computational expense for each line search iteration while there is no such requirement. Hence, this work proposes employing partially converged values to accelerate




Figure 3.2: Research blade CFD domain

the step-size optimisation of the algorithm.

3.2 Application to a jet engine fan blade

In this section, the optimisation of a modern, low speed, high bypass ratio research fan blade is approached with the global strategy described in Ch. 2, and the adjoint-based strategy from Sec. 3.1. The test case under study is called VITAL, hereby referred to as the *research blade*, shown in Fig. 3.2. As documented in [89], for high bypass ratio fans, a 1.4% increase in efficiency yields a 1.0% reduction in the engine's specific fuel consumption, making this component an ideal test case to study the benefits of global optimisation on high-dimensional parametrisation schemes. The span of the research blade is about two-thirds smaller than that of a conventional fan blade making it suitable for rig tests. The rotational speed has been adjusted to emulate the flow physics at cruise condition and the thickness has been increased to achieve the mechanical integrity required of a large aero-engine composite fan blade.

3.2.1 Computational Tools

The Rolls-Royce proprietary CFD code Hydra [90] was used throughout this study to simulate the flow about the research blade. Hydra is an unstructured solver employing 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. Hydra's discrete adjoint capability was employed in this work to estimate the gradients of cost function and constraints with respect to the design parameters.

The computational domain used in this study is shown in Fig. 3.2. It is a single-passage, single-blade row model with the downstream splitter. The whole domain is modelled on a rotating frame, with the casing, splitter, inlet and exit surfaces set as stationary. The rotor, hub and splitter surfaces are set as viscous walls. At the inlet, a one-dimensional Boundary Condition (BC) is enforced, specifying a radial distribution of total pressure, total temperature, whirl and pitch angles and turbulence intensity, where the values for these quantities are obtained from experimental analyses. For the bypass and core exit surfaces, a non-reflecting, radial-equilibrium capacity exit BC is enforced.

The domain is discretised using the Rolls-Royce proprietary geometry and meshing software, PADRAM [91]. The blocking strategy employed by PADRAM consists of an H-O-H topology, with H blocks for the upstream and downstream regions, as well as the upper and lower periodic boundaries. The blade is enveloped in an O-mesh while a C-mesh is employed for the splitter. A mesh convergence study was previously undertaken to identify the optimal distribution of nodes [26], 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. CFD-experimental validation of this set-up have previously been reported achieving a good match for a number of different operating conditions [26, 27].

3.2.2 Optimisation Problem

The definition of the geometry parametrisation is a critical factor dictating much of the success that can be obtained through optimisation, since the search will only comprise geometries that are attainable with it. Many parametrisation schemes have been introduced in the literature, such as Free-Form Deformation or B-Splining. However, most of these are abstract in the sense that the user has little knowledge on how a particular parameter affects the geometry, and indeed much less on how the interaction between various parameters affects the final shape. In this work, the geometry parametrisation is defined through PADRAM's *Engineering Design Parameters* (EDP), which comprise a set of intuitive geometry manipulation handles based on first principles. Each EDP, illustrated in Fig. 3.3 for an aerofoil section, con-




Figure 3.3: Geometry parametrisation through EDP. The magnitude of the perturbations has been enlarged for clarity

trols 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*. Two additional DOFs controlling the locality of the LE and TE recambering are also introduced, such that low values of these parameters cause very localised camber line alterations, and vice-versa. Sufficiently large values can propagate the perturbations through the aerofoil, thus providing complete control over the camberline. The EDP are applied on five aerofoil *control sections* uniformly distributed through the blade span - at 0%, 25%, 50%, 75%, 100% - providing a total of 35 DOFs, arranged in the design vector, \mathbf{x} . The value of the deformation applied as a function of the blade span is achieved through

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

The upper and lower optimisation ranges for each DOF employed in this study, denoted by the vectors \mathbf{x}^U and \mathbf{x}^L respectively, were defined based on previous experience with the parametrisation scheme. Let χ be the space of possible designs attainable with the 35 DOFs described, the design space is thus defined as follows:

$$\{\mathbf{x} \in \chi \mid \mathbf{x}^L \leq \mathbf{x} \leq \mathbf{x}^U\} \quad (3.1)$$

The cost function (CF) for the optimisations carried out in this work is the fan bypass isentropic efficiency, defined in Eq. 3.2, where the total pressure and total temperature quantities employed are extracted from the CFD solution by performing double averaging over the surface of interest: mix-out circumferentially and mass-mean radially.

$$CF = \eta(\mathbf{x}) = \frac{\left(\frac{p_{0_{exit}}}{p_{0_{inlet}}}\right)^{\frac{\gamma-1}{\gamma}} - 1}{\frac{T_{0_{exit}}}{T_{0_{inlet}}} - 1} \quad (3.2)$$

$$0.99 PR_{datum} \leq PR(\mathbf{x}) \leq 1.05 PR_{datum} \quad (3.3)$$

Through the optimisation, upper and lower constraints are enforced for the fan pressure ratio (PR) according to Eq. 3.3. The lower bound is specified to prevent new geometries from maximising efficiency by greatly reducing the PR, which would cause the LP shaft to operate at a higher speed to meet the engine's thrust requirement. Similarly, the upper bound is enforced to avoid designs with overly large PR, that operate at significantly lower shaft speeds and can lead to overloading of the LP turbine blade. It is worthy to note that the EDP employed do not alter the thickness of the blade, thus preventing the search for an aerodynamically optimum design to lead to overly thin blades which would significantly affect the mechanical integrity of the blade.

3.2.3 Global Optimisation Approach

In this section, the research blade is optimised employing the global strategy detailed in Ch. 2. The iterative process from Fig. 2.4 is solved for efficiency and PR, employing $n = 105$ initial function evaluations with an increment of $p = 17$. The sampling was done following a Design Of Experiment approach employing Sobol' sequence. The convergence of the algorithm for the first two dominant directions is shown in Fig. 3.4. The trend demonstrates a decreasing magnitude of the angular variation between successive iterations,




Figure 3.4: Eigenvector convergence

suggesting the C-matrix estimation via the neural network predictions is converging to its true form. By the eighth cycle, corresponding to 250 total samples, a sufficiently low angular variation has been achieved. Certainly, the process could be continued with increasing number of samples leading to an ameliorated resolution in the active directions, however, the trend shown in the figure suggests that only minor variations in the directions can be expected thereafter. Hence, the process is deemed converged at the eighth iteration.

Figures 3.5a and 3.5b plot the eigenvalue decay of the final efficiency and PR gradient covariance matrices. The rapid decrease noticeable in the eigenvalues is positive because it denotes that most of the system’s variance is captured by just a few dominant directions. In fact, through the cumulative energy plots shown in Fig. 3.5c and 3.5d it is noticeable that almost 100% of the total system’s energy can be represented by the first 4 eigenvalues. Thereby, for the optimisation task at hand, the ADS was constructed employing the first four dominant directions.

Through Eq. 2.8, the 250 samples were mapped to the PR and efficiency active subspaces, and a final neural network was trained and tuned for both functions. Table 3.1 gathers the prediction accuracy obtained with these networks and compares it against the ones trained on the high-dimensional space. A noticeable increase in prediction accuracy is achieved for both functions when training the networks with the low-dimensional space, due to the fact that the data is structured and the ratio of samples to dimensions is greatly increased. The ameliorated network performance provides an ad-

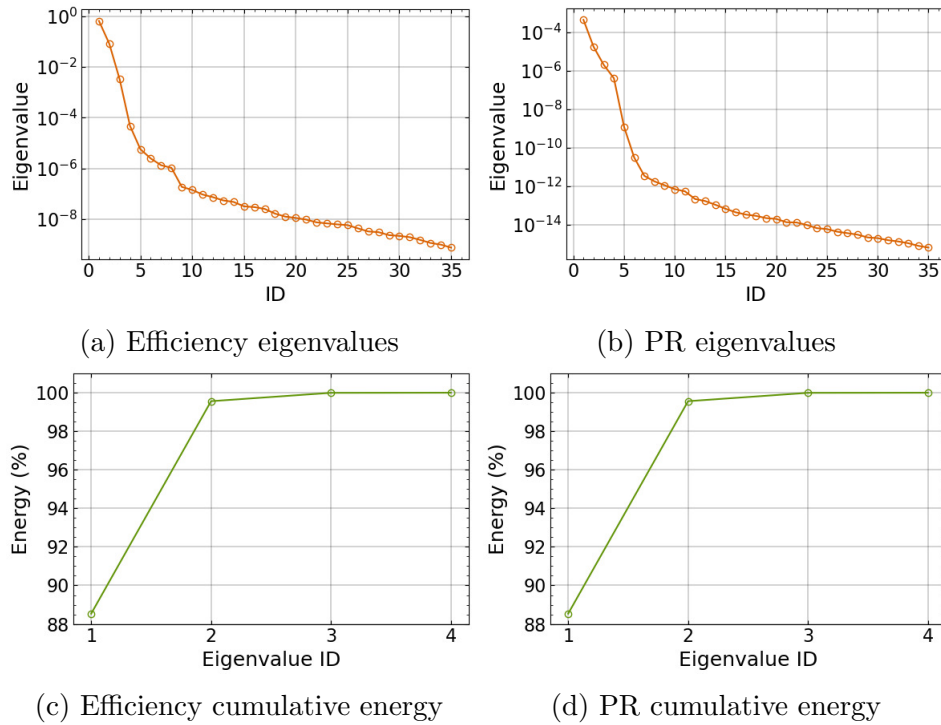


Figure 3.5: Eigenvalue decay and cumulative energy plots for efficiency and pressure ratio gradient covariance matrices

ditional level of confidence in the subsequent surrogate-based optimisation process, since the performance benefit of any optimised designs estimated from the networks should be close to what CFD would predict.

For visualisation purposes, the data can be condensed even further. The cumulative energy plots show that over 98% of the energy is captured by just the first two eigenvalues and this allows building 2-dimensional active subspace performance maps, following [66], without much loss of information. These maps, shown in Fig. 3.6, enable visualising the behaviour of the objective functions in their active subspaces. The active vector coordinates (0,0) in these plots corresponds to the datum design. The efficiency map shows that this design lies in a corridor of high-performance, but could still be improved quite significantly by moving in the second active variable's positive direction. However, such design changes could affect the PR beyond the specified tolerances. The active subspace on which the PR map is constructed is different than that for efficiency, preventing a direct combination of both plots. Through the reformulation of the optimisation problem, discussed in Sec. 2.4, the search algorithm will be able to navigate through these subspaces searching for the best trade-off design.

Table 3.1: Network prediction accuracy trained with high-dimensional inputs and active vectors

Function	$R^2 [\hat{f}(\mathbf{x})]$	$R^2 [\hat{f}(\mathbf{y})]$
Efficiency	0.892	0.937
PR	0.987	0.990

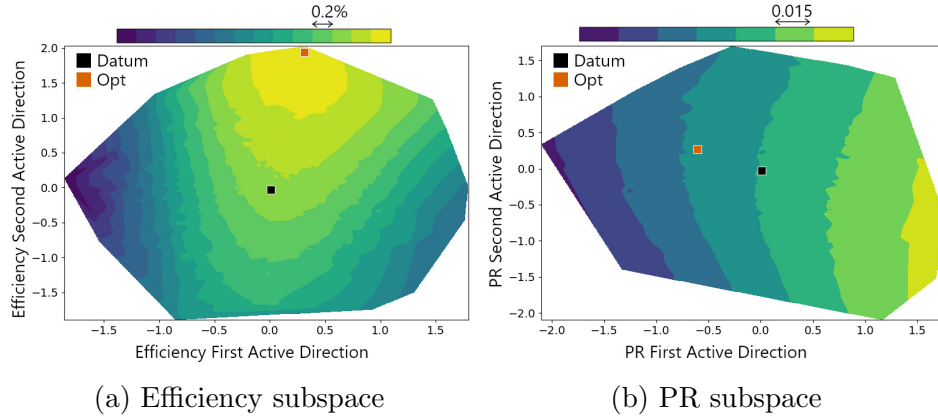


Figure 3.6: 2D active subspace maps for the optimisation functions

In this work, a genetic algorithm was employed to perform the search for a globally optimum design using Eq. 2.16. Upon convergence, the optimal design point was simulated with CFD leading to a significant 0.47% improvement in efficiency, while the PR constraint was achieved.

Figure 3.7a plots the radial profile of circumferentially mixed-out values of efficiency (as per Eq. 3.2) for the datum and optimised designs. This graph reveals that the performance increase for the optimum design arises largely from the upper 20% of the blade's span, while a slight reduction in efficiency is noticeable for the mid span region, due to the radial adjustment of the flow. The isentropic Mach number distribution for the 90% span section, plotted in Fig. 3.7b, suggests that the primary mechanism for this performance benefit is an improved shock behaviour via the reduction of the pre-shock Mach number. Additionally, the pressure side spike has been mitigated, leading to a more uniform loading at the leading edge (LE). The suction side LE loading, however, has slightly been worsened for the optimum, noticeable by a minor peak followed by a sudden decay.

Geometrically, the ameliorated flow behaviour is achieved through a slightly more convex camber line in the 20% to mid-chord region, as shown in Fig. 3.8a. This causes a stronger pressure gradient that increases the supersonic diffu-

sion of the flow, as shown in the contour plots of Fig. 3.9, reducing the pressure difference across the shock and thereby the entropy creation. In addition, the shock is spilled out of the passage, increasing the bow shock stand-off distance. This enables a pressure recovery in the region after the shock downstream to the LE. The effect is achieved in spite of the LE being sharper for the optimised design, as can be noted from the increased curvature shown in Fig. 3.8b, which produces a smoother expansion, mitigating the Mach number spike at the pressure side.

3.2.4 Adjoint-based Approach

The efficiency improvement at each iteration of the optimisation is shown in Fig. 3.10. It can be appreciated that within the first three iterations, the optimiser identified an optimal step size, which was employed throughout the optimisation. This suggests that the function's Hessian matrix –estimated with each gradient evaluation– does not significantly change for the design points assessed during the process. Thus, the SLSQP routine employed in this section is able to identify the dominant direction for each iteration.

Additional understanding of the process can be gained by analysing the evolution of the gradients through the optimisation. Each gradient evaluation provides information regarding the sensitivity of the cost function with respect to each of the parameters at a particular iteration. A gradient component with a large magnitude reveals a parameter that is more influential than the others, and dictates the primary geometrical modification of that iteration. This information is shown in Fig. 3.11, where each column encodes the efficiency gradient at a particular iteration. Throughout the process LE and TE recambering were the most influential parameters, with negative

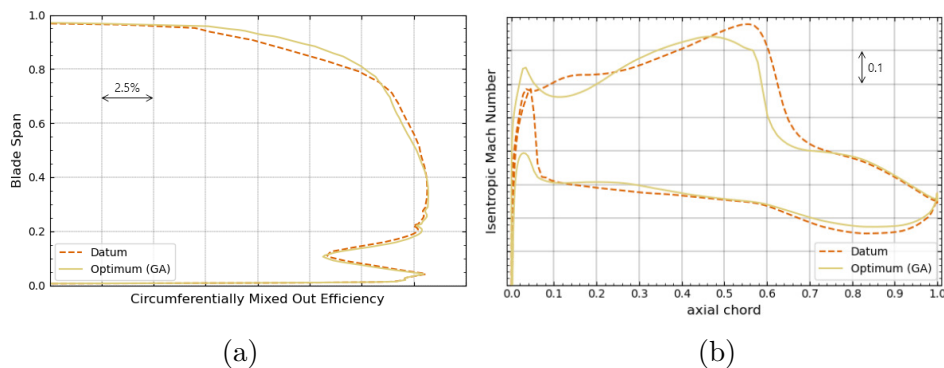


Figure 3.7: Radial efficiency profile for datum and optimised designs in (a); lift plot at 90% span in (b)

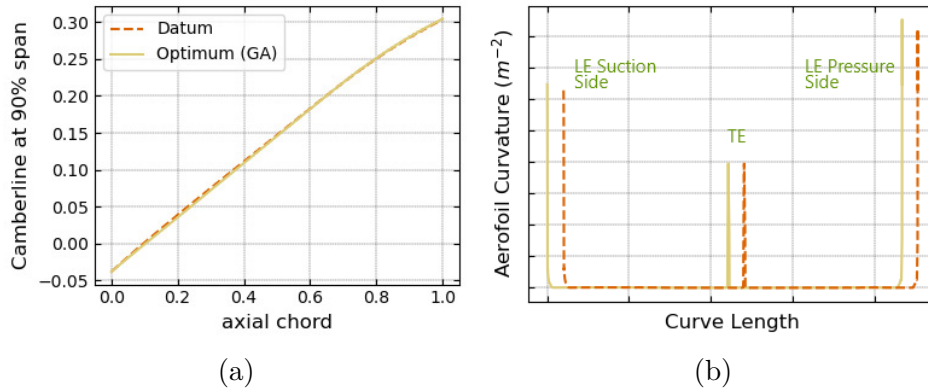


Figure 3.8: Aerofoil geometry at 90% span. (a): camberline distribution; (b): curvature vs arc length. The datum curve has been shifted for clarity




Figure 3.9: Static pressure contours at 90% span. (a): datum blade; (b): optimised design

gradient values near the root of the blade, and positive towards the tip.

Geometrically, this implies that the camberline curvature was reduced for the tip section, resulting in an s-shape which is more convex towards the pressure side, as shown in Fig. 3.12a, with a similar outcome to the global optimum design from Sec. 3.2.3. This mitigates the expansion along the suction surface that delays and weakens the shock, as shown in the lift plot in Fig. 3.12b. For the midspan regions the aerodynamic outcome was similar, but the geometrical mechanism was different. As shown in Fig. 3.13a, the blade inlet angle was increased, which causes a reduction in the effective flow incidence, thus reducing the loading at the LE. Additionally the camberline concave curvature was also increased, as shown in Fig. 3.12c, causing a smoother expansion which delayed the shock further downstream, reducing the shock-induced separated region. This ameliorated shock behaviour enabled increasing the flow turning through a reduction in the exit angle (towards more negative angles), shown in Fig. 3.13b, thus increasing the

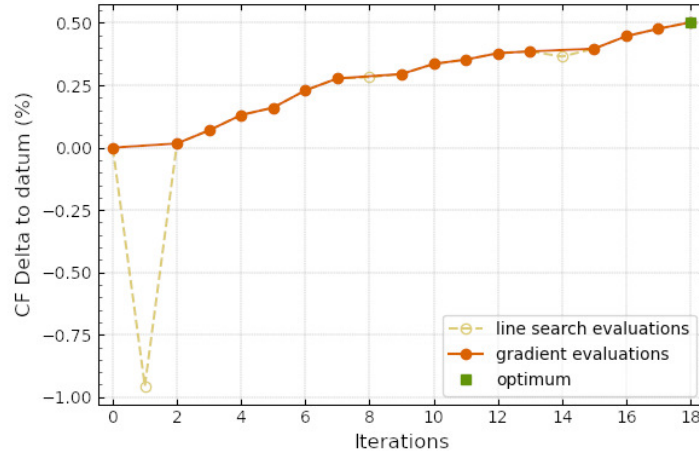


Figure 3.10: Adjoint-based optimisation history. The optimum found is highlighted in green

performance of the fan blade at these sections.

The gradient evolution in Fig. 3.11 reveals another interesting behaviour of the optimisation process. Primarily, that sweep, lean and skew had little effect on the objectives. Since the research blade is already high-performing, it is likely that the geometrical features controlled by these parameters were already at an optimum setting. Additionally, it can be noted that, after the initial design cycles where predominantly the recambering parameters were modified, the search direction became affected by the parameters controlling the locality of the recambering. This suggests that the optimiser was trying to fine-tune the camberline distribution, resulting in the final shape previously described, which was found to be tailored at mitigating the shock.

This work was centred on the blade design point. However, off-design performance is an important consideration when assessing fan blade designs. This information is presented in Fig. 3.14 for the optimised and datum blades. The characteristic curves show that stall margin is maintained for both GA and SQP optimums, while the choke margin has been slightly reduced. Additional objective or constraint functions could be introduced in the optimisation problem to maintain the datum off-design performance. Additionally, bulk-skewing the optimised blades could be considered to recover the choke margin [26].




Figure 3.11: Evolution of efficiency gradient through optimisation. The gradient has been normalised by l_2 norm.

3.2.5 Remarks on Computational Efficiency

Table 3.2 summarises the computational expense and improvement achieved by each optimisation method assessed in this study. The global optimisation of the research fan blade, through the novel strategy proposed, required 250 CFD samples. The incremental number of samples employed for the iterative loop from Fig. 2.4 was adopted to exploit parallel computing capabilities such that each iteration ran all samples in batch and had the run time requirement of a single CFD computation. For generating the initial dataset, 6 of such batches were required. Thus, the time requirement for generating the complete dataset was approximately 14 times the running time of a single CFD. The time requirement of fitting and optimising the neural networks is not considered since it was significantly lower than that of CFD runs.

In contrast, the adjoint-based optimisation required 19 iterations, of which 16 were B-matrix updates and the remaining 3 were for line search. As discussed, the update of the B-matrix requires a primal CFD with 4 times the cost of a normal run, and two adjoint calculations with 3 times the cost of the primal –hence, 12 times the cost of a normal run. This optimisation amounts to a total computational expense of approximately 387 CFD runs. In terms of the time requirement, only the adjoint calculation for the cost function and constraint can be parallelised, since the optimiser is sequential

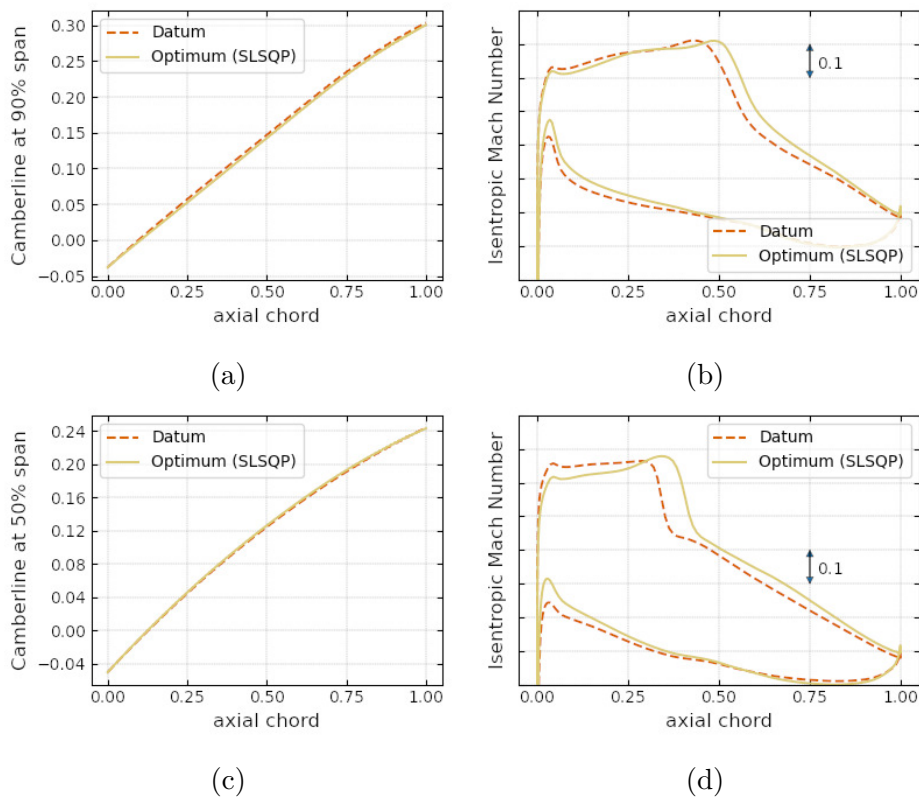


Figure 3.12: Camberline distribution and lift plots of datum and optimum designs at (a), (b): 90% span; and (c), (d): 50% span

and decisions are made based on the previous solution. This leads to a total time requirement of approximately 256 CFD runs.

The behaviour of the cost and constraint functions in the design space was such that both algorithms were able to converge to similar regions and produce optimised designs with comparable performance. The slightly lower benefit achieved with the global approach arises from the fact that, although highly accurate, the meta model built on the active subspace is a simplification of the true function shape. To overcome this, additional samples can be taken in the vicinity of the optimised design to refine the neural network in said region and improve the optimisation accuracy.

The global optimisation strategy achieved, with a reduced computational cost and runtime, an efficiency improvement comparable to a state-of-the-art adjoint-based approach. Additionally, the inherent implementation enables exploitation of parallel computing capabilities, which can significantly reduce the time requirement of the process. Moreover, the cost function and con-

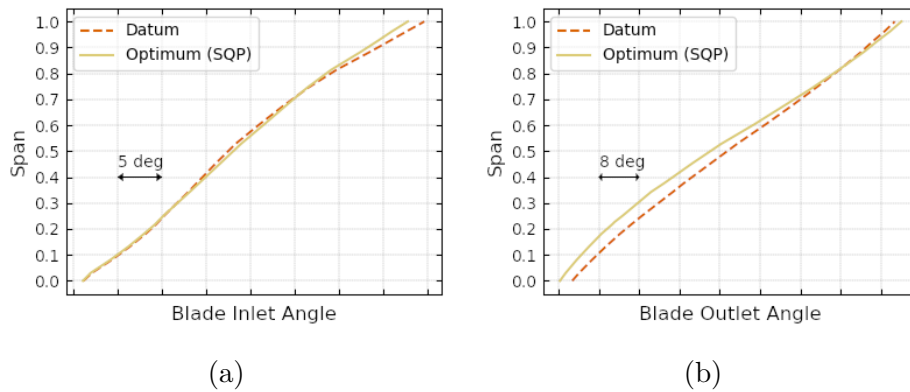


Figure 3.13: Blade metal angle distribution. 3.13a: Inlet angle, 3.13b: Exit angle

Table 3.2: Summary of computational efficiency for the optimisation approaches

	Global Strategy	Adjoint-based
Total Cost (CFD)	250	387
Total Run Time (CFD)	14	256
CF Improvement	0.47%	0.5%
Constraint Achieved	yes	yes

straint values can all be obtained with a single CFD run, thus providing better scalability of the global approach to multi-objective or multi-constrained optimisation frameworks. This could enable the introduction of additional constraints designed to maintain the off-design performance, or the specification of an exit total pressure radial profile via multiple constraints.

3.3 Conclusion

A novel global optimisation strategy has been developed that leverages the capabilities of ANNs for regressing complex functions while coupling them with ADS to reduce the number of samples required. This strategy was applied on the efficiency optimisation of a modern jet engine fan blade with constrained PR and compared, both in terms of overall improvement and computational expense, to an adjoint-based approach employing the same parametrisation. The global strategy achieved an efficiency increase com-




Figure 3.14: Fan characteristic curves for optimised designs

parable to the adjoint approach, with a reduced computational cost. In addition, it was demonstrated that adequate scalability to multi-objective or multi-constrained optimisation applications is achieved.

Chapter 4

Understanding Fan Stall Range through AInAS

Nomenclature

ADP	Aerodynamic Design Point
ADS	Active design subspace
BC	Boundary Condition
BOGV	Bypass Outlet Guide Vane
C	Function's gradient covariance matrix
CFD	Computational Fluid Dynamics
ESS	Engine Side Stator
FEA	Finite Elements Analysis
LASSO	Least absolute shrinkage and selection operator
LE	Leading edge
LS	Least squares
MAE	Mean absolute error
PR	Pressure ratio
QoI	Quantity of interest
RSM	Response surface mode
SA-H	Helicity-corrected Spalart-Allmaras turbulence model
SM	Stall margin
S-SA	Standard Spalart-Allmaras turbulence model
SV	Straightener Vane
TE	Trailing edge
TOC	Top of Climb condition
R^2	Coefficient of determination

W	Eigenvector matrix of C
Λ	Eigenvalue matrix of C
Φ	Flow capacity = $\dot{m}\sqrt{T_0}/p_0$
\dot{m}	Mass flow
T	Mass-averaged total temperature
p	Mass-averaged total pressure
C_μ	Tip leakage flow axial momentum flux
η	Isentropic efficiency
λ	LASSO regularization coefficient

4.1 Introduction

Rotating stall is an unstable regime of fan and compressor operation whereby the presence of non-axisymmetric disturbances prevents the uniform flow of mass through the machine, limiting the extent to which it can be safely operated. The stall margin of a fan or compressor defines the range of stable operation and is a major safety criterion that designers need to take into account for a safe operation of the jet engine and also for a successful certification.

Upon first appearance, the flow disturbances grow and propagate around the annulus and stall is said to occur when their effect is sufficient to promote flow breakdown. Two main paths into breakdown have been identified, distinguishable by their signatures in the pressure/velocity traces and commonly referred to as *modal oscillations* and *spikes* [92, 93]. Modal oscillations are employed to describe the occurrence of low-amplitude periodic axial velocity fluctuations with length scales the order of the compressor circumference. Such oscillations cause localized increases in blade incidence at certain parts of the annulus that can drive the loading beyond critical values and lead to the formation of stall cells. Modal flow oscillations can be detected in a machine many revolutions prior to stall and can develop smoothly into stall cells that span a large sector of the annulus. This sequence is commonly associated with conditions at, or to the left of, the peak of the total-to-static pressure characteristic [94, 95] and the initial stages of its evolution, where the amplitude remains small, can be well explained through linear models [96, 97].


However, modern compression systems employ highly optimized blades with reduced losses and blockage and typically do not experience a peak of the pressure-rise characteristic. In addition, design trends favour highly loaded blades, for which minor distortions in the incoming flow, or blade shape deviations, can drive the incidence beyond critical values when operating at a point on the the right side of the characteristic [98, 99]. On such

scenarios pre-stall waves are not observed [93]. Thus, the more frequent path into instability for modern fans and compressors is characterised by smaller scale (a few blade pitches) transients, called spikes, which appear as sharp oscillations in the pressure traces and propagate circumferentially at speeds about 70% to 80% of the rotor speed [100]. The amplitude of these distortions is larger compared to the mean flow velocity and they can grow into fully-formed stall cells within about three rotor revolutions.

Many research efforts have been dedicated to defining the characteristics of spikes and the fluid mechanisms that lead to their formation in core compressors [101–104] and fan blades [105–107]. The onset of these disturbances has been shown to be primarily linked to the tip aerodynamics, where flow separation from the leading edge (LE) begins due to high incidence. This condition needs only to be present at a single blade to initiate the unstable evolution towards rotating stall. The vorticity shed from the separated flow rolls up into a vortex tube that is bounded by the blade suction surface on one end and the casing on the other, and it propagates circumferentially [108]. The convecting vortex creates blockage in the passage, which increases incidence of the adjacent blade and thus constitutes a positive feedback loop for the propagation and growth of the instability.

Vo *et al.* [101], suggest two criteria that promote high incidence and the formation of a spike-type disturbance, both linked to over tip leakage flow. The first condition is spillage of the leakage jet ahead of the adjacent blade's LE and below the tip radius. As the flow coefficient decreases, so does the axial momentum of the incoming flow. The blade loading, in turn, rises, driving more fluid through the tip gap and consequentially increasing the momentum of the leakage jet. These conditions result in an upstream movement of the interface between incoming and leakage fluids. The flow behaviour in the tip region can be studied through entropy contours, as per Fig. 4.1a, which highlight the location of the interface through a sharp gradient when the low-entropy incoming flow mixes with the high-entropy leakage fluid; or through radial vorticity contours, as per Fig. 4.1b, which are useful to isolate the tip clearance jet. The first criteria indicates that stall onset occurs when the incoming/tip clearance flow interface is aligned with the rotor leading edge plane, or, equivalently, when the leakage jet becomes tangential. Any mass flow reductions from this point would tip the leakage jet over and cause spillage ahead of the leading edge. The occurrence of leakage fluid travelling upstream of the leading edge further increases the incidence on the tip, which promotes a detachment of the boundary layer, leading to corner separation and heavy blockage.

The second condition describes a situation where spike-initiated stall occurs before the interface aligns with the leading edge. In this case, the tip



(a) Contour of entropy for a near-stall flow field [101].

(b) Radial vorticity contour at near-stall [107].

Figure 4.1: Useful markers for analysing flow mechanisms present at stall.




Figure 4.2: Illustration leakage jet impingement at the adjacent blade's pressure surface.

leakage fluid from an adjacent blade passage impinges on the pressure surface of the adjacent blade at the trailing edge, which creates a backflow that induces corner separation, as illustrated on Fig. 4.2.

Pullan *et al.* [102], showed that while tip leakage spillage is a primary mechanism contributing to the formation of spikes, these disturbances can also occur in the absence of any leakage flow. The study detected spikes on rotors without tip clearance, suggesting a secondary mechanism that drives their formation. Shock-boundary layer interactions at the tip region cause separation downstream of the shock that create blockage. As the mass flow is reduced, blockage in the passage grows and reaches the casing, forming a corner separation and triggering the formation of a spike.

By analysing an array of compressor configurations with varying tip gaps, Hewkin-Smith *et al.* [104] found interactions between the two competing mechanisms and their correlation to the tip gap size. At low clearance values, casing corner separation is the leading cause of spike formation. However,

as the gap is opened up, leakage fluid energises the boundary layer and interrupts the passage blockage from reaching the casing, delaying stall onset. Further increases in the tip gap size, however, increase the axial momentum of the leakage flow and contributes to early spillage, initiating the formation of spikes. They found thus a non-zero tip clearance that is optimal in extending the stability range to lower mass flows.

Kim *et al.* [107], studied the different spike-inducing mechanisms on fan blades and found that the radial distribution of loading has an effect on the final size of the stall cell but not necessarily the path to spike formation. The study assessed two fan blades that operated at similar pressure ratios at design point, but presented different loading distributions, one approximately uniform and the second with a marked decay in the tip region. Their results show that the fan with uniform loading was more prone to spikes initiated by casing-corner separation, while the configuration with a reduced tip loading presented a stronger tip leakage flow that dominated the path to instability.

Significant efforts to improve the stability range of fans and compressors have been made, mostly through the application of casing treatments [109, 110] or tip injection [111, 112]. These methods have been proved successful in extending the range of safe operation, but their introduction is commonly associated with performance deficits. As a result, their practical application has been limited.

From a blade design perspective, the increased understanding of spike formation mechanisms has not been thoroughly exploited. The tip clearance remains considered as the primary lever that designers must consider for managing the over tip fluid. While it is indeed a critical parameter, its value also responds to other considerations, such as the desire to minimise the tip rubs experienced in service, the material and stiffness of the blade or manufacturing precision. Such other considerations currently limit the designer's control over stall margin and hence, the design point performance may have to be compromised to achieve sufficient stall range.

The axial momentum flux of the tip leakage fluid across the gap (C_μ) has been found to be a valuable metric to capture the effects of over tip flow on different designs [104, 107]. To enable more control over operability range, in this study, we present a systematic way to design for a target C_μ by considering full-span blade shaping whilst keeping the tip clearance and the design pressure ratio unchanged. By sampling a distribution of different C_μ values, the behavior of stall margin with respect to this quantity is derived, which reinforces the notion that significant control over the operability range (including the onset mechanism) can be achieved through blade design. Based on this concept, it is proposed that tip gap variations are purely a strong driver for C_μ , but that stall characteristics are defined by the latter metric.

Moreover, the effect of C_μ on the design point efficiency is also assessed and the trade-off between the stability range and efficiency is discussed in depth. A sensitivity analysis on the design parameters that control C_μ , design point efficiency and design pressure ratio is also performed, leading to an increased understanding of their exchange and which design levers to pull in order to achieve a better overall performance.

4.2 Computational Model

This study considers a low speed, high bypass ratio jet engine fan blade, representative of future composite geared turbo-fan, hereby referred to as the *research blade*, as the basis for analysis. This rotor is chosen because it embodies modern design trends resulting in high performance and flow range and it exhibits spike-initiated route to rotating stall. The span of the research blade is about one fourth of a conventional engine-size blade span, as it has been designed to be suitable for rig tests. The rotational speed has been adjusted to emulate the flow physics at the cruise condition and the blade thickness-to-chord ratio is representative of that necessary to maintain the structural integrity of the full scale engine blade.

An experimental campaign has been performed for this blade and the results were employed to construct a computational domain that would recreate as accurately as possible the test conditions. A description of the approach taken to maximise the fidelity to which the model simulates the real system is reported in this section.

4.2.1 Blade Geometry

A digital representation of the geometry tested on the rig was constructed by laser scanning the manufactured part and converting the resulting point cloud to CAD-file format through a novel inverse mapping process [113]. Such a geometry profile is representative of the blade as it was manufactured, in other words, cold and static. To obtain the shape of the blade as it is subjected to centrifugal and gas loads, representative of running conditions, a multi-physics iterative process involving a structural finite elements analysis (FEA) model and a computational fluid dynamics (CFD) model was employed for all speeds assessed in this study. A scheme of the process is presented in Fig. 4.3. Details of the CFD model are of special relevance to the current work and will be addressed at a following section. A description of the FEA model, however, is not provided for the sake of brevity. For every iteration in the process (other than the first one, for which no gas loads

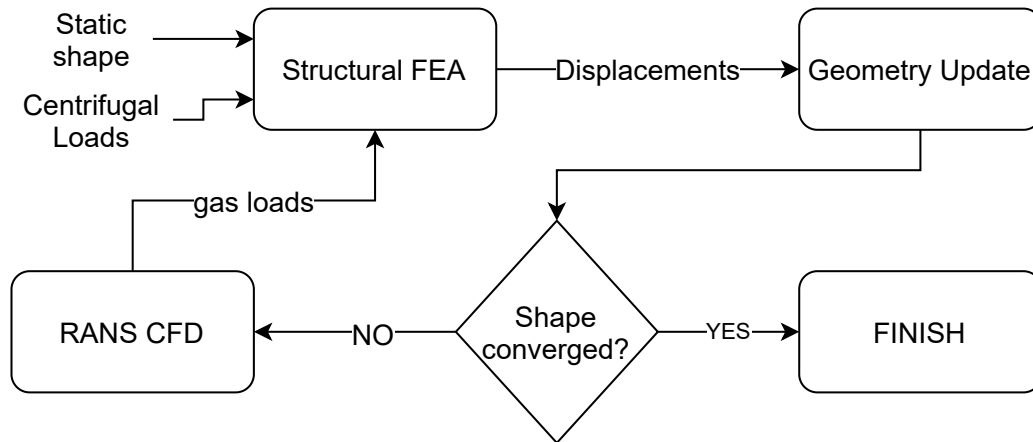


Figure 4.3: Running up process of cold static blade shape.

are applied), the FEA model is solved to find the shape displacements that result from the application of centrifugal and gas loads to the static shape. The approximating *hot running* geometry that stems from this is simulated with CFD to find the corresponding pressure distribution on the blade, information which is subsequently fed back to the FEA model for the next iteration. The L_2 norm of the distance field between successive geometries and the blade untwist is monitored to determine the convergence.

As it can be seen from Fig. 4.4, the process converges well and a criterion of $\text{distance} \leq 10\mu\text{m}$ was achieved in under 7 iterations for the speeds assessed in this study.

As discussed in the introduction, the size of the tip gap is a particularly critical parameter defining the stalling mechanism. The rig was configured with a tip clearance of approximately 1.7% chord, sufficient to guarantee the blade tip would not rub. To ensure the computational model captures the correct behavior, the clearance was measured on the rig and the digital domain applies the corresponding value.

4.2.2 CFD Domain

The computational domain used in this study is shown in Fig. 4.5. It consists of a single passage for every subdomain employing periodic boundaries. The model includes the rotor blade, bypass outlet guide vane (BOGV), engine side stator (ESS) and straightener vane (SV). The rotor subdomain is modeled on a rotating frame, with the casing, splitter, inlet, exit and a patch of the hub set as stationary surfaces. All blades, hub, splitter and casing surfaces are set as viscous walls. At the inlet, a non reflecting radial profile of total

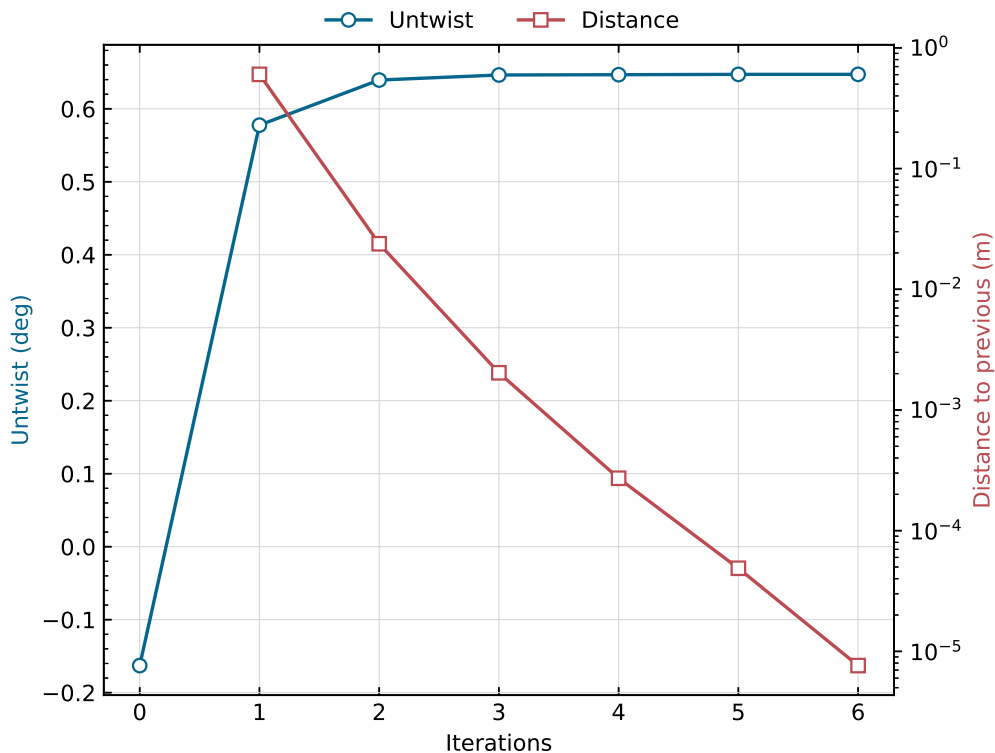


Figure 4.4: Convergence of the state to state process at a representative speed.

pressure, total temperature, whirl and pitch angles, and turbulence intensity is specified, where the values for these quantities were obtained from the rig test. For the bypass and core exit surfaces, non-reflecting, radial-equilibrium capacity (Φ) exit boundary conditions (BC) are enforced, while mixing planes are employed at the zone interfaces.

The CFD pre-processing makes use of the Rolls-Royce proprietary geometry and meshing software, PADRAM [91]. The structured mesh blocking strategy employed by PADRAM for each zone consists of an H-O-H topology for the blade passage, with H blocks for the upstream and downstream regions, as well as the upper and lower periodic boundaries. The blades are enveloped in an O-mesh while a C-mesh is used for the splitter. For the rotor tip, a butterfly topology is employed. The grid strategy was tailored for the aerodynamic design point (ADP) and employed for all subsequent operating points assessed, details of the rotor mesh are shown in Fig. 4.6.

The grid spacing on the solid walls is adjusted to produce y^+ values between 1 and 2 on the blade surfaces and lower than 5 on the rotor tip




Figure 4.5: Illustration of the CFD domain for the research blade. Image distorted and not to scale.

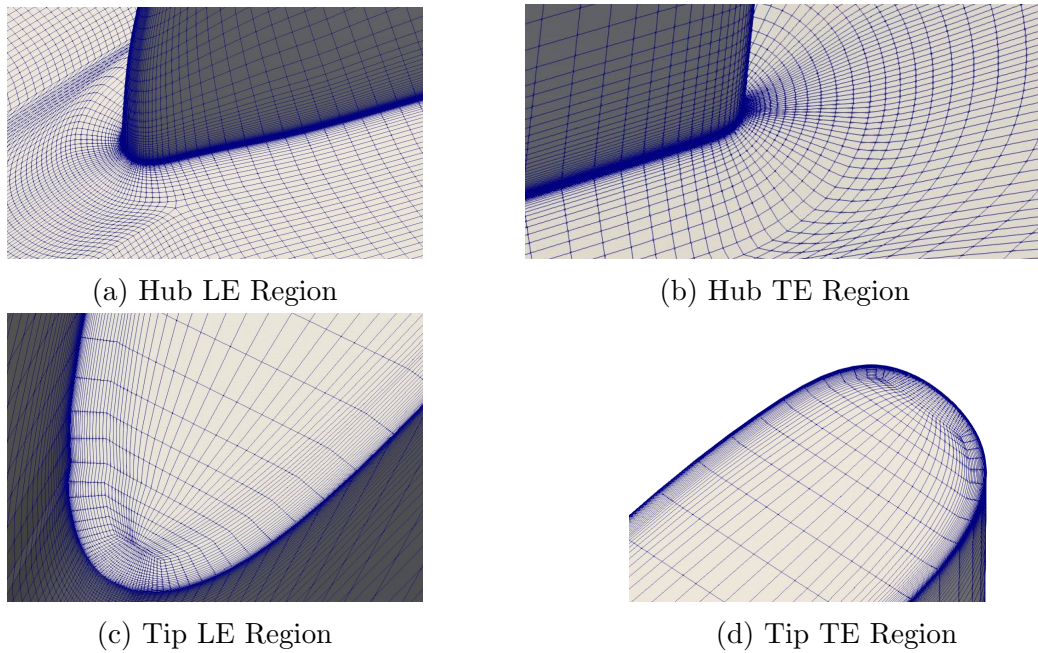


Figure 4.6: Rotor mesh details

and casing. The radial distributions of nodes in the rotor zone guarantees a proper grid density not only in the proximity of the hub and casing walls but also at the mid height span, being set such that 90 nodes were placed at the core exit and 160 at the bypass exit, with 40 of those being located

in the tip gap. The downstream stators, in turn, copy the node distribution from the rotor at the interfaces to avoid any radial mesh discontinuities at the mixing planes. Figure 4.7 shows the radial distribution of nodes at the interfaces between the rotor and downstream stators. The O-mesh applies 30 nodes around the blade surface and 30 additional nodes are placed in the H-mesh blade-to-blade direction. The number of axial nodes for each zone varies based on their actual distance from inlet to exit. The spacing criteria employed is such that the average axial spacing is about 2% axial chord at mid-span.

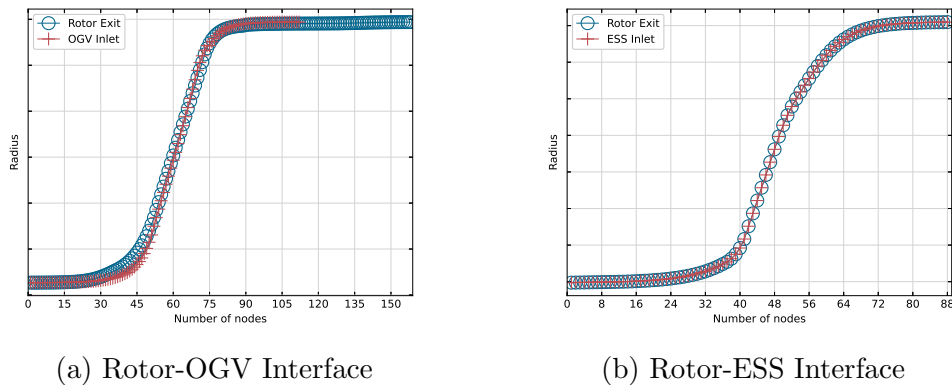


Figure 4.7: Radial distribution of nodes at mesh zone interfaces

4.2.3 CFD Solver

While fan stall is inherently an unsteady phenomenon, recent turbulence closure formulations enable performing steady-state numerical simulations until just before the onset of instability [114], leading to a cost-effective means to identify the stall point. This work employs the helicity-corrected Spalart-Allmaras turbulence closure model [115]. It was found that the scaling of the production term based on the local flow helicity promotes an increased resolution of the leakage flow and corner separation, stabilising the computations at lower mass flows and enabling more accurate predictions of the stall point. The improved accuracy in stall point definition was also observed on similar fan blade computations with this turbulence model [107]. Interestingly, in this work, it is also shown that it is possible to accurately predict the different spike formation mechanisms through steady-state simulations employing this turbulence model, providing an accurate characterisation of stall onset and minimising the need for time-accurate computations. Additional infor-




Figure 4.8: Schematic of the function implemented for stall identification via Fibonacci search method.

mation regarding this turbulence model can be found in the Appendix at the end of this chapter.

The model is implemented within the Rolls-Royce proprietary CFD code Hydra [90], which was used for all numerical computations in this study. Hydra is an unstructured solver employing 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 (RANS).

4.2.4 Identification of stall point

The fan operating point is controlled in the model through the bypass capacity BC. We define the numerical stall point as the minimum exit capacity for which the CFD solution reaches a steady convergence, as defined by inlet capacity, pressure ratio and bypass isentropic efficiency monitors. The emergence of small amplitude oscillations in the convergence history of any quantity was deemed sufficient to infer the presence of unsteady features and consider the operating point as the onset of instability. The simulations at those conditions are deemed as not-converged. Therefore, the computation of the stall point is performed at the last numerically stable operating condition, just before the onset of instability, and the stalling mass flows reported in this study are representative of the minimum value that would produce a

fully steady flow field, where oscillations in the aforementioned monitors are not observed.

To identify the stall point as efficiently as possible, a bracketing optimisation approach was employed following a Fibonacci search method [116]. To implement this, the rotor inlet capacity is modelled as a function of the bypass exit capacity BC value, $\Phi_{in} = f(\Phi_{BC})$. Approaching stall from the right, namely, for $\Phi_{BC} > \Phi_*$, where Φ_* is the BC value at exactly the stalled point, the function can be well represented by a linear model and its behaviour predicted with CFD using discrete function evaluations. At stall and for $\Phi_{BC} < \Phi_*$ the function experiences a discontinuity and CFD can no longer provide its value. Therefore, it is modelled as a constant with a penalty value $Penalty \gg f(\Phi_*)$, as in Fig. 4.8. This representation ensures that there is a minimum at $f(\Phi_*)$ and hence, any bracketing algorithm is capable of identifying it. The final size of the interval was selected to be 0.2% of the ADP capacity. Most stall point estimations converged within 5 to 7 CFD runs.

4.2.5 Model Validation

The aim of this work is to assess the behaviour of different fan designs in terms of their ADP performance and stall margin. Additionally, the top of climb (TOC) condition was also selected to monitor the stall margin as it has typically constituted a challenging condition both from an operability and a numerical modelling point of view. To provide confirmation that the model is capturing the correct physics with sufficient accuracy, bypass characteristics were run at two speeds, 95% and 103%, corresponding to the ADP and TOC speeds respectively, and were compared against experimental results. To throttle from choke to stall the bypass capacity BC was gradually reduced, whilst the core BC was kept at the cruise working line value.

The characteristic curves are presented in Fig. 4.9, showing a close agreement between the helicity-corrected Spalart-Allmaras (SA-H) computations and experimental results in pressure ratio and temperature rise ratio for both speeds assessed. The performance variation as the shaft is sped up from 95% to 103% is captured in the model with negligible deviations with respect to experimental data, denoting an adequate treatment of the gas and centrifugal loads on the blade shape. For the 95% curve, the calculations follow closely the rig results from choke down to a normalised mass flow of 0.8, indicating that the unstalled flow fields are replicated adequately. The stalling mass flow is over-estimated, with the CFD prediction being about 2.4% higher than the experimentally determined value. For the 103% speed, however, the stall margin prediction is closer and the model stalled within




Figure 4.9: Validation of performance characteristics for model against experimental data.

0.6% of the experimental stalling mass flow, while maintaining a close agreement at higher mass flows. The stall margin computations employing the standard Spalart-Allmaras model (S-SA) are also included in the figure, i.e., the last numerically stable point from the simulation with that model, highlighting the significantly increased accuracy obtained in the computation of this quantity when employing the helicity-corrected term.

The radial profiles of pressure ratio and temperature rise ratio, extracted at the BOGV LE, are shown in Fig. 4.10 for 95% speed at two conditions, ADP and the CFD stall point. Near the hub, there is a slight under prediction of both metrics at the conditions assessed, whereas the distribution

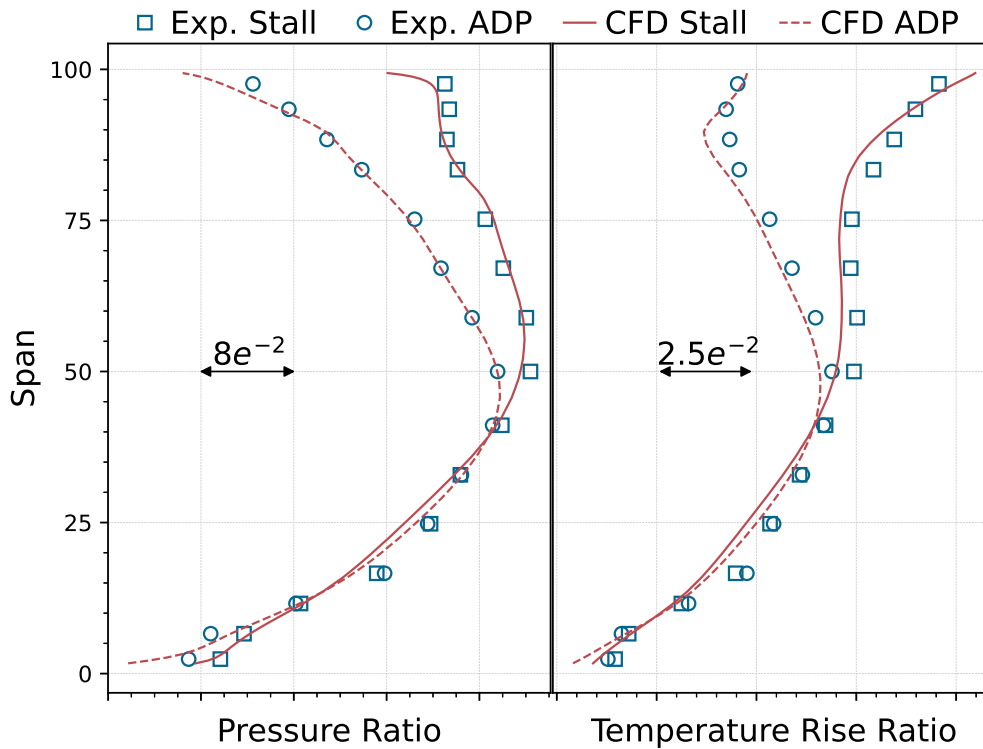


Figure 4.10: Radial profiles at the BOGV LE for 95% speed at the ADP and near stall conditions.

outboard of 25% highlights the close agreement between the CFD model and experimental data. The loading increase on the outer 50% span experimented as the mass flow is reduced from ADP to stall is captured successfully on the pressure and temperature profiles. This is relevant for the current investigation as near stall conditions the shock moves closer to the LE and promotes a thickening of the boundary layer downstream, which in turn, and due to a low axial momentum, tends to migrate outwards. The ability to adequately predict this radial migration of the flow has historically been associated with turbulence modelling. Thus, the validation of the profiles provides confidence that the turbulence model employed is producing adequate radial mixing and that the near stall three-dimensional flow field is captured with sufficient accuracy in steady state simulations.

4.3 Design Space Exploration

The hypothesis of this research work is that the stall margin and the stall onset mechanism are primarily driven by the axial momentum flux of the tip leakage jet across the gap. To test it, a systematic method to control this feature and characterise its behavior through three-dimensional design is sought. To this end, a flexible control over the geometry under study is essential, while at the same time it remains necessary to maintain adequate generalisation to other problems. While free-form deformation or B-splining-based parametrisations would provide increased geometry control, a subsequent analysis would be inherently biased and the location of nodes difficult to reproduce. Therefore, this work makes use of PADRAM's *Engineering Design Parameters*, which are composed of intuitive manipulation handles based on first principles. The degrees of freedom employed, shown in Fig. 4.11 for an aerofoil section, include *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*, as well as two additional parameters that control the *locality* of the recambering. The locality parameters control the chord-wise extent of the camberline alterations such that low values concentrate the changes on the LE/TE and high values propagate them throughout the aerofoil, providing full control over the camberline. While an important design parameter, the blade chord is kept unchanged in the current study to avoid inducing changes that might affect the weight of the system. 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 35 degrees of freedom. The value of the deformation applied as a function of the blade span is achieved through smooth cubic B-spline interpolation, with multiple control points via the control sections. The design parameters are gathered in the design vector, \mathbf{x} , and its dimensionality is denoted by the letter m , where $m = 35$. The lower and upper bounds that define the ranges of the design space are prescribed based on prior experience with the parametrisation, which has been used on previous optimisation studies providing sufficient geometry control [2, 117].

4.3.1 AI-Enabled Active Subspaces

To characterise the response of the system to the geometry changes introduced by this parametrisation and perform exploratory analyses, this work makes use of AI-Enabled Active Subspaces, as described previously in Chapter 2.

We wish to derive a relationship between the design parameters and

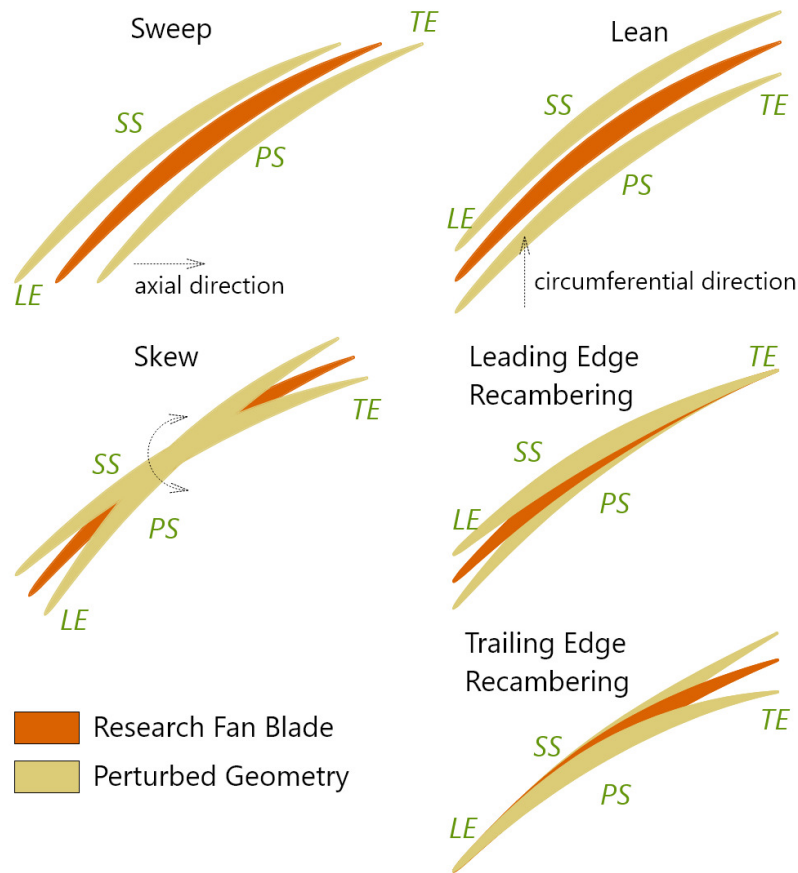


Figure 4.11: Degrees of freedom employed to manipulate the geometry of the fan blade [2].

the quantities of interest (QoIs). While the parametrisation selected comprises aerodynamically intuitive shapes, the number of degrees of freedom is nonetheless large and regressing the behaviour of any QoI that depends on them is challenging. Therefore, dimension reduction is desired to narrow the scope of the problem and identify a few *dominant* parameters which cause active changes in the QoI. Because of their ability to handle non-conditioned datasets, Active Design Subspaces (ADS) are selected to learn a low-dimensional representation for each QoI.

By performing the eigenvalue decomposition of the QoI's gradient covariance matrix, C , as $C = W\Lambda W^T$, the ADS approach identifies linear combinations of all the components of \mathbf{x} that best describe the variability in the function. Essentially it performs a rotation of the coordinate system such that the new directions explain more efficiently the observed variability.

The new basis is captured in the eigenvectors of the covariance matrix, W . The eigenvalues, on the other hand, denote the relative importance of each direction and allow partitioning Λ and W as per Eq. 4.1, such that the most dominant directions are captured in the *active subspace* W_1 . Thus, ADS first identifies directions where the function changes most actively and secondly, performs screening of these directions to keep only the k most dominant ones, achieving an efficient dimensionality reduction whilst maintaining all of the original degrees of freedom. The transformation expressed in Eq. 4.2, called *forward map*, is employed to map the high-dimensional inputs to their low-dimensional active representation, \mathbf{y} .

$$\Lambda = \begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix}, W = [W_1 \ W_2], W_1 \in \mathbb{R}^{m \times k} \quad (4.1)$$

$$\mathbf{y} = W_1^T \mathbf{x} \quad (4.2)$$

However, learning the ADS requires knowledge on the gradient of the QoI at a large number of locations in the design space, which has traditionally limited the applicability of the method due to the high cost associated with gradient computations. To mitigate this cost, the AI-enabled approach first employs zeroth order information to train a neural network regressor that learns a map from the high-dimensional input space to the QoI and exploits it to estimate the gradients using finite differences, which are subsequently used to construct the ADS covariance matrix. An iterative loop is employed to ensure the subspace is detected with the least amount of samples. Upon identification of the ADS, the forward map is performed for all samples and a secondary neural network is trained on the active directions. This leads to increased predictive accuracy as the ratio of samples to dimensions is increased and the complexity of the function is minimised through the use of \mathbf{y} . Therefore, the final network can be employed as a response surface to characterise the response of the system with respect to the design parameters through an enhanced resolution.

4.4 Exploratory Design Analysis

The QoIs considered for exploration are the fan total isentropic efficiency, η , the total pressure ratio, PR, and the axial momentum flux of the tip leakage flow normalised by the inlet axial momentum, C_μ . The axial momentum flux is calculated from Eq. 4.3, where V_{Xjet} is the mass-mean axial velocity of the leakage jet as it leaves the gap through the suction side, \dot{m}_{jet} is the mass flux through the tip gap, calculated at the suction surface, and \dot{m}_{inlet} and

V_{Xinlet} are the mass flow and mass-mean axial velocity at the passage inlet. This quantity was adopted to investigate the stability of the fan blade, since, as described in [104, 107], it correlates well with SM. Moreover, unlike SM, whose definition requires identification of the stall point (achieved through 5 to 7 CFD simulations), the value of C_μ is obtained at ADP. As a result, for each design assessed, all metrics were obtained by running a single CFD computation at the ADP.

$$C_\mu = \frac{\dot{m}_{jet} V_{Xjet}}{\dot{m}_{inlet} V_{Xinlet}} \quad (4.3)$$

To efficiently explore the design space and obtain the samples necessary to perform AI-Enabled ADS, a trust region strategy was followed based on maximising η , since this is the only QoI that with certainty we wish to maximise. The approach, shown schematically in Fig. 4.12, involved constructing sub-regions of the design space where the ranges of the parameters were reduced. By considering a reduced sector, the complexity of the function is inherently mitigated and an accurate regressor can be constructed with a lower number of samples. Initially the sub-region was centered on the datum blade design and spanned 25% of the design space in each direction. A quasi-random DOE based on Sobol' sequences was employed to draw 100 samples. The neural network predictor was constructed for efficiency and a first optimisation was performed to maximise it within the sub-region. Next, the center was moved to the location of the optimum and the process repeated. For each iteration, the samples from the previous runs were added and the ranges enlarged to cover the previous sub-regions, such that each step increased the overall sample count and the extension of the design space that was being explored. This strategy avoided placing too many samples on regions of the design space where the fan efficiency was low, while providing an adequate coverage of the remaining section. The optimisations for each sub-region step were performed with a standard elitist genetic algorithm and specified a tight double-bounded constraint on the PR to within $\pm 0.15\%$ of the datum value. The C_μ value was allowed to float freely. The process was stopped once the optimised design fell well within the trust region, i.e. away from its boundaries, conveying little gains from further exploration. Overall, three iterations were required to identify an optimum efficiency point, leading to 300 sampled designs which were employed to construct the final response surfaces.

The performance of the efficiency optimum is of relevance to the current investigation and its design will be addressed accordingly at a later section. For now, focus is set on analysing the response of the QoIs and deriving generic design rules to provide control over each.

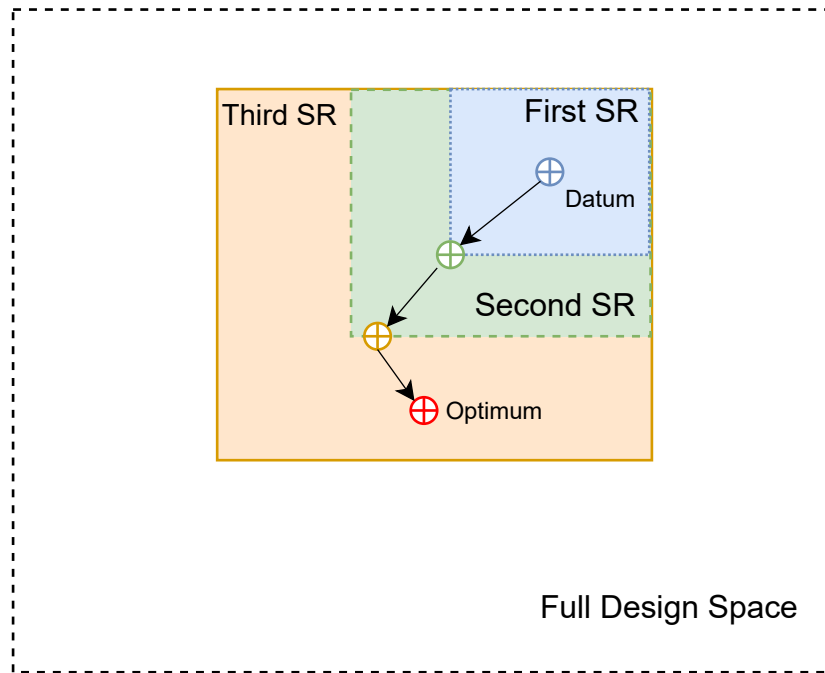


Figure 4.12: Sketch of the sub-region approach followed for exploration of the Efficiency Design Space.

As previously discussed, the ADS approach produces an eigenvalue matrix which can be used to identify the number of active dimensions to employ. Cumulative energy plots are a useful tool to convey how much of the energy is captured by the first k eigenvalues and thus visualise the potential for dimensionality reduction. It is worthy to highlight that, even though they may stem from the same high-dimensional space, the active subspaces for different functions are indeed different and, therefore, so are the number of dimensions required to express them.

Figure 4.13 shows the cumulative energy distribution for each of the metrics considered. For PR, it can be observed that with the first eigenvalue (or active direction) almost 100% of the energy is captured. This denotes that this metric has an approximately unidirectional behaviour in the active subspace and the 35 original parameters can be replaced by the first active direction without significant loss of information. C_μ in turn, is approximately two-dimensional, as suggested by the fact that the first two eigenvectors are required to capture the same amount of energy as a single one does for PR. Finally, η requires 6 directions to capture over 99% of the energy, implying this function's behaviour is more complex to regress and more dimensions are needed to explain the variability it experiences within the design space.

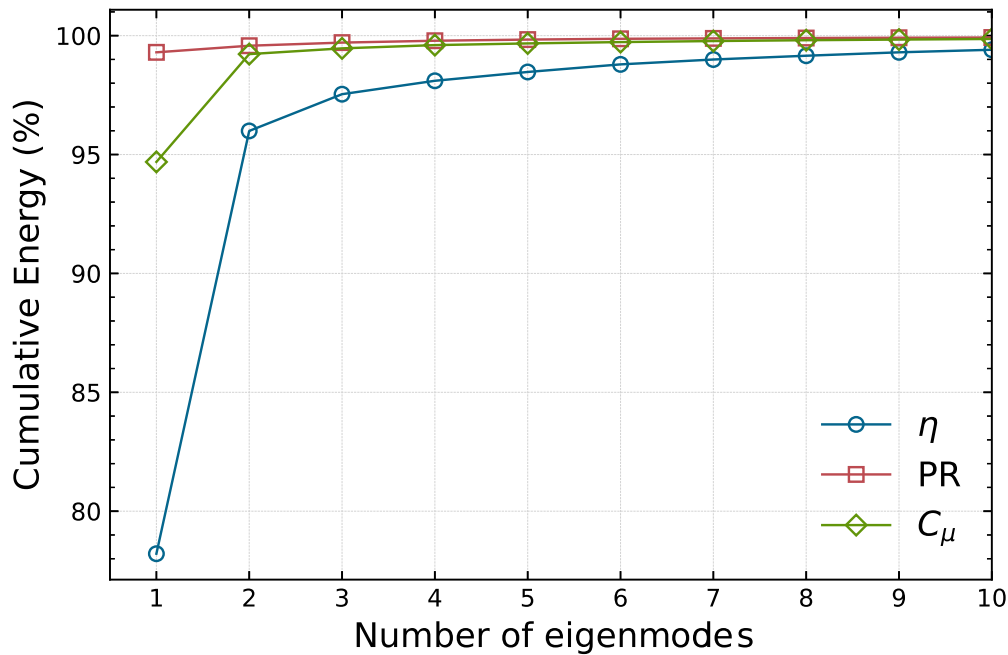





Figure 4.13: Cumulative energy plots for the covariance matrix of Efficiency, PR and Tip Leakage Axial Momentum Flux. Only the first 10 eigenmodes are shown.

As such, 1, 2 and 6 active dimensions are employed for PR, C_μ and η , respectively, to build their response surface model (RSM), thus providing a sufficient dimensionality reduction from the 35 original parameters.

Based on the decay of their respective eigenvalues, the number of active dimensions to use for each QoI was defined and the response surface models (RSMs) constructed. The predictive performance of each RSM, evaluated through cross-validation, is presented in Tab. 4.1, where R^2 stands for coefficient of determination and MAE is the mean-absolute-error. It can be observed that high R^2 values were obtained for all metrics, confirming that the responses of the models follow very close the ones of the real system and therefore any inferences made on them translates well. In addition, the low MAE values measured provide confidence that the true performance of any design is close to the value predicted with these RSM. It is worthy to highlight that generating the training data for the networks required solving a computationally expensive multi-row CFD analysis per sample. The methodology described thus enabled constructing accurate response surfaces with a limited cost. Due to a reduced number training samples, the neural




(a) Efficiency

(b) C_μ 

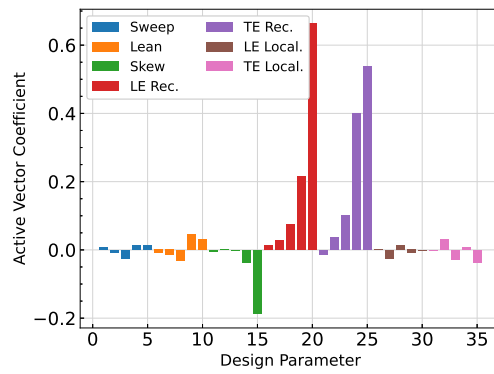

(c) PR

Figure 4.14: Performance maps showing the response of the QOIs in the design space.

networks employed are *shallow*, consisting of just two hidden layers. As such, their training can be completed without much expense. For the current task,



(a) Efficiency

(b) C_μ 

(c) PR

Figure 4.15: Coefficient activation for the first active direction of the QOIs. Higher magnitude coefficients highlight parameters with increased effect.

this was completed within 10 minutes on a 16-core CPU.

The fact that all QoIs present such low-dimensional structure, enables vi-

Table 4.1: Performance metrics for the response surfaces constructed through AI-Enabled ADS.

Metric	Active Dimensions	R^2	MAE
PR	1	0.998	$1.90e^{-4}$
C_μ	2	0.998	$1.77e^{-5}$
η	6	0.995	$9.77e^{-5}$

sualising their response through two-dimensional performance maps, shown in Fig. 4.14, without significant loss of information. The original samples used to construct the response surfaces have been included in the maps and their clustering responds to the movement of the sub-region center. In the efficiency map, to the left of the figure, the movement is particularly evident as this metric was employed to drive the exploration. In addition, the efficiency map has been constructed only in the explored region, as this sector of high-performance was discovered through the trust-region approach. For the other two metrics, the neural networks are employed to extrapolate the data to the corners of the design space, which is represented as a solid line enveloping the contours. The fact that they are low-dimensional and the trend is adequately captured with the ADS promotes accurate extrapolation.

It is worthy to note that the performance maps of Fig. 4.14 reveal that the change in η and C_μ experienced by the samples does not necessarily induce a variation on the PR, and this was intentionally enforced through the tight constraint specified on the trust-region exploration. The fact that the model employs a BC for the ADP bypass exit capacity implies that the mass flow is constrained as well. Therefore, by constraining the PR and enforcing its value to be effectively *unchanged*, the designs assessed are comparable as they provide the same thrust and share very similar operating points.

The efficiency map further shows that there is a corridor of high efficiency (towards the center-bottom) where the optimum lies. The exploration driven by the maximisation of efficiency has caused a movement across the C_μ subspace, particularly along the first dominant direction, that results in designs with highly variable magnitudes for this quantity. To conceptualise the changes that have caused the performance variation, the active directions can be further exploited to assess the effect that each design parameter has on the variability of the QoIs. The bar graphs in Fig. 4.15 show the activation coefficient of each parameter in the most dominant direction identified for each function. Recall that the dominant directions are a linear combination of all the original parameters, hence higher magnitude activations highlight parameters with increased effect. Each type of parameter has five bars as-

sociated with it, corresponding to the perturbation at the various span-wise 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.

A first glance through Fig. 4.15 reveals that the active vector coefficients for efficiency are almost the mirror-image of the coefficients for C_μ . This suggests that the two metrics are closely coupled and variations in one translate to variations in the other. This is due to the fact that the interactions of tip leakage jet (expressed by C_μ), with the incoming flow and passage shock are a primary source of losses and blockage [118], thus this feature has a strong effect on efficiency. A closer look at the coefficients for η and C_μ reveal that the primary design lever that drives their value is the tip LE recambering. This implies that the incidence angle and camberline distribution at the LE are responsible for much of the variations detected and have a strong effect on these quantities. The locality of the LE recambering was not detected as significant for either, implying that the camberline distribution is of primary significance only in a small region around LE and full chord control is not particularly necessary. For η , the LE recambering at 75% is the second driver, whilst this parameter does not cause a relatively strong effect on C_μ . In fact, parameters placed at spans other than 100% have a reduced effect on C_μ , highlighting the local nature of the tip leakage flow.

The analysis of the active coefficients reveal that close control over C_μ , and hence the leakage jet, can be achieved by setting focus on the tip aerofoil, while the design of the remainder of the blade can be driven by other objectives. Consequently, it is noted that the coefficients for LE and TE tip recambering are positive for C_μ and this implies that to move in the positive direction of the map (from left to right in Fig. 4.14), and thus *decrease* C_μ , positive values must be applied to the recambering parameters. With the parametrisation employed, this means *opening up* the tip and taking camber away from the LE and TE. Inversely, closing the tip airfoil and adding camber *increases* the axial momentum of the leakage jet. Some of these changes may appear counter-intuitive. To illustrate, by closing the blade down, the incidence angle at the ADP is reduced and thus the peak suction of the airfoil is mitigated, minimising lift. This effect can therefore be thought to reduce the mass flow of the leakage jet. However, the higher stagger indicates that the leakage jet is more axial, thereby increasing its axial momentum. In addition, increasing camber also increases lift, which drives more over tip flow.

For other design considerations, it is noted that the most active PR coefficients are mostly contained within the TE recambering type. This is known from first principles since the TE camberline distribution and exit angles

control the amount of turning. Blade skew at all spans is secondary in significance for PR while it is not so for the other metrics. Thus, for PR, design focus can be set on defining these parameters. Likewise, blade lean has a moderate effect on efficiency and a very small effect on the other metrics, hence it can be exploited without affecting the PR and leakage jet. Therefore, such parameters that affect mostly a single QoI can be exploited to tune the design and provide a higher degree of control over the desired metric without affecting the performance on other considerations. The following section proposes employing a regularised inverse-mapping methodology to achieve this tight control.

4.5 Effect of Tip Leakage Flow on Stability Range

The response surface models constructed based on the active vectors provide increased control over each QoI. This enables producing designs that achieve a desired value of η , PR or C_μ . In this section, an array of uniformly spaced designs is sampled from the C_μ subspace with varying values for this quantity, ranging from the minimum to the maximum possible values attainable with the chosen parametrisation. The location of these samples in the zonotope is shown in Fig. 4.14.

The forward map from Eq. 4.2 is used to map samples in the high-dimensional space to their corresponding location in the low-dimensional subspace. This is a well-posed problem yielding a single vector \mathbf{y} for a given vector \mathbf{x} . However, the inverse is not true. There are infinitely many \mathbf{x} s that solve the inverse map for a given \mathbf{y} . Therefore, to find the high-dimensional vectors, \mathbf{x}^* , that map to the chosen locations in the C_μ subspace, regularisation is needed. This work employs an inverse-map as expressed through Eq. 4.4.

$$\begin{aligned} \mathbf{x}^*(\mathbf{y}) = \underset{\mathbf{x}}{\operatorname{argmin}} \quad & \frac{1}{2} \left\| W_{1C_\mu}^T \mathbf{x} - \mathbf{y} \right\|_2^2 + \lambda \|\mathbf{x}\|_1 \\ \text{subject to} \quad & \left| \frac{PR(\mathbf{x}) - PR_{datum}}{PR_{datum}} \right| \leq 1.5e^{-3} \end{aligned} \quad (4.4)$$

The first term of the minimisation cost function in Eq. 4.4 is designed to identify points that, when forward-mapped, fall on the chosen location \mathbf{y} . The introduction of the constraint, to produce designs where the PR is within $\pm 0.15\%$ of the datum value, leads to the possibility that there might not be feasible designs that forward map to \mathbf{y} . Therefore, the cost

function will produce the feasible point that maps *closest* to \mathbf{y} as per the L_2 norm. The cost function of Eq. 4.4 employs also a LASSO regularisation term [119]. This term biases the selection to designs where the least important parameters in the subspace (the ones that do not affect the value of the first term) are pushed to zero. By changing the value of the coefficient λ , this shrinkage effect is strengthened or weakened, thus providing control over the final design. This study assessed an array of λ values, ranging from 10^{-1} to 10^{-4} .

For each sample, the inverse map was solved and the resulting design was simulated with CFD at the ADP condition and at the stall point. The Stall Margin (SM) was subsequently calculated as per Eq. 4.5.

$$SM = \frac{\Phi_{ADP} - \Phi_{stall}}{\Phi_{ADP}} * 100 \quad (4.5)$$

Figure 4.16 shows the distribution of stall margin as a function of C_μ for the samples assessed. It can be observed that C_μ is a very strong driver for SM and that it has a quadratic response, as indicated by a high coefficient of determination (0.96) for a least-squares (LS) second order polynomial fit to the data. Three regions of distinctive response are noted and representative designs indicated by the numbers 1, 2 and 3 in the figure. At low C_μ values, corresponding to a mitigated leakage jet (point 1), the SM is initially low. Shock-boundary layer interactions take precedence and, due to its low energy, the boundary layer separates downstream of the shock, causing blockage that extends well into the passage, as it can be visualised by the radial vorticity contours of Fig. 4.17a, which are taken at 95% span. While the leakage jet cannot be entirely eliminated (for non-zero tip gaps), its effect can be greatly mitigated causing the spike inducing mechanism to be separation-dominated. As the C_μ is increased beyond this region, the leakage jet is enhanced and energises the boundary layer, suppressing the separation and alleviating the blockage, thus contributing to an increase in SM. The mechanism switches from being primarily separation-dominated to being leakage-jet dominated, although there is still some separation-induced blockage, as it can be seen from Fig. 4.17b, which corresponds to point 2 in Fig. 4.16. This is consistent with stall onset at the optimal tip gap value identified in [104]. Further increases in the leakage jet strength promote this feature to become dominant and impinge on the adjacent blade's pressure side, promoting flow reversal and early spillage, as shown in Fig. 4.17c for point 3.

The ideal Mach number distributions, shown in Fig. 4.18, further clarify how the design parameters are affecting the behaviour of the flow in these distinctive regions. Going from design 1 to design 2, the camber has been increased at the LE, strengthening the suction spike, which in turn drives

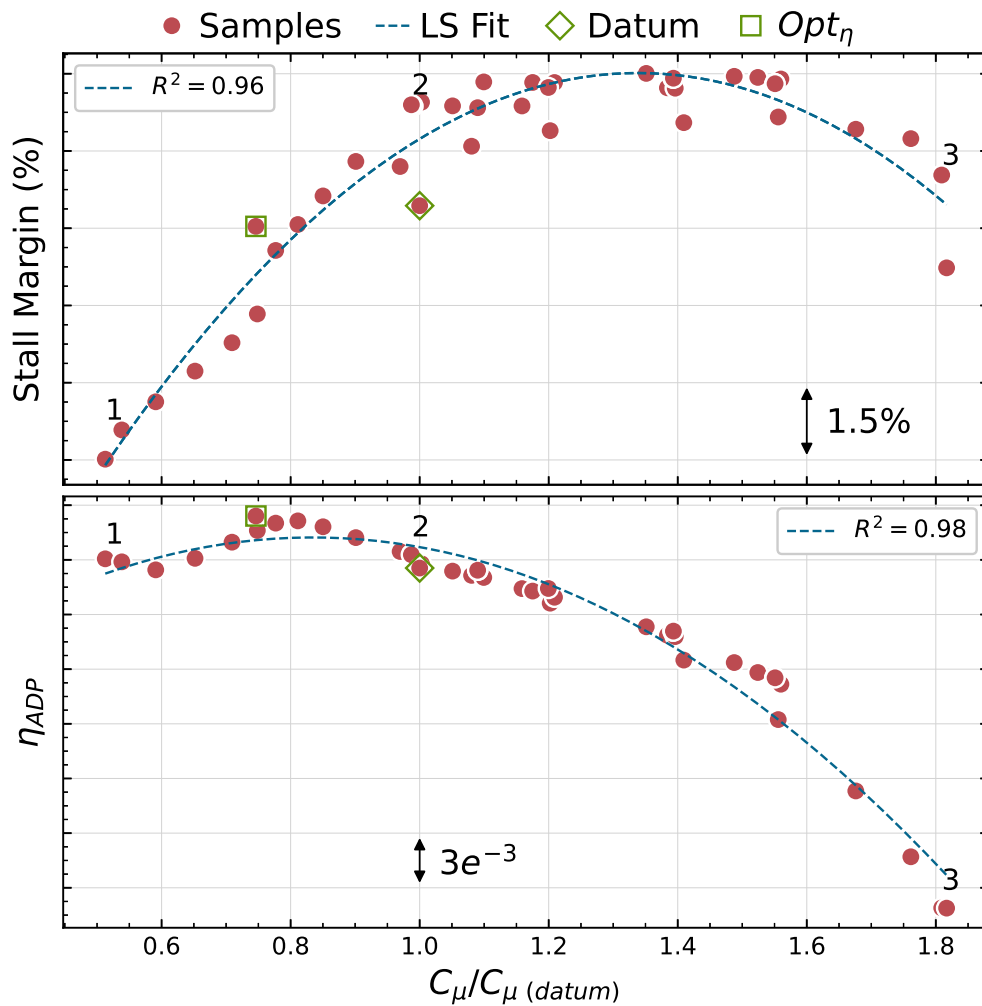




Figure 4.16: Stall margin and ADP efficiency versus tip leakage axial momentum flux, normalised by datum value. The QOIs present a quadratic response to changes in axial momentum.


more flow over the tip. In addition, the blade has been closed down, resulting in an overall lower incidence that reduces the pre-shock Mach number. The combination of a weaker shock and a more energetic boundary layer, produces a more uniform loading and increases lift. Additional mass flow through the gap, driven by local LE recambering, and a larger axial jet resulting from closing the blade lead to design 3, which promotes the impingement of the adjacent blade's leakage jet on the pressure surface, causing a significant loss of pressure at the LE and greatly accentuating the pressure side spike. The lower incidence delays the peak suction point and the shock gets sucked



(a) Design 1



(b) Design 2



(c) Design 3

Figure 4.17: Contours of radial vorticity at stall for the three representative designs at 95% span. The enhancement of the leakage jet initially suppresses the corner separation leading to increased stall range. Further increases in its strength cause impingement on the adjacent blade's pressure surface and early spillage.




Figure 4.18: Isentropic Mach number distributions at 95% span.

into the passage, developing a double-shock structure which is detrimental to stability and performance.

The effect of C_μ on ADP efficiency is also shown in Fig. 4.16. As previously mentioned, these two metrics are tightly coupled and this figure confirms this. The effect of C_μ is again quadratic with a coefficient of determination of 0.98. The region of low jet momentum is associated with intermediate efficiency values, as the shock-induced separation causes loss. The suppression of the separation promotes an efficiency increase, while designs with overly strengthened leakage jet perform the poorest, due to the detrimental effects induced by the mixing of the leakage jet with the incoming flow, as well as passage shock-jet interactions. Moreover, without having actively controlled the leakage jet during the process, the efficiency optimisation has produced a design where the C_μ was naturally placed in the region of high-performance, which corresponds to normalised C_μ values of around 0.75, as it can be seen from Fig. 4.16. While this design is high-performing at the ADP, its SM is slightly lower than the datum and indeed far from optimal in that sense. There is thus a trade-off between efficiency and stability range, which can be mostly driven by controlling the axial momentum of the leakage flow.

The η -optimised design achieves an efficiency improvement of 0.3% and a reduction in SM of 0.4%. Figure 4.19 presents a top view of the datum and the efficiency-optimised design. It can be appreciated that, predominantly, the optimised design features a lean of the tip towards the pressure side and the LE has been opened-up. As such, the optimiser has focused on mitigat-

ing the leakage jet (through LE recambering) and has further enhanced the performance by applying a parameter which did not affect this flow feature (lean).

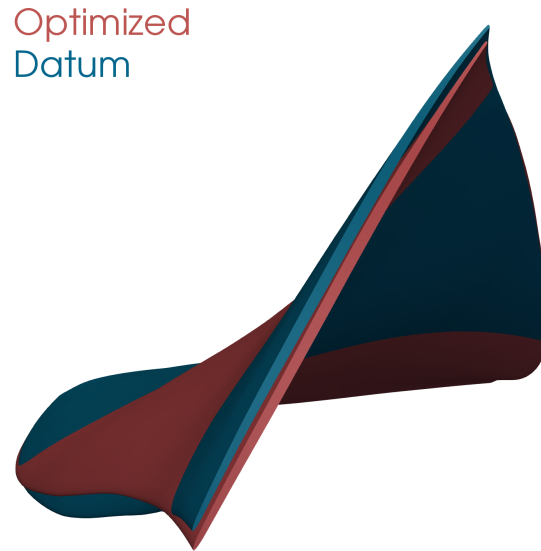


Figure 4.19: Top view of the datum and efficiency-optimised designs.

It can be noticed from Fig. 4.16 that, while slightly lower than the datum's value, the optimised design's SM is higher than other points that have the same C_μ . In fact, the figure reveals a number of designs that have the same C_μ but different SM values. This suggests that C_μ is a strong driver for SM, but there are other parameters that also cause a significant effect. The ADP efficiency, on the other hand, remains approximately constant for constant C_μ values. As such, there are designs (such as point 2 in Fig. 4.16) that have efficiency values very close to the datum's, but a significantly increased stability. The different designs at constant axial momentum values were constructed by employing varying levels of the regularisation coefficient in Eq. 4.4, and thus incorporate different design considerations. By analysing the different samples that yield higher SM, further design principles can be derived.

Focus is now set on the region of normalised C_μ between 1.0 and 1.6, where the additional designs have been constructed. It was found that the main difference between these was the tip loading, measured as the average pressure ratio on the upper 5% span, and the level of lean applied to the tip aerofoil. Figure 4.20 shows the correlation of these two design considerations with SM. It can be noticed that tip lean presents the higher correlation,




Figure 4.20: Secondary design drivers for stall margin. For a given axial momentum, off-loading the tip and applying lean towards the pressure side produce increased stability.

and increasing levels promote an increase in SM. Tip lean is calculated as a circumferential rotation of the tip section (Fig. 4.11), where positive values are measured in the direction of rotation of the blade. Therefore, the correlation reveals that leaning the tip aerofoil towards the pressure side leads to increased stability for given C_μ values. This is observed for the efficiency-optimised design as well (Fig. 4.19). With respect to the tip loading, the correlation with SM is not as strong but it remains considerable. The results suggest that off-loading the tip generally improves the stability margin for constant C_μ values.

4.5.1 Results at Top of Climb

The previous analysis has been performed at 95% speed, which corresponds to the ADP speed. However, as already mentioned, the TOC is a key condition to assess the operability range. Therefore, this condition was simulated for a

few designs from Fig. 4.16 ranging from low to high values of C_μ . It is worth noting that the derivation of these designs was performed for 95% speed and therefore the ranges of C_μ quoted in this section are not necessarily the minimum and maximum possible values that can be attained at this speed. Moreover, the blade shapes at 103% have been obtained by employing the *running-up* process described in Sec. 4.2.1.




Figure 4.21: Variation of fan performance versus axial momentum flux as the speed is increased.

Figure 4.21 shows the variation in SM and efficiency with respect to C_μ for some designs as they are sped up from 95% to 103% speed. To make the results at the two speeds comparable, the data has been normalised by

the datum value at the corresponding speed. This figure reveals that the leakage jet still promotes a quadratic response on both metrics at increased speeds. However, the variations at 103% are much more pronounced and the trends fall-over more rapidly at high C_μ values than they do at the ADP speed, suggesting the system is less robust to variable C_μ at high speeds. Moreover, for SM, the data indicates that operating on normalised C_μ values that provide optimal stability at ADP actually leads to operating at the condition “3”, identified previously, at TOC speed, with a poor stability performance. There is thus a trade-off in the tip leakage strength for optimal stability at different speeds. For efficiency, on the other hand, the optimal axial momentum for maximum performance is consistent at the two speeds, indicating there is no real trade-off in this metric as the fan is sped-up with respect to the axial momentum.

4.6 Conclusions

This study has focused on the effect of tip leakage axial momentum flux on design point efficiency and stability range for an axial fan representative of future composite geared turbo-fan engines. The analysis has been driven by the use of a high-fidelity computational model which has been validated against experimental data. The simulations performed were steady-state and involved the use of the helicity-corrected Spalart-Allmaras turbulence model. It has been shown that, through this model, sufficient resolution can be obtained for the calculation of the stalling mass flow and the onset mechanism, minimising the need for time-accurate computations.

Furthermore, a design and analysis methodology has been employed, which enabled building highly accurate meta-models for efficiency, pressure ratio and stall margin with a reduced computational cost. Through its application, the characterisation of the system and the interactions between the quantities of interest was achieved using only 300 samples for 35-dimensional input parameters.

A quadratic response has been detected for both efficiency and stability with respect to the tip leakage jet’s axial momentum. Three regions of distinctive behaviour are observed. At low axial momentum values, the leakage jet is mitigated and shock-boundary layer interactions dominate the exchange. The stall onset mechanism is driven by separation-induced blockage and the efficiency is intermediate, responding to losses originated through the separation. As the leakage jet is strengthened, however, it initially suppresses the separation and leads to an increase in both stability and efficiency. The stall onset mechanism switches to being primarily leakage jet-dominated.

Nonetheless, further strengthening of the leakage flow promote an increased effect of this feature which becomes detrimental to stability, by impinging on the adjacent's blade pressure side and early spillage ahead of the leading edge, and efficiency, by inducing increased mixing losses. The optimal setting for efficiency differs from the one for optimal stability but the trade-off can be exploited through controlling the axial momentum. This trend has been verified for higher speeds corresponding to the engine's top of climb condition. However, it was revealed that the optimal leakage momentum for maximum stability varies as the shaft speed is increased, and there is also a trade-off for the stability margin at different speeds, indicating compromises are required to balance ADP and TOC stability.

Blade design rules for increased over-tip flow control have been derived which promote local tuning of the tip aerofoil while the remainder of the blade can be designed for a target pressure ratio. It has been shown that opening the tip and taking camber away from the leading edge are strong levers to mitigate the leakage jet. Additional stability enhancements beyond optimal axial momentum have been shown to be attainable through off-loading the blade tip and leaning the tip towards the pressure side.

Appendix

The helicity-corrected Spalart-Allmaras model was first introduced by Liu *et al.* [115] in 2011. Different flavours of the correction have been proposed in the literature, but the implementation employed in this work stems from the original definition. The source term in the S-SA model is augmented as per Eq. 4.6, where the helicity is computed from Eq. 4.7. The velocity, \mathbf{u} , and vorticity, $\boldsymbol{\omega}$, in the equations are measured in the absolute frame of reference. The value of the constants is taken from [115].

$$\tilde{S}_{SA-H} = (1 + C_{h_1} h^{C_{h_2}}) \omega + \frac{\tilde{\nu}}{\kappa^2 d^2} f_{\nu_2} \quad (4.6)$$

$$h = \frac{|\mathbf{u} \cdot \boldsymbol{\omega}|}{\|\mathbf{u}\| \|\boldsymbol{\omega}\|} \quad (4.7)$$

Chapter 5

Robust Design Optimisation

Nomenclature

\mathbf{x}	design vector
χ	design space
ω	uncertain random variables
Ω	event space
PCE	Polynomial Chaos Expansion
RDO	Robust Design Optimisation
ARMOGA	Adaptive Range Multi-Objective Genetic Algorithm
BC	Boundary Conditions
DOF	Degree of Freedom
PR	Pressure Ratio
η	Fan isentropic efficiency
μ	Mean efficiency
σ	Standard deviation of efficiency
DGO	Deterministic Global Optimum

5.1 Introduction

The performance of engineering components is inherently variable. Performance metrics therefore drift, not only as a result of the conditions under which the components currently operate, but most importantly, due to their their previous history. Manufacturing variability [120, 121], hysteresis [122, 123] and erosion [124, 125] introduce permanent changes in the as-designed machine that can prevent it from operating optimally.

In fact, the as-designed machine lives only digitally, as the real world is stochastic and imperfections are inevitable. Safety factors and design tolerances are introduced to set bounds to the allowable variability and it is commonly the case that the less information there is available about the system, the tighter these can be, which has significant implications to the cost of manufacturing and maintenance of the machine. The treatment of uncertainties therefore plays a significant role in the success (or lack thereof) of any engineering application and their modelling at the design stage has been the topic of many studies [126–130].

In chapter 4, significant prior information about the system and expensive processes were employed to closely replicate a real manufactured fan blade, which led to a very accurate computational model that captured the response of the physical system to a high degree of certainty. Following, a *deterministic* optimisation approach was performed whereby the quantities of interest were improved and their response characterised. However, in the absence of such vast information, deterministic approaches can be misleading, as the effects of uncertainties are ignored in the model. Robust Design Optimisation (RDO) aims to incorporate the effect of input variability in the quantities of interest during the design process.

Let $f = f(\mathbf{x}, \omega)$ represent the response of a quantity of interest for a design parameterised by the design vector $\mathbf{x} \in \chi \subseteq \mathbb{R}^m$, and subjected to uncertain random variables $\omega \in \Omega \subseteq \mathbb{R}^k$. Let \mathbf{x}_1 , \mathbf{x}_2 and \mathbf{x}_3 be three arbitrary design points in χ . The response of f can be thought to follow three different random variables in the event space Ω :

$$\begin{aligned} f(\mathbf{x}_1, \omega) &= F_1(\omega) \\ f(\mathbf{x}_2, \omega) &= F_2(\omega) \\ f(\mathbf{x}_3, \omega) &= F_3(\omega) \end{aligned} \tag{5.1}$$

The distributions of such functions are illustrated in Fig. 5.1. Assuming the convention that optimum designs evaluate to the minimum function value, in a deterministic setting, \mathbf{x}_1 would dominate the other designs, since it can be noted that F_1 has the lowest central tendency. This conclusion illustrates the perils of deterministic optimisation, as F_1 also experiences the largest variability introduced by the uncertain parameters. Consequentially, the true value for $f(\mathbf{x}_1, \omega)$ measured in practise will often be higher than the response observed for the other design vectors considered. Likely, if only sensitivity to uncertainty were to be considered, \mathbf{x}_3 would be adopted, even though the response promoted is highest. It follows that RDO seeks the best trade-off between mean performance and sensitivity to uncertainty.




Figure 5.1: Illustration of probability distribution for the response of the quantities of interest in the event space Ω .

5.2 Uncertainty Quantification

The formulation and resolution of a framework for design under uncertainty poses many challenges. In the previous illustration, the probability distributions for the response to uncertain parameters at the different design points was known. However, in practice, estimating the distributions is a major field on its own, called *Uncertainty Quantification*. A brief description is provided here, the interested reader is referred to [131] and references therein.

In applications where function evaluations are inexpensive, the propagation of uncertainties can be performed by random sampling in the event space, such as by Monte Carlo simulations [132], which yield the probability distributions for the quantities of interest. However, this work deals mostly with high-fidelity models for which such approaches are unfeasible. In such cases, surrogate modelling is employed to predict the performance of the functions. Polynomial Chaos Expansion (PCE) [133] has become the dominant approach for estimating the statistical moments. Following an approach analogous to Eq. 2.1, PCE aims to construct an analytical model as per (5.2), where ω is a k -dimensional vector of independent random variables, and $\{\psi_j^{(i)}\}_{j=0}^{\infty}$ is a family of polynomials L^2 -orthogonal.

$$F(\omega) \approx \sum_{j=1}^n k_j \psi_j(\omega) \quad (5.2)$$

The inherent use of orthogonal polynomials make PCE regression easily integratable through quadrature rules, and the evaluation of the statistical moments is straightforward through the expansion's coefficients. However, PCE is particularly sensitive to the curse of dimensionality (Sec. 1.1), and computing the coefficients of the expansion requires significant computational effort.

To mitigate this effect, first order information of the functions can be incorporated into the PCE regression, greatly reducing the sampling requirements. The traditional approach is to employ a least-squares formulation to derive the coefficients for the PC expansion by stacking the function and gradient evaluation matrices. However, as detailed in Ghisu *et al.* [117], in the presence of noisy gradient information, this approach is prone to enlarge numerical inaccuracies, leading to poor estimation of the statistics. Instead, [117] shows that the inaccuracies encountered with the stacked approach can be greatly mitigated through constrained least-squares approximations using a linear inequality approach via a null-space method. Such an approach 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 optimising under uncertainty have typically required designers to employ models of reduced complexity (e.g. 2D) to lessen the expense of functional evaluations [134], or consider a reduced number of uncertain parameters to limit the overall number of functional calls required [135, 136], the null-space approach provides greater flexibility in both these considerations through full exploitation of gradient information. In this section, this approach is combined with AInADS to perform the Robust Design Optimisation (RDO) of an industrially relevant problem. Moreover, to demonstrate the differences between deterministic and robust approaches, the case considered is the same as in Sec. 2.1.

5.3 Optimisation Framework

For a function $f = f(\mathbf{x}, \omega)$ representing the response of a quantity of interest for a design parameterised 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 optimisation problem as per Eq. 5.3, where $\mu(\mathbf{x}, \omega)$ represents the mean value and $\sigma(\mathbf{x}, \omega)$, the standard deviation of f , respectively.

$$\begin{aligned}
& \underset{\mathbf{x}}{\text{minimise}} \quad \{\mu(\mathbf{x}, \omega), \sigma(\mathbf{x}, \omega)\} \\
& \mathbf{x} \in \chi \subseteq \mathbb{R}^m \\
& \omega \in \Omega \subseteq \mathbb{R}^k
\end{aligned} \tag{5.3}$$

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 (5.3) can be obtained by a least-squares PCE augmented with gradient evaluations through the null-space method as per [117], using p quadrature sample points of f and \mathbf{g} .

Equation 5.3 involves two design considerations which may present conflicting optimality criteria. While the optimisation of μ is aimed at identifying designs with statistically improved values of f , optimising for σ identifies designs where the variability encountered in f due to the uncertain parameters is minimised. 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 (5.3) 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) optimiser, present in the optimisation library SOFT [137], 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 [138]. In [138], the approach was applied to the optimisation of the nominal performance of the Vital fan, while here it is extended to a robust optimisation problem, using the same strategy to regress the optimisation 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 [138].

5.4 Computational Tools

The function and gradient evaluations in this section were performed with the Rolls-Royce proprietary CFD suite Hydra [90]. 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 used here is equivalent to the one used in Chapter 3, shown in Figure 3.2, and it 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 [91], 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 [26], 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 has previously been reported [26, 27], achieving a good match for a number of different operating conditions.

5.4.1 Geometry Parametrisation

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.

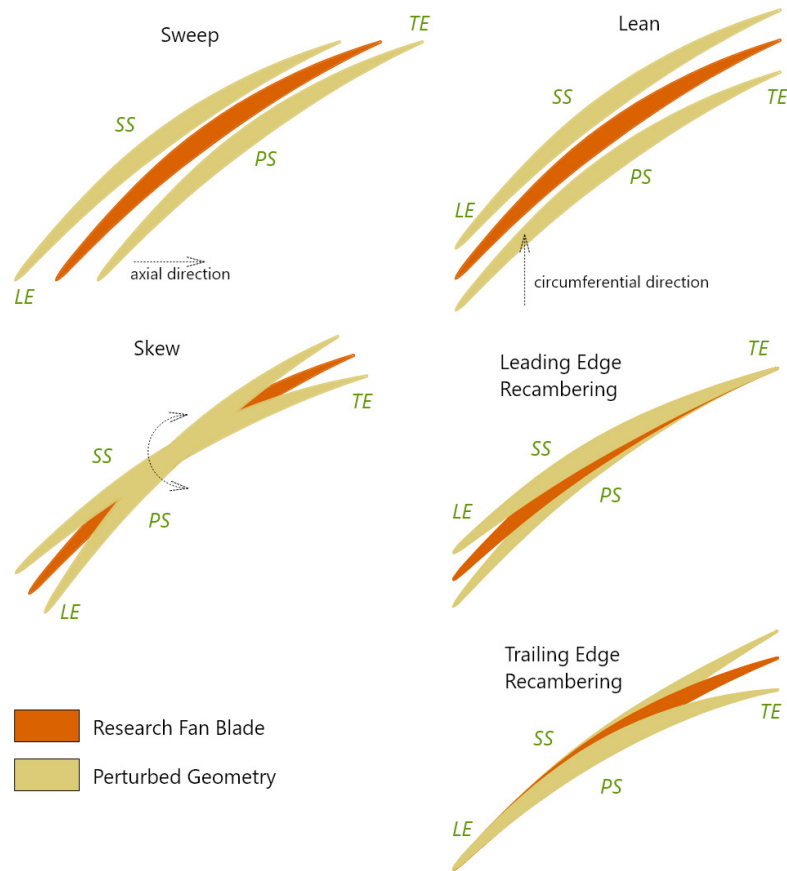


Figure 5.2: Parametrisation of Manufacturing Variability

Identification of the geometrical deviations in manufactured blades can be achieved through reverse engineering processes, which usually decode 3D scans of the part into more meaningful data formats [139]. 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 5.2 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 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 5.3. 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 5.3. This knowledge was employed to define the quadrature collocation points in the UQ of each design assessed in the optimisation process.

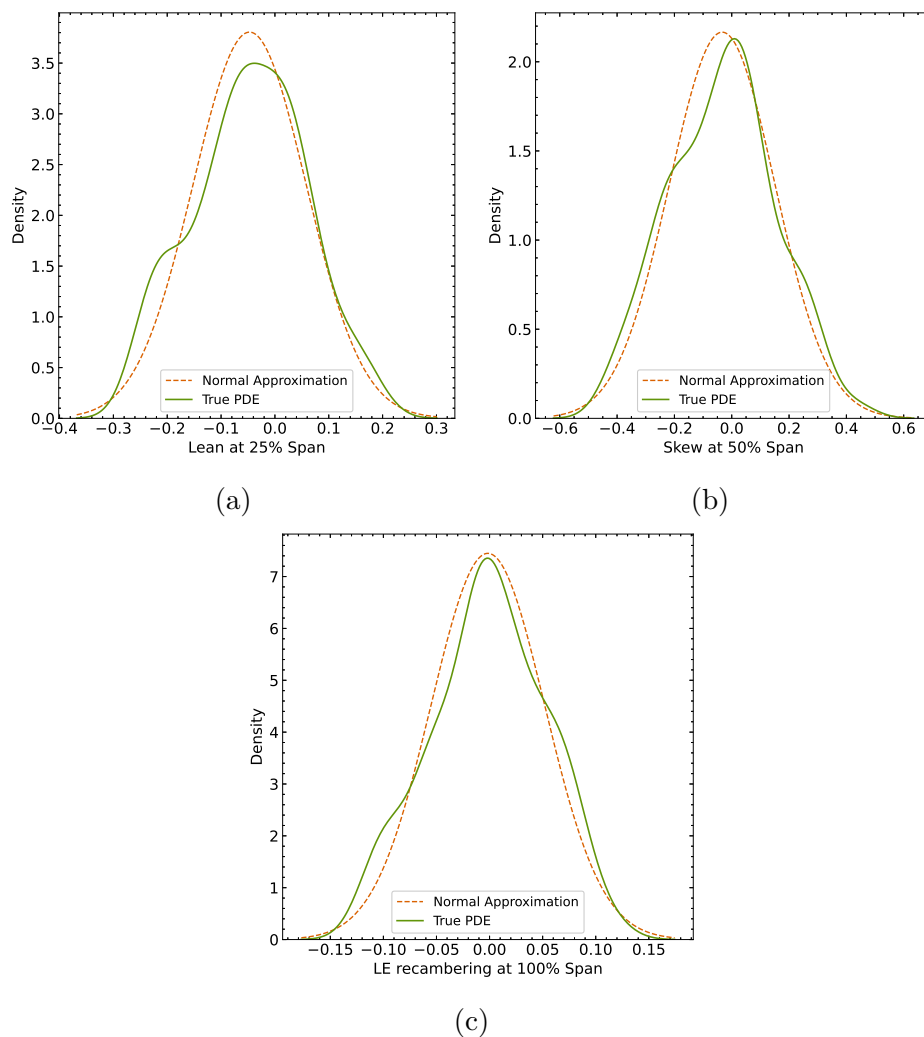


Figure 5.3: Typical PDFs for uncertain parameters

For the generation of new designs, the design vector, \mathbf{x} , was constructed

using the same parametrisation. 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 *recambering* 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 through the aerofoil, thus providing complete control over the camberline. Thus, the resulting dimensionality of the design vector is 35. Such a parametrisation has been previously used for design optimisation studies [138].

5.5 Robust Design Optimisation

The quantity of interest for the optimisation carried out in this section is the fan bypass isentropic efficiency, defined in Eq. 5.4. 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 optimisation cost functions defined in (5.3). A deterministic constraint on the Pressure Ratio (PR) was enforced as per Eq. 5.5 to discourage the optimiser 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.

$$\eta(\mathbf{x}, \omega) = \frac{\left(\frac{p_{0_{exit}}}{p_{0_{inlet}}}\right)^{\frac{\gamma-1}{\gamma}} - 1}{\frac{T_{0_{exit}}}{T_{0_{inlet}}} - 1} \quad (5.4)$$

$$0.99 PR_{datum} \leq PR(\mathbf{x}) \leq 1.05 PR_{datum} \quad (5.5)$$

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 Ch. 2. The application of this methodology also enabled reducing the dimensionality of the parametrisation, 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.

Table 5.1: 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%	-	-

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 5.4 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 first principles, as they mostly control flow turning.

An adaptive sampling strategy based on optimising 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




Figure 5.4: Coefficients of the most dominant direction for (a): mean of the efficiency; (b): standard deviation and (c): pressure ratio

the Pareto front. The results are shown in Figure 5.5, 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 only a slight trade-off between mean and variance of efficiency, but mostly the dominant directions appear to be aligned for the parametrisation considered, and optimising 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 5.4, 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 5.1. Relative to the datum design, the optimised 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 behaviour of the cost functions in the uncer-




Figure 5.5: Pareto frontier for the robust designs

tainty 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 optimisation 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 the work conducted with the same parameterisation and PR constraints in Chapter 3. 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 optimisation 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.

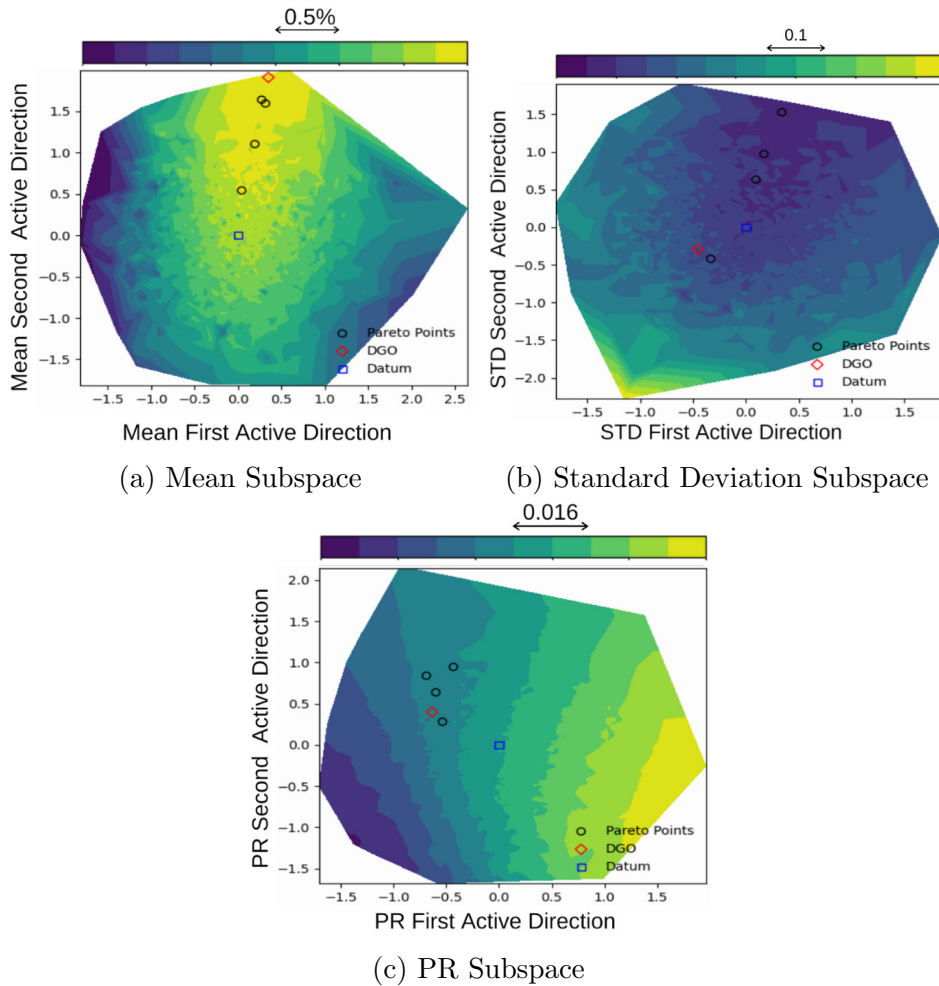


Figure 5.6: 2D Performance Maps for the Quantities of Interest

For visualisation purposes, the performance of the optimum points can be visualised in 2-dimensional active subspace performance maps, following [66]. These maps, shown in Figure 5.6, demonstrate that all the designs lie in a corridor of high mean efficiency, which the optimisation process has been able to identify, and that could be further exploited to devise 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.



5.6 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 [138] for a more in-depth analysis of the DGO shape.

Figure 5.7 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 minimising 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 5.8a. 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 5.8b, thus minimising 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 behaviour, 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 5.8.

The geometry of DGO is presented in Fig. 5.9. 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 it can be observed from the lift plot in Fig. 5.10b. In turn, this reduces the shock-induced separated



(c) View From the Top

Figure 5.7: Geometrical Features of Opt-1 (orange) relative to Datum (grey)

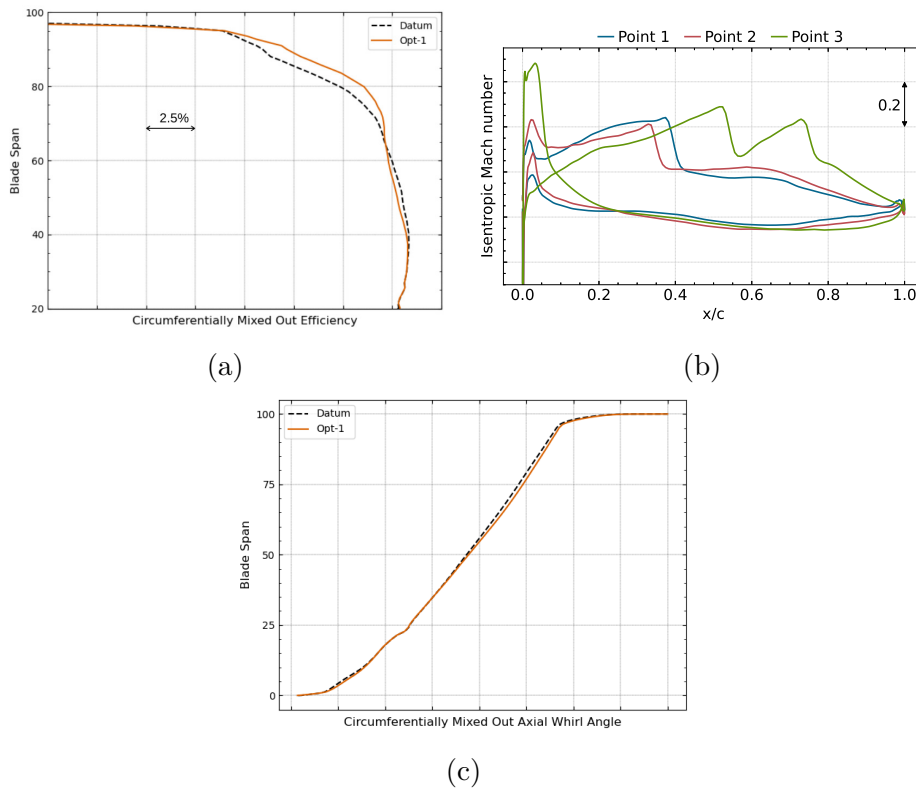
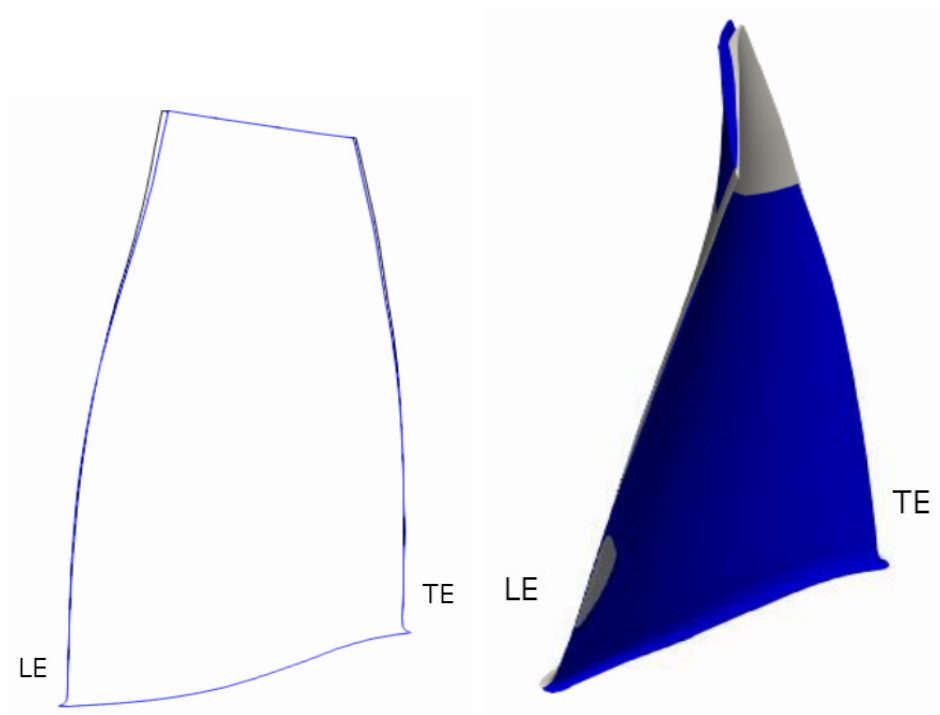


Figure 5.8: Aerodynamic behaviour of Opt-1 relative to Datum. (a): radial efficiency distribution; (b): lift plot at 80% span; (c): exit axial whirl angle distribution

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. 5.10a.


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 optimised geometry are of a highly three-dimensional nature, and their identification has been possible due to the application of a 3D CFD setup employed in the optimisation 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.



(a) Meridional View

(b) Isometric View



(c) View From the Top

Figure 5.9: Geometrical Features of DGO (blue) relative to Datum (grey)

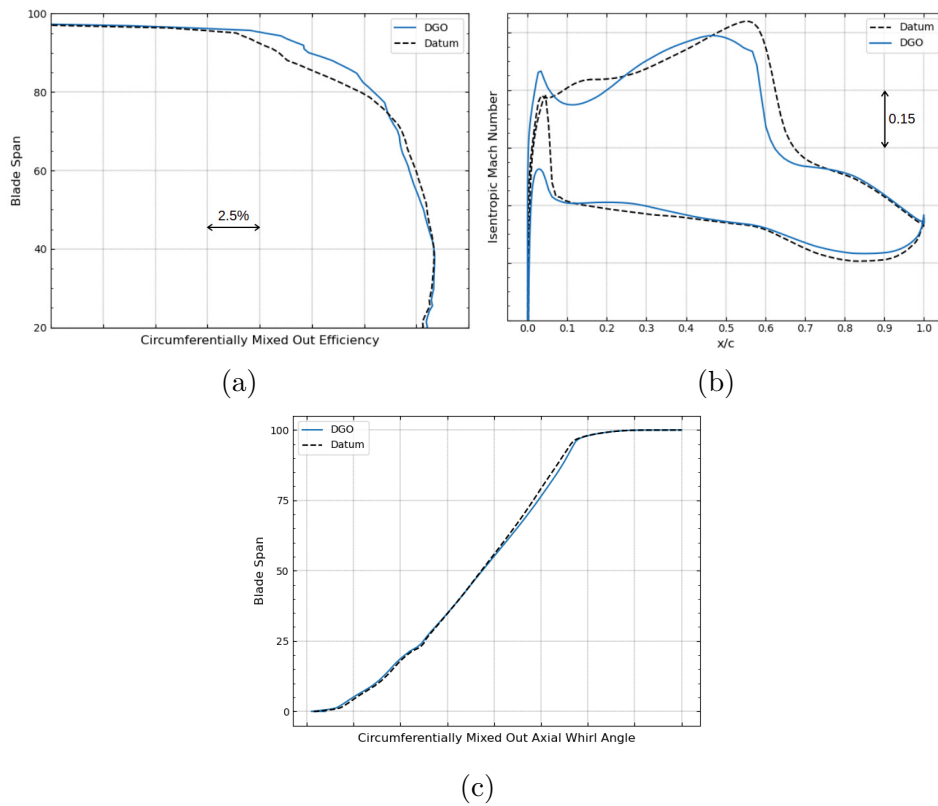


Figure 5.10: Aerodynamic behaviour of DGO relative to Datum. (a): radial efficiency distribution; (b): lift plot at 90% span; (c): exit axial whirl angle distribution

5.7 Conclusions

This study made use of Polynomial Chaos Expansion to derive the statistical behaviour of efficiency and pressure ratio of a fan blade under uncertain perturbations in its shape, due to manufacturing variability. AInADS has been used to perform dimensionality reduction and reduce the computational expense of performing uncertainty quantification. Through accurate response surfaces, a multi-objective, robust-based optimization has been performed, which led to the identification of the trade off between mean and standard deviation for the isentropic efficiency of a modern aero-engine representative fan called Vital. The designs in the extreme of the Pareto frontier were analysed in depth and compared against the result of a deterministic optimization. The analysis shows that the reliability of deterministic optimums may suffer in the presence of manufacturing variability and hence, it is advisable to incorporate uncertainties in the design process.

Bibliography

- [1] D.S. Lee, D.W. Fahey, A. Skowron, M.R. Allen, U. Burkhardt, Q. Chen, S.J. Doherty, S. Freeman, P.M. Forster, J. Fuglestvedt, A. Gettelman, R.R. De León, L.L. Lim, M.T. Lund, R.J. Millar, B. Owen, J.E. Penner, G. Pitari, M.J. Prather, R. Sausen, and L.J. Wilcox. The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. *Atmospheric Environment*, 244:117834, 2021.
- [2] Diego I. Lopez, Tiziano Ghisu, and Shahrokh Shahpar. Global Optimization of a Transonic Fan Blade Through AI-Enabled Active Subspaces. *Journal of Turbomachinery*, 144(1), 09 2021.
- [3] Sri Addepalli, Gerardo Pagalday, Konstantinos Salonitis, and Rajkumar Roy. Socio-economic and demographic factors that contribute to the growth of the civil aviation industry. *Procedia Manufacturing*, 19:2–9, 2018. Proceedings of the 6th International Conference in Through-life Engineering Services, University of Bremen, 7th and 8th November 2017.
- [4] Sarah Freeman, David S. Lee, Ling L. Lim, Agnieszka Skowron, and Ruben Rodriguez De León. Trading off aircraft fuel burn and nox emissions for optimal climate policy. *Environmental Science & Technology*, 52(5):2498–2505, 2018.
- [5] Rolls-Royce plc. Ultrafan: The ultimate turbofan. <https://www.rolls-royce.com/innovation/ultrafan.aspx>. Accessed: 2020-09-14.
- [6] J. D. Denton and W. N. Dawes. Computational fluid dynamics for turbomachinery design. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 213(2):107–124, 1998.

-
- [7] Antony Jameson and John Vassberg. Computational fluid dynamics for aerodynamic design-its current and future impact. In *39th Aerospace Sciences Meeting and Exhibit*, page 538, 2001.
- [8] J. Horlock and J. Denton. A review of some early design practice using computational fluid dynamics and a current perspective. *Journal of Turbomachinery*, 127, 01 2005.
- [9] R. Bellman, Rand Corporation, and Karreman Mathematics Research Collection. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957.
- [10] David Scott and James Thompson. Probability density estimation in higher dimension. *Computer Science and Statistics: Proceedings of the Fifteenth Symposium on the Interface*, 01 1983.
- [11] G. V. Trunk. A problem of dimensionality: A simple example. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-1(3):306–307, 1979.
- [12] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, page 604–613, New York, NY, USA, 1998. Association for Computing Machinery.
- [13] Stefan Berchtold, Christian Böhm, and Hans-Peter Kriegel. The pyramid-technique: Towards breaking the curse of dimensionality. *SIGMOD Rec.*, 27(2):142–153, jun 1998.
- [14] C.R. Taylor. *Applications Of Dynamic Programming To Agricultural Decision Problems*. CRC Press, 2019.
- [15] Antony Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3, 12 1988.
- [16] Antony Jameson and Sangho Kim. Reduction of the adjoint gradient formula for aerodynamic shape optimization problems. *AIAA journal*, 41(11):2114–2129, 2003.
- [17] Mike Giles and Niles Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65, 04 2000.

- [18] Charles A Mader, Joaquim RRA Martins, Juan J Alonso, and Edwin Van Der Weide. Adjoint: An approach for the rapid development of discrete adjoint solvers. *AIAA journal*, 46(4):863–873, 2008.
- [19] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- [20] L Lasdon, S Mitter, and A Waren. The conjugate gradient method for optimal control problems. *IEEE Transactions on Automatic Control*, 12(2):132–138, 1967.
- [21] Shuzhi Gao, Zhiming Pei, Yimin Zhang, and Tianchi Li. Bearing fault diagnosis based on adaptive convolutional neural network with nesterov momentum. *IEEE Sensors Journal*, 21(7):9268–9276, 2021.
- [22] Dimitrios Papadimitriou and Kyriakos Giannakoglou. Compressor blade optimization using a continuous adjoint formulation. *Proceedings of GT2006 ASME Turbo Expo*, 01 2006.
- [23] Mihai Duta, Shahrokh Shahpar, and Mike Giles. Turbomachinery design optimization using automatic differentiated adjoint code. *Proceedings of the ASME Turbo Expo*, 6, 01 2007.
- [24] D. Wang, Lintong He, Y. Li, and R. Wells. Adjoint aerodynamic design optimization for blades in multi-stage turbomachines: Part ii—validation and application. *Journal of Turbomachinery*, 132, 04 2010.
- [25] Shenren Xu, David Radford, Marcus Meyer, and Jens-Dominik Mueller. Cad-based adjoint shape optimisation of a one-stage turbine with geometric constraints. *Proceedings of GT2015 ASME Turbo Expo*, page V02CT45A006, 06 2015.
- [26] Andrea Giugno, Shahrokh Shahpar, and Alberto Traverso. Adjoint-based optimization of a modern jet engine fan blade. *Proceedings of the ASME Turbo Expo*, 06 2020.
- [27] Alistair John, Qin Ning, and Shahrokh Shahpar. The influence of parametrisation setup on the constrained adjoint optimisation of transonic fan blade. *Proceedings of the ASME Turbo Expo*, 06 2020.
- [28] Can Ma, Xinrong Su, and Xin Yuan. An efficient unsteady adjoint optimization system for multistage turbomachinery. *Journal of Turbomachinery*, 139(1):011003, 2017.

- [29] Benjamin Walther and Siva Nadarajah. Constrained adjoint-based aerodynamic shape optimization of a single-stage transonic compressor. *Journal of turbomachinery*, 135(2), 2013.
- [30] Dimitri J Mavriplis. Discrete adjoint-based approach for optimization problems on three-dimensional unstructured meshes. *AIAA journal*, 45(4):741–750, 2007.
- [31] Randy Haupt. Comparison between genetic and gradient-based optimization algorithms for solving electromagnetics problems. *IEEE Transactions on Magnetics*, 31(3):1932–1935, 1995.
- [32] Geoffrey E Hinton and Sam Roweis. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15, 2002.
- [33] D.E. Goldberg, G. David Edward, D.E.G. Goldberg, and V.A.P.H.D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley series in artificial intelligence. Addison-Wesley, 1989.
- [34] James Kennedy and Russell Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95-international conference on neural networks*, volume 4, pages 1942–1948. IEEE, 1995.
- [35] D. Simon. *Evolutionary Optimization Algorithms*. Wiley, 2013.
- [36] Andy Connor, P. Clarkson, Shahrokh Shahpar, and Paul Leonard. Engineering design optimisation using tabu search. *Proceedings of Design for Excellence: Engineering Design Conference*, 06 2016.
- [37] Tiziano Ghisu, Geoffrey T. Parks, Jerome P. Jarret, and John Clarkson. Accelerating design optimisation via principal components' analysis. *Proceedings of the 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 09 2008.
- [38] Akira Oyama, Meng-Sing Liou, and Shigeru Obayashi. Transonic axial-flow blade optimization: Evolutionary algorithms/three-dimensional navier-stokes solver. *Journal of Propulsion and Power*, 20(4):612–619, 07-08 2004.
- [39] B. Yang, Q. Xu, L. He, L. H. Zhao, Ch. G. Gu, and P. Ren. A Novel Global Optimization Algorithm and Its Application to Airfoil Optimization. *Journal of Turbomachinery*, 137(4), 04 2015. 041011.

- [40] Long Wu, Jae-Wook Kim, Alexander George Wilson, and Shahrokh Shahpar. Automatic Design Optimization of a Transonic Compressor Rotor for Improving Aeroacoustic and Aerodynamic Performance. *Journal of Turbomachinery*, 144(8), 03 2022. 081016.
- [41] C. De Maesschalck, V. Andreoli, G. Paniagua, T. Gillen, and B. Barker. Aerothermal Optimization of Turbine Squealer Tip Geometries With Arbitrary Cooling Injection. *Journal of Turbomachinery*, 143(11), 06 2021. 111010.
- [42] Bo Chen and Xin Yuan. Advanced aerodynamic optimization system for turbomachinery. *Journal of Turbomachinery*, 130(021005-1), 04 2008.
- [43] Stefano Caloni, Shahrokh Shahpar, and John Coull. Numerical investigations of different tip designs for shroudless turbine blades. *Journal of Power and Energy*, 230, 08 2016.
- [44] Thomas O. Meyer, Sybrand J. van der Spuy, Christiaan J. Meyer, and Alessandro Corsini. Optimization of a Tip Appendage for the Control of Tip Leakage Vortices in Axial Flow Fans. *Journal of Turbomachinery*, 144(10), 04 2022. 101006.
- [45] Yigang Luan, Lanyi Yan, Yue Yin, and Hao Fu. Optimization of Rib Shape and Inclination in a Square Duct With Response Surface Methodology. *Journal of Turbomachinery*, 144(5), 01 2022. 051007.
- [46] Richard Amankwa Adjei, Chengwei Fan, WeiZhe Wang, and YingZheng Liu. Multidisciplinary Design Optimization for Performance Improvement of an Axial Flow Fan Using Free-Form Deformation. *Journal of Turbomachinery*, 143(1), 12 2020. 011003.
- [47] Andrew Gelsey, Mark Schwabacher, and Don Smith. Using modeling knowledge to guide design space search. *Artificial Intelligence*, 101(1):35–62, 1998.
- [48] Timothy W. Simpson, Timothy M. Mauery, John J. Korte, and Farrokh Mistree. Kriging models for global approximation in simulation-based multidisciplinary design optimization. *AIAA Journal*, 39(12):2233–2241, 2001.
- [49] G. Gary Wang. Adaptive Response Surface Method Using Inherited Latin Hypercube Design Points . *Journal of Mechanical Design*, 125(2):210–220, 06 2003.

- [50] Hongbing Fang and Mark F. Horstemeyer. Global response approximation with radial basis functions. *Engineering Optimization*, 38(4):407–424, 2006.
- [51] Manolis Papadrakakis, Nikos D Lagaros, and Yiannis Tsompanakis. Structural optimization using evolution strategies and neural networks. *Computer methods in applied mechanics and engineering*, 156(1-4):309–333, 1998.
- [52] Farrukh Mazhar, Abdul Munem Khan, Imran Ali Chaudhry, and Mansoor Ahsan. On using neural networks in uav structural design for cfd data fitting and classification. *Aerospace Science and Technology*, 30(1):210–225, 2013.
- [53] M. Venturini. Simulation of compressor transient behavior through recurrent neural network models. *Journal of Turbomachinery*, 128, 01 2005.
- [54] *Axial-Flow Compressor Model Based on a Cascade Stacking Technique and Neural Networks*, volume Volume 5: Turbo Expo 2002, Parts A and B of *Turbo Expo: Power for Land, Sea, and Air*, 06 2002.
- [55] *The Use of Artificial Neural Networks for Performance Prediction of Return Channels for Industrial Centrifugal Compressors*, volume Volume 5: Turbo Expo 2002, Parts A and B of *Turbo Expo: Power for Land, Sea, and Air*, 06 2002.
- [56] Stephane Pierret and R. Van den Braembussche. Turbomachinery blade design using a navier–stokes solver and artificial neural network. *Journal of Turbomachinery*, 121, 04 1999.
- [57] Man Rai and Nateri Madavan. Aerodynamic design using neural networks. *AIAA Journal*, 38:173–182, 01 2000.
- [58] Krzysztof Kosowski, Karol Tucki, and Adrian Kosowski. Application of artificial neural networks in investigations of steam turbine cascades. *Journal of Turbomachinery*, 132, 01 2010.
- [59] Tom Verstraete, Sergio Amaral, René Van den Braembussche, and Tony Arts. Design and Optimization of the Internal Cooling Channels of a High Pressure Turbine Blade—Part II: Optimization. *Journal of Turbomachinery*, 132(2), 01 2010. 021014.

- [60] Karl G Jöreskog. Factor analysis as an error-in-variables model. *Principles of modern psychological measurement*, pages 185–196, 1983.
- [61] Steven Berguin and Dimitri Mavris. Dimensionality reduction in aerodynamic design using principal component analysis with gradient information. *10th AIAA Multidisciplinary Design Optimization Specialist Conference*, 01 2014.
- [62] Peter A Lachenbruch and Matthew Goldstein. Discriminant analysis. *Biometrics*, pages 69–85, 1979.
- [63] I.M Sobol'. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation*, 55:271–280, 02 2001.
- [64] Tan Bui-Thanh, Karen Willcox, and Omar Ghattas. Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM Journal on Scientific Computing*, 30(6):3270–3288, 2008.
- [65] P. G. Constantine. *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*. SIAM Spotlights, 2015.
- [66] Pranay Seshadri, Shahrokh Shahpar, Paul Constantine, Geoffrey Parks, and Mike Adams. Turbomachinery active subspace performance maps. *Journal of Turbomachinery*, 140(041003-1), 04 2017.
- [67] Ashley Scillitoe, Bryn Ubald, Shahrokh Shahpar, and Pranay Seshadri. Design space exploration of stagnation temperature probes via dimension reduction. *Proceedings of the ASME Turbo Expo*, 09 2020.
- [68] David J Munk, Timoleon Kipouros, Gareth A Vio, Geoffrey T Parks, and Grant P Steven. Multiobjective and multi-physics topology optimization using an updated smart normal constraint bi-directional evolutionary structural optimization method. *Structural and Multidisciplinary Optimization*, 57(2):665–688, 2018.
- [69] M.J. Kochenderfer and T.A. Wheeler. *Algorithms for Optimization*. MIT Press, 2019.
- [70] Simon Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice-Hall, 1999.
- [71] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2:303–314, 1989.

-
- [72] C.M. Bishop, P.N.C.C.M. Bishop, G. Hinton, and Oxford University Press. *Neural Networks for Pattern Recognition*. Advanced Texts in Econometrics. Clarendon Press, 1995.
- [73] C.M. Bishop. *Pattern Recognition and Machine Learning: All "just the Facts 101" Material*. Information science and statistics. Springer (India) Private Limited, 2013.
- [74] A.I. Galushkin. *Neural Networks Theory*. Springer Berlin Heidelberg, 2010.
- [75] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- [76] D Randall Wilson and Tony R Martinez. The general inefficiency of batch training for gradient descent learning. *Neural networks*, 16(10):1429–1451, 2003.
- [77] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012.
- [78] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60, 06 2016.
- [79] Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Constructive Approximation*, 26:289–315, 08 2007.
- [80] A. N. Tikhonov and V. Y. Arsenin. *Solution of Ill-posed Problems*. Washington: Winston & Sons, 1977.
- [81] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- [82] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.

- [83] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [84] Antonio Gulli and Sujit Pal. *Deep learning with Keras*. Packt Publishing Ltd, 2017.
- [85] M. Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):111–133, 1974.
- [86] Trent Lukaczyk, Francisco Palacios, Juan Alonso, and Paul Constantine. Active subspaces for shape optimization. *10th AIAA Multidisciplinary Design Optimization Specialist Conference*, 01 2014.
- [87] T. E. Oliphant. Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20, 2007.
- [88] Dieter Kaft. A software package for sequential quadratic programming. *Tech. Rep. DFVLR-FB 88-28, DLR German Aerospace Center – Institute for Flight Mechanics*, 1998.
- [89] Giulio Zamboni and Liping Xu. Fan root aerodynamics for large bypass gas turbine engines: Influence on the engine performance and 3d design. *Journal of Turbomachinery*, 134(6):061017, 01 2009.
- [90] Leigh Lapworth. Hydra-cfd: A framework for collaborative cfd development. *International Conference on Scientific and Engineering Computation (IC-SEC)*, 07 2004.
- [91] Andrea Milli and Shahrokh Shahpar. Padram: Parametric design and rapid meshing system for complex turbomachinery configurations. *Proceedings of the ASME Turbo Expo*, 8, 06 2012.

- [92] N. M. McDougall, N. A. Cumpsty, and T. P. Hynes. Stall Inception in Axial Compressors. *Journal of Turbomachinery*, 112(1):116–123, 01 1990.
- [93] I. J. Day. Stall Inception in Axial Flow Compressors. *Journal of Turbomachinery*, 115(1):1–9, 01 1993.
- [94] F. K. Moore and E. M. Greitzer. A Theory of Post-Stall Transients in Axial Compression Systems: Part I—Development of Equations. *Journal of Engineering for Gas Turbines and Power*, 108(1):68–76, 01 1986.
- [95] T. R. Camp and I. J. Day. A Study of Spike and Modal Stall Phenomena in a Low-Speed Axial Compressor. *Journal of Turbomachinery*, 120(3):393–401, 07 1998.
- [96] V. H. Garnier, A. H. Epstein, and E. M. Greitzer. Rotating Waves as a Stall inception Indication in Axial Compressors. *Journal of Turbomachinery*, 113(2):290–301, 04 1991.
- [97] H. J. Weigl, J. D. Paduano, L. G. Frechette, A. H. Epstein, E. M. Greitzer, M. M. Bright, and A. J. Strazisar. Controls and Diagnostics Committee: Active Stabilization of Rotating Stall and Surge in a Transonic Single-Stage Axial Compressor. *Journal of Turbomachinery*, 120(4):625–636, 10 1998.
- [98] Z. S. Spakovszky, C. M. van Schalkwyk, H. J. Weigl, J. D. Paduano, K. L. Suder, and M. M. Bright. Rotating Stall Control in a High-Speed Stage With Inlet Distortion: Part II—Circumferential Distortion. *Journal of Turbomachinery*, 121(3):517–524, 07 1999.
- [99] Tim S Williams, Cesare A Hall, and Mark Wilson. Low pressure ratio transonic fan stall with radial distortion. *Journal of the Global Power and Propulsion Society*, 4:226–237, 2020.
- [100] Stephanie Weichert and Ivor Day. Detailed Measurements of Spike Formation in an Axial Compressor. *Journal of Turbomachinery*, 136(5), 09 2013.
- [101] Huu Duc Vo, Choon S. Tan, and Edward M. Greitzer. Criteria for Spike Initiated Rotating Stall. *Journal of Turbomachinery*, 130(1), 01 2008.

- [102] G. Pullan, A. M. Young, I. J. Day, E. M. Greitzer, and Z. S. Spakovszky. Origins and Structure of Spike-Type Rotating Stall. *Journal of Turbomachinery*, 137(5), 05 2015.
- [103] J. Dodds and M. Vahdati. Rotating Stall Observations in a High Speed Compressor—Part I: Experimental Study. *Journal of Turbomachinery*, 137(5), 05 2015.
- [104] M. Hewkin-Smith, G. Pullan, S. D. Grimshaw, E. M. Greitzer, and Z. S. Spakovszky. The Role of Tip Leakage Flow in Spike-Type Rotating Stall Inception. *Journal of Turbomachinery*, 141(6), 02 2019.
- [105] Minsuk Choi, Mehdi Vahdati, and Mehmet Imregun. Effects of Fan Speed on Rotating Stall Inception and Recovery. *Journal of Turbomachinery*, 133(4), 04 2011.
- [106] Minsuk Choi, Nigel H. S. Smith, and Mehdi Vahdati. Validation of Numerical Simulation for Rotating Stall in a Transonic Fan. *Journal of Turbomachinery*, 135(2), 11 2012.
- [107] S. Kim, G. Pullan, C. A. Hall, R. P. Grewe, M. J. Wilson, and E. Gunn. Stall Inception in Low-Pressure Ratio Fans. *Journal of Turbomachinery*, 141(7), 02 2019.
- [108] M. Inoue, M. Kurooumaru, T. Tanino, and M. Furukawa. Propagation of Multiple Short-Length-Scale Stall Cells in an Axial Compressor Rotor . *Journal of Turbomachinery*, 122(1):45–54, 02 1999.
- [109] Mert Cevik, Huu Duc Vo, and Hong Yu. Casing Treatment for Desensitization of Compressor Performance and Stability to Tip Clearance. *Journal of Turbomachinery*, 138(12), 06 2016. 121008.
- [110] Matthias Rolfes, Martin Lange, Konrad Vogeler, and Ronald Mailach. Experimental and Numerical Investigation of a Circumferential Groove Casing Treatment in a Low-Speed Axial Research Compressor at Different Tip Clearances. *Journal of Turbomachinery*, 139(12), 10 2017. 121009.
- [111] Jichao Li, Feng Lin, Zhiting Tong, Chaoqun Nie, and Jingyi Chen. The Dual Mechanisms and Implementations of Stability Enhancement With Discrete Tip Injection in Axial Flow Compressors. *Journal of Turbomachinery*, 137(3), 10 2014. 031010.

- [112] Wei Wang, Boxing Liu, Jinling Lu, Jianjun Feng, Wuli Chu, and Yanhui Wu. Comparative Study of Tip Injection in a Transonic and Subsonic Compressor. *Journal of Turbomachinery*, 144(6), 02 2022. 061009.
- [113] Shahrokh Shahpar. Building Digital Twins to Simulate Manufacturing Variation. *Proceedings of the ASME Turbo Expo*, 8, 09 2020. Paper No. GT2020-15263.
- [114] Kuen-Bae Lee, Mark Wilson, and Mehdi Vahdati. Validation of a Numerical Model for Predicting Stalled Flows in a Low-Speed Fan—Part I: Modification of Spalart–Allmaras Turbulence Model. *Journal of Turbomachinery*, 140(5), 04 2018.
- [115] Yangwei Liu, Lipeng Lu, Le Fang, and Feng Gao. Modification of spalart–allmaras model with consideration of turbulence energy backscatter using velocity helicity. *Physics Letters A*, 375(24):2377–2381, 2011.
- [116] David E. Ferguson. Fibonacci searching. *Commun. ACM*, 3(12):648, dec 1960.
- [117] Tiziano Ghisu, Diego I. Lopez, Pranay Seshadri, and Shahrokh Shahpar. Gradient-enhanced least-square polynomial chaos expansions for uncertainty quantification and robust optimization. *AIAA AVIATION 2021 FORUM*, 08 2021.
- [118] Roberto Biollo and Ernesto Benini. Shock/boundary-layer/tip-clearance interaction in a transonic rotor blade. *Journal of Propulsion and Power*, 25(3):668–677, 2009.
- [119] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [120] Mark W Hilburger, Michael P Nemeth, and James H Starnes Jr. Shell buckling design criteria based on manufacturing imperfection signatures. *AIAA journal*, 44(3):654–663, 2006.
- [121] Ronald S. Bunker. The Effects of Manufacturing Tolerances on Gas Turbine Cooling. *Journal of Turbomachinery*, 131(4), 07 2009. 041018.
- [122] Hsin-Hsiung Wang, Miroslav Krstic, and Michael Larsen. Control of deep-hysteresis aeroengine compressors. i. a moore-greitzer type model. In *Proceedings of the 1997 American Control Conference (Cat. No. 97CH36041)*, volume 2, pages 998–1002. IEEE, 1997.

- [123] Longhui Zhang, Antonio Pellegrino, David Townsend, and Nik Petrinic. Strain rate and temperature dependent strain localization of a near α titanium alloy. *International Journal of Impact Engineering*, 145:103676, 2020.
- [124] W. Tabakoff. Compressor Erosion and Performance Deterioration. *Journal of Fluids Engineering*, 109(3):297–306, 09 1987.
- [125] A Hamed, W Tabakoff, and D Singh. Modeling of compressor performance deterioration due to erosion. *International journal of rotating machinery*, 4(4):243–248, 1998.
- [126] E. K. Antonsson and K. N. Otto. Imprecision in Engineering Design. *Journal of Mechanical Design*, 117(B):25–32, 06 1995.
- [127] Wei Chen, J. K. Allen, Kwok-Leung Tsui, and F. Mistree. A Procedure for Robust Design: Minimizing Variations Caused by Noise Factors and Control Factors. *Journal of Mechanical Design*, 118(4):478–485, 12 1996.
- [128] Jian Su and John E Renaud. Automatic differentiation in robust optimization. *AIAA journal*, 35(6):1072–1079, 1997.
- [129] Daniel A DeLaurentis and Dimitri Mavris. Uncertainty modeling and management in multidisciplinary analysis and synthesis. In *38th Aerospace sciences meeting and exhibit*, page 422, 2000.
- [130] Xiaoping Du and Wei Chen. Towards a better understanding of modeling feasibility robustness in engineering design. *J. Mech. Des.*, 122(4):385–394, 2000.
- [131] F. Canavero. *Uncertainty Modeling for Engineering Applications*. PoliTO Springer Series. Springer International Publishing, 2018.
- [132] Gianluca Geraci, Michael S Eldred, and Gianluca Iaccarino. A multifidelity multilevel monte carlo method for uncertainty propagation in aerospace applications. In *19th AIAA non-deterministic approaches conference*, page 1951, 2017.
- [133] Thierry Crestaux, Olivier Le Maitre, and Jean-Marc Martinez. Polynomial chaos expansion for sensitivity analysis. *Reliability Engineering & System Safety*, 94(7):1161–1172, 2009.

-
- [134] T. Ghisu, G. T. Parks, J. P. Jarrett, and P. J. Clarkson. Robust design optimisation of gas turbine compression systems. *Journal of Propulsion and Power*, 27(2):282–295, 2011.
- [135] Pranay Seshadri, Shahrokh Shahpar, and Geoffrey Parks. Robust compressor blades for desensitizing operational tip clearance variations. *Proceedings of the ASME Turbo Expo*, 2, 06 2014.
- [136] Xiaodong Wang, Charles Hirsch, Zhiyi Liu, Shun Kang, and C. Lacor. Uncertainty-based robust aerodynamic optimization of rotor blades. *International Journal for Numerical Methods in Engineering*, 94:111–127, 04 2013.
- [137] Shahrokh Shahpar. Soft: A new design and optimisation tool for turbomachinery. *Evolutionary Methods for Design Optimisation And Control*, K. Giannakoglou et al., ed., 01 2002.
- [138] Diego I. Lopez, Tiziano Ghisu, and Shahrokh Shahpar. Global Optimization of a Transonic Fan Blade Through AI-Enabled Active Subspaces. *Journal of Turbomachinery*, 144(1), 09 2021.
- [139] *Building Digital Twins to Simulate Manufacturing Variation*, volume Volume 2A: Turbomachinery of *Turbo Expo: Power for Land, Sea, and Air*, 09 2020. V02AT32A049.