Probabilistic analysis of manufacturing variability and wear of core engine components

A. Kotnik

Technische Universiteit Delft

TU Delft

Probabilistic analysis of manufacturing variability and wear of core engine components

by

A. Kotnik

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Monday November 6th, 2017 at 14:30 AM.

Student number:	4415973	
Project duration:	October 1, 2016 – August 31, 2017	
Thesis committee:	Dr.ir. M.I. Gerritsma	
	Dr. R.P. Dwight,	TU Delft, supervisor
	Dr.ir. M. Pini,	TU Delft

This thesis is confidential and cannot be made public until November 07, 2022.

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Summary

In the framework of LuFo R&T projects a number of Rolls-Royce Deutschland Low and High Pressure Compressor rotor blades were geometrically scanned at the beginning and end of their life cycle. Based on this data component wear, ageing and manufacturing variability across the life cycle of the engine can be derived. Using Computational Fluid Dynamics (CFD) such geometric component variability can be converted into key parameter changes of aerothermal performance of the compressor. The aim of this thesis is to post-process the available component data and to quantify the geometric changes versus the nominal design using existing analysis tools already being applied in probabilistic assessments of compressor components. In a next step, a probabilistic investigation can be performed using the Rolls-Royce aerothermal analysis tool Hydra in order to quantify the impact of the component variability onto compressor aerothernmal performance. Doing so an assessment and analysis of key geometric component parameters onto compressor performance across the component life cycle can be done.

The results show that there are sources of geometrical uncertainty of the blades throughout their life cycles. The beginning of life geometrical variability has been shown to be larger than the end of life geometrical variability. Furthermore, it has been shown that the geometrical variability due to operational wear is too small to be accurately captured by the current CFD evaluation methodology. This methodology focused on evaluating single blade geometries without the effects of fillet and tip gap variabilities.

Acknowledgements

This work has been performed in partial fulfilment of the Aerospace Engineering MSc programme requirements at TU Delft in partnership with Rolls-Royce Deutschland Ltd Co & KG. I would like to thank Rolls-Royce for funding this thesis, allowing its publication, and all the learning and development opportunities presented to me during the last two years.

I started my studies at TU Delft in autumn of 2014 as a part of a truly international group. I count myself as lucky to have had such a varied mix of flatmates and friends, as they made the experience much more enjoyable and richer. I would like to especially thank my friends Anand, Andrea and Laurence.

A very large part of my development and learning during my MSc studies is therefore tightly tied to my work and experiences at Rolls-Royce. I would like to thank my colleagues from both the UK and Germany for their guidance and sharing their knowledge with me. Furthermore, I would like to thank them for all the tennis "lessons", and interesting lunch and coffee breaks.

I would also like to thank Richard Dwight and Marcus Meyer, my thesis supervisors from TU Delft and Rolls-Royce respectively. I would like to thank them for their support at all stages of preparing this thesis, and for allowing me creative freedom while doing my work. Furthermore I have managed to learn a lot of valuable knowledge from them, which I will no doubt use in the future.

Lastly, I would like to thank my parents for their support and love. Without them arriving at this point would have been much more difficult.

A. Kotnik
Berlin, August 2017

Contents

List of Figures

List of Tables

1

Introduction

World War II was a profound test of humanity whose resolution was an enormous industrial challenge that required efficient use of resources. In response, the use of statistical quality-control concepts in manufacturing has been greatly expanded. Wartime experience made it apparent that statistical techniques were necessary to control and improve product quality. After the war Japan, with the help of US experts such as Deming and Juran, developed its own ideas regarding quality control [44][7][8].

The benefits followed quickly: "The result is now legendary. From a standing start, by the 1970s manufacturing industry in Japan was producing cars, motorcycles and domestic goods cheaper than anyone else. Added to that, such items were becoming nicer to use and of better quality" [60]. The increased competitiveness drove the USA to develop similar quality-control projects such as Total Quality Management, Six Sigma and Lean manufacturing. Interestingly, all these projects pursue the goals of "minimising waste, improving customer satisfaction and improving financial results" [3] and all originate from the Japanese projects [8].

Six Sigma particularly became popular in the 1990s when it was championed by companies such as Motorola and General Electric, which have benefited from it dramatically [7]. General Electric specifically attributed hundreds of millions in savings per year to Six Sigma activities and Motorola claimed billions in savings overall [13][14][26].

It is clear that assuring quality has enormous benefits to companies in various fields and it does so on different fronts, such as minimising waste, attracting and satisfying customers, and maximising product value over its life cycle. Furthermore, it has become clear that quality assurance is vital in keeping up with the competition.

Almost two decades after launching Six Sigma, the market is evolving from quality of product to quality of service [15]. This is especially true for aircraft engine companies which are moving into service based contracts, such as the Rolls-Royce TotalCare®[5][48]. Because of increased quality demands from the customers and to secure future profits, it is more important than ever for the engines to be designed in a way that reduces waste and ensures reliability and performance throughout the engine life cycle.

At the core of this, as well as at the core of the engine, are the gas path components such as compressor and turbine blades, whose performance have major implications on the performance of the entire product. To ensure reliability and performance of the engine throughout its life cycle it is therefore important to ensure that the level of quality of compressor and turbine blades is high throughout their life cycle.

1.1. Opportunities for quality improvement of compressor and turbine blades

Improving the quality of these components can be done in multiple ways. Assuming that the manufacturing process has already been optimised from the economical perspective, the quality of the end product can now be primarily changed in the design phase. The following opportunities, or combinations of them, are available.

- The blade design can be adjusted such that the manufacturing process will produce less blades that do not conform to the geometric requirements.

- The blade design can be adjusted such that the blade mechanical performance will be less sensitive to manufacturing variability and deterioration.

- The blade design can be adjusted such that the blade aerodynamic performance will be less sensitive to manufacturing variability and deterioration.

The first point addresses the material and economical waste that occurs directly as a result of the manufacturing process. The manufacturing of every blade incurs costs while only the conforming blades make a return profit. The return profit must outweigh the incurred costs and therefore the price of a conforming blade rises with the number of non-conforming ones. The activities the first point are linked with the manufacturing process and its techniques.

The second and third points are interconnected as both arise from the same source - geometrical variability due to manufacturing and deterioration processes. The impact of geometrical variability on the blades can be viewed from multiple perspectives. It can be seen in terms of stress bearing capability or in terms of aerodynamic performance. These are addressed in points two and three respectively.

The stress bearing capability determines how many operating cycles the blade can mechanically withstand. Blades that have high mechanical stress variability require higher safety factors to ensure a given safety level. This then translates to an increase of the weight of an individual blade and consequently the entire engine. However, the blades are expected to be very similar to one another in terms of shape mass and volume and therefore any variability in mechanical stresses is expected to be small.

On the other hand the aerodynamic performance is the main driver in blade shape design and small deviations in the blade shape are expected to have a relatively large impact on performance. In the case of turbine blades this also includes blade cooling.

Note that all three opportunities involve aerodynamic modelling of the blades at some point. In the end, aerodynamic interaction of the blade with the working fluid is its purpose. No matter which direction the cost reduction activities take, this purpose must be preserved within some performance requirements. However, from an aerodynamics perspective it is point three that is the crown jewel of the listed opportunities. This is because this activity focuses directly on aerodynamic performance as opposed to the other two that feature aerodynamics only tangentially.

1.2. Conclusions of the literature review

To give a visual overview of the current state of the art the work presented in sections 3.1 and 3.2 is summarised in table 1.2, with the work focused on holistic modelling omitted. In table 1.2 some of the references are color coded with either gray, blue or cyan. The grey references represent work where no CFD has been performed and the analysis was based on the parameters of the blades. Cyan denotes the work performed on centrifugal compressors. The results of references highlighted in blue contain the effects of manufacturing variability and deterioration combined. Their separation has not been performed within the scope of the respective works.

		Manufacturing		**Deterioration**	
		Compressor	**Turbine**	**Compressor**	**Turbine**
Quantification	**2D**	[16]	[12]	[28],[18],[52]	
	3D	[37],[35],[36], [38] ,[39],[19], [56]	[21],[27], [22], [9],[46]	[40],[53]	[55]
Optimisation	**2D**	[29],[30],[11] [6]	[42]	[31]	
	3D	[24]	[61],[4]	[25]	

To summarise the results, the effects of manufacturing variability on compressor and turbine performance have already been studied significantly. The variation has been described both by Principle Component Analysis (PCA) approaches and more intuitive NACA aerofoil parameters. It was reported that the most important modes of the PCA resemble the NACA aerofoil parameters. Furthermore, the approach using NACA like aerofoil parameters has been used in more publications than the PCA approach.

The uncertainty quantifications of manufacturing scatter have already been performed using 2D and 3D Monte Carlo (MC) CFD on individual blade, stage, and whole module levels for compressors. For turbines most applications focused on single blade analyses, but a synthesis approach was also demonstrated that was able to give results for a full turbine stage. The most important parameters were different in different studies, but usually limited to stagger angle, maximum camber and blade thickness at various positions. This suggests that the identity of the most important parameters depends on the blade design itself. An important lesson is also that the performance scatter reduces with an increasing amount of uncertain blades, but the performance offset increases.

In the field of deterioration quantification, the initial investigations have been performed on a holistic, whole engine level. The investigations focused on individual blade components have been performed in 2D. These works simulated the deterioration without using measured data. In 3D the geometrical variations between blades have already been analysed. Some work has also been performed to quantify these variations, however the effects of manufacturing and deterioration were not separated. Because of this it is not known which process has a larger impact on performance and the effects cannot be modelled separately. The 3D work here was also used NACA like aerofoil parameters to characterise the blades.

The next step based on uncertainty quantification of compressor and turbine blades is to determine how do the effects of manufacturing variability and deterioration compare. This can be done in terms of geometrical parameters as well as performance scatter. Furthermore, the performance evaluations could be enhanced by incorporating all known deterioration effects such as, erosion, deformations, dirt thickness accumulation, surface roughness changes and tip height reduction in the 3D domain.

Significant work has also been performed in the field of robust optimisation of compressor and turbine blades. This work has also mostly focused on 2D applications. The manufacturing and deterioration scatter in 2D were simulated using Hicks-Henne functions. The optimisation efforts were not based on measured uncertainty data and used arbitrary assumptions to model it.

Some 3D methodologies tackling robust design against manufacturing variability were also presented. Here again the NACA like parameters were used. In the case of deterioration work on robust optimisation of compressors to take into account surface roughness increases has already been made. The opportunities here remain to apply the robust optimisation based on measured manufacturing uncertainty. This could initially be performed on an individual blade level, followed by work on the stage and entire module levels.

The next steps of robust optimisation of compressor and turbine blades include combining of its multiple effects, for example deterioration geometry changes, surface roughness increase and tip gap behaviour. Robust design of blades against deterioration is also an open opportunity. This could also be performed on the individual blade, stage and module levels.

1.3. Research question

This thesis is focused on understanding the impact of deterioration on engine performance within the the demarcated framework presented in subsequent sections 4.2, 4.3, and 4.3 . The research question is:

> *The research question is to investigate the impact of in service deterioration on blade (1) geometry, (2) performance at the cruise design point, and (3) how these compare to the corresponding impacts of manufacturing variability by analysing and comparing the blade geometrical parameters and 3D CFD results of a set of blades measured before assembly into the engine and after removal during an overhaul activity.*

Answering this question will give an insight into how individual blades deteriorate geometrically, how to appropriately model this process, and how it compares to geometric variability due to manufacturing. The models capable of simulating manufacturing variability and deterioration effects can support subsequent robust optimisation activities.

1.4. Thesis outline

To introduce the reader to the topics discussed in subsequent chapters, chapter 2 briefly presents the major terms discussed in these chapters and the ideas that underpin them. A very general introduction is given to uncertainty quantification and robust optimisation. Furthermore, the effects of engine deterioration are briefly discussed and an example calculation of possible benefits of reducing performance variability is worked through.

In order to be able to constructively add to the effort of improving aerodynamic quality of compressor and turbine blades, the current state of the art needs to be reviewed. This is done in chapter 3, which outlines the steps that need to be taken to improve quality and what work has been done on these steps so far. Furthermore, within the review of literature different methodologies of these steps are briefly presented. Lastly, the possible next steps are presented.

In chapter 4 the steps selected to be performed in the scope of this thesis are presented. The dataset and the general assumptions made are also discussed. Research objectives and goals are formulated for each of these chapters, and the individual steps that make up the subsequent chapters are outlined. Lastly, the methodology used is selected based on the findings of the literature review.

Chapter 5 presents the geometrical evaluation of manufacturing variability and deterioration. The comparisons are done by comparing measurements directly, as well as characterising the blades and comparing the characterisations.

Chapter 6 presents the quantification of the effects of geometrical variability on blade performance. The measured blades are evaluated using 3D CFD and their results are compared to isolate effects due to individual uncertainty sources. Lastly, the influence of individual geometric parameters on the performance Quantities of Interest (QoI) are analysed.

Chapter 7 presents how to model the geometrical effects of manufacturing and deterioration. First the process that rebuilds the blade geometries from a set of geometrical parameters is discussed. Secondly the modelling of these parameters such that they represent manufacturing and deterioration is presented. Lastly, the convergence of the modelling to the measured results is presented.

2

Useful background concepts

This section aims to introduce the basic ideas used within this thesis. As outlined in chapter 1, the main opportunities that can be taken are the adjustment of the blade design such that the end product quality levels are increased.

The concept of uncertainty quantification is presented in section 2.1. Adjustment of the blade design for increased quality level relies heavily on the ability to predict these quality levels of a proposed blade design. This prediction includes the evaluation of the impact of uncertainty on a selected parameter of interest, which is called uncertainty quantification.

With the ability to predict the performance scatter of a specific design, the quality level of a product series can be increased by altering the design itself. This is called robust design and is discussed in section 2.2.

Combined, the concepts of uncertainty quantification and robust design allow to account for the uncertainty of processes in the blade design. Deterioration is one of the most important uncertain processes as it heavily impacts engine performance. A short overview of engine deterioration is given in section 2.3.

Lastly, for sake of clarity, the naming convention for a blade is outlined in section 2.5. This section describes the prominent features of the blade and provides the names used for these features in the following chapters.

2.1. Uncertainty quantification

An illustration of an example Uncertainty Quantification (UQ) problem can be seen in figure 2.1. This figure contains several features, such as a 2D Gaussian distribution, a response surface and several point clouds. These elements together show the impact of two uncertain parameters of a sample product on a selected performance parameter of this product.

The 2D Gaussian distribution seen in figure 2.1 represents the joint probability distribution of the two uncertain parameters. The two magenta lines plotted on the joint distribution surface show the location of the deterministic design parameter. In other words, the design intent requires both parameters to be zero, but due to uncertainty they span a range of values.

The uncertain parameters could be many things, ranging from geometrical, operational, environmental parameters to physical model configuration settings such as specific heats, fuel calorific values or other property coefficients. In the case considering the geometrical changes of blades, this distribution characterises the geometrical variability of a batch of manufactured blades.

The response surface characterises the performance of the product, and is a function of the uncertain parameters. The response too could be many different things, that range from simple area calculation between two uncertain lengths, complex 3D CFD and other models. Furthermore, this surface is not necessarily smooth and can feature non-linearities, e.g. shocks of a CFD solution. In the case of blade evaluation the response is the aerodynamic, or other multidisciplinary performance of the blade geometry.

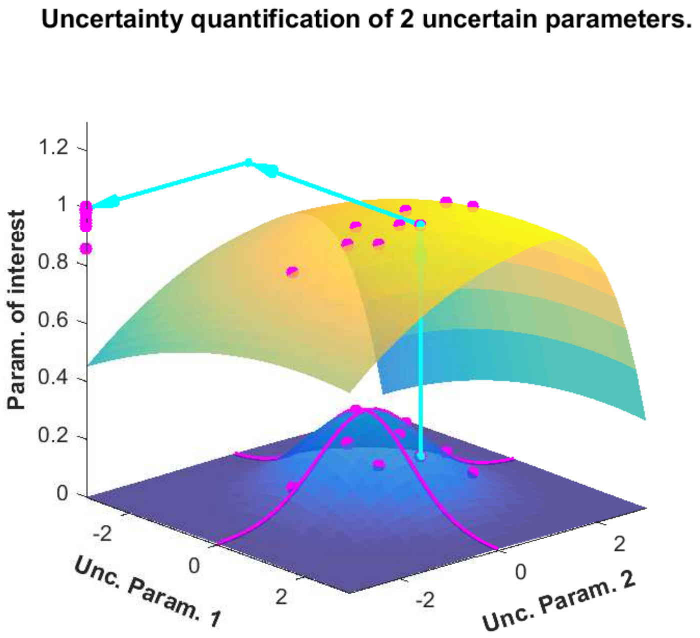

Figure 2.1: Example illustrating the principle of uncertainty quantification.

The point clouds are comprised of 10 sample points, which represent 10 product realisations. There are three point clouds in total, connected with a series of arrows. The first is on the 2D Gaussian distribution, which shows where the points have been sampled. The second is on the response surface and shows the performance of the sampled realisations. This point cloud shows the performance scatter in 3 dimensions. The interest lies only in the performance scatter, therefore this cloud is compressed into 2D and finally into 1D with two projections that remove the scatter due to the uncertain parameters. The 2D and finally 1D projections are shown as the last two point clouds. The arrows further illustrate the steps in the procedure.

The actual costs connected to the performance scatter can then be evaluated by calculating the cost of producing non-conforming products. Alternately, the impact of the remaining scatter can be quantified based on the company's business model.

2.2. Robust optimisation

Based on the illustration of an example UQ problem in section 2.1, an illustration of the concept of robust optimisation is given in figure 2.2. This figure shows the 2D Gaussian distribution, response surface and sample points from figure 2.1. Additionally, another 2D Gaussian distribution that is offset by 2 along one of the axis and plotted in yellow, has been added along with its corresponding sample.

For the sake of the following discussion, the 2D Gaussian distributions in blue and yellow are named design 1 and design 2 respectively. In this example, both design 1 and design 2 uncertain parameters have standard deviations of unity. Furthermore, it is assumed that higher values of performance parameter are more optimal.

The first thing that can be noticed about design 2 is that the intended performance point lies at a higher positions on the response surface than the intended performance point of design 1. However, the performance scatter of design 2 is also larger, which can be seen by comparing the blue and magenta points plotted on the vertical axis.

Notice that beyond the intended design point of design 2 the response surface drops quickly. Some of the samples sampled from the design 2 joint Probability Density Function (PDF) therefore fall within this region, which then causes the performance scatter to increase. Essentially, relative to design 1, design 2 trades low performance scatter for higher nominal performance.

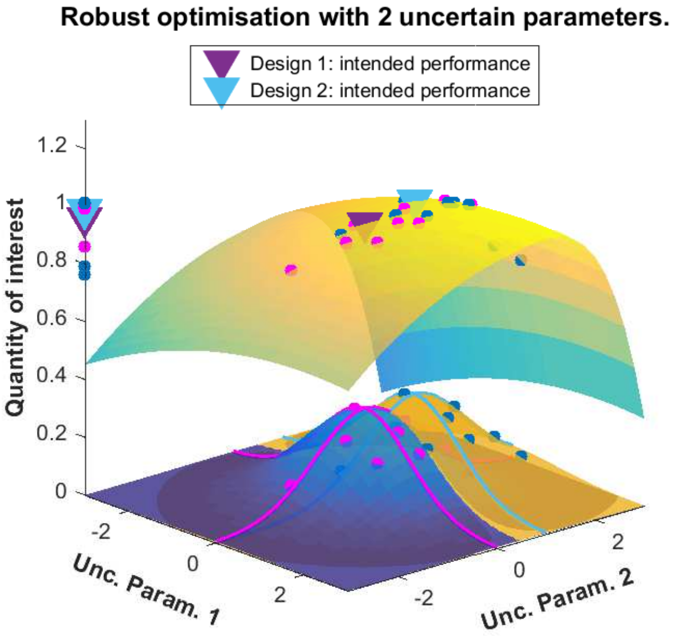

Figure 2.2: Example illustrating the principle of robust design.

Design 1 is the more robust design. In this case, relative to design 2, realisations of design 1 will perform better on average, and have less realisations non-conforming to minimum performance requirements. However, design 1 is not necessarily the optimal robust design point. This depends largely on the formulation of the cost function.

An example cost function to be minimised during robust optimisation might be:

$$f(x) = -A_1\mu(p) + A_2\sigma(p) \tag{2.1}$$

Equation 2.1 will perform a basic trade-off between the mean of the performance and its variability, represented by the standard deviation. In this case p is a vector of realisation performances, with one value for each of the realisations. Notice that due to construction, the function can be minimised by minimising the performance variability, and maximising the mean performance. The exact location of the optimal point is then governed by the response surface and trade-off coefficients A_1 and A_2. With the variation of the selections for A_1 and A_2 a Pareto front can be obtained.

Notice that the selection of A_1 and A_2 can lead to unexpected results. Imagine the extreme case where p is an efficiency and $A_2 >> A_1$. The efficiency will be in the range of 0 to 1 by definition. As the standard deviation dominates the mean performance, the optimisation could end in a location where the efficiency is 0, but is very consistent over all the realisations. In essence, this would be a robust design, however not one that is very interesting.

Furthermore, even though the input uncertain parameters are normally distributed, this is not necessary true for the output performance parameter. Specifically, the performance parameter will only be normally distributed when the inputs are normally distributed and the response surface is linear or close to linear in the region where the samples are evaluated.

The standard deviation is well suited for evaluation of symmetrical samples, but does not give any information of asymmetry if it is present. An alternate approach could be to use the first and third quartiles to assess this. In cases where minimal performance requirements are imposed, the first quartile of the performance should be used instead of the standard deviation.

Using the first quartile effectively changes the nature of the cost function. If the standard deviation is used in the cost function, the performance scatter will be minimised. Note that products that perform better than intended also increase the standard deviation. However, optimising to reduce the scatter caused by better products is counter productive. If the first quartile is used, the performance of the three quarters of products is maximised, thereby avoiding the situation where better performing products could cause an increase in cost.

2.3. Deterioration

The aim of this section is to give a rough overview of the effects that deteriorate the performance of compressor blades with in service use. These effects can change the geometry of the gas path by changing the aerofoil metal geometry, changing the aerofoil geometry by particle deposition, and displace it through the degradation of the alignment surfaces. These effects result in a reduced stage and overall compressor performance, as well as lower operational surge margins.

Subsection 2.3.1 presents the effects that impact the aerofoil metal geometry, and subsection 2.3.2 presents the change of tip gaps with in service use. Subsection 2.3.3 discusses the change of aerofoil geometry due to particle deposition and its related surface roughness increase with in service use. Lastly, subsection 2.3.4 gives a general explanation of engine performance over the span of its life.

2.3.1. Geometrical changes of blades

There are several effects that directly change the metal geometry of the blade aerofoils. In the case of compressor blades the most prominent effects that pertain to the aerofoil section itself are:

- erosion - mechanical reactions with deposition particles
- corrosion - chemical reactions with deposition particles
- Foreign Object Damage (FOD)- damage caused by larger ingested objects

Erosion is caused by ingested abrasive particles in the air such as ash, and sand. These particles then cause a reduction of aerofoil geometry, and subsequent surface roughness increase. The reduction of geometry can be seen, among others, as the blunting of blade LE, and reduction of the blade chord length [20].

Corrosion causes pitting on the compressor blades, which causes small indentations in the blade surface as a result of material loss due to chemical reactions. "Experience has shown that deposits on compressor blades often contain sodium and potassium chlorides. These combine with water to form aggressive solution that promotes pitting corrosion of the blades" [41]. Other molecules, such as sulfur dioxide (SO_2) and hydrocloric acid (HCl) also create acidic conditions that promote pitting [41].

FOD is typically caused at the forward compressor stages, as these are some of the first components along the gas path. However, in some cases the foreign objects reach the rear stages as well, and cause damage. The extent of the damage depends on the foreign object size and composition, blade construction, and impact location. It can lead to direct or secondary failure. An example of foreign objects that can cause damage is ice [41].

2.3.2. Tip gaps

The performance of compressor blades is dependent on the size of the tip gaps. Because of this gap air from the pressure side can flow over the tip to the suction side, which introduces losses and reduces the blade efficiency. Intuitively, the larger the gap the worse the performance is. The gap is also subject to uncertainty and should therefore be evaluated along with other geometrical parameters.

This size of the tip gap also depends on the mechanical and thermal loads the components are subject to during operation. For example, on take-off the engines are run to the maximum power setting and therefore maximum revolutions per minute (RPM), which translates into maximum centripetal forces stretching the blades. This reduces the tip gap and can make the blades rub with the liner attached to the casing. Other examples include temperature expansion of the blades and the casing.

Furthermore, the tip gap is expected to change due to deterioration effects during the course of an engines life. Of all the aforementioned compressor components the drum and the casing are not expected to deteriorate geometrically. On the other hand, the liner and the blade are both expected to have their geometries change. The intuitive expectation is that the blades and liner have their heights decreased as a consequence of geometric interaction with each other. The degree of this interaction and its distribution between both these components depends on the specific blade and liner combination as well as the operational use of the engines.

Two types of liners are available: abrasive and abradable. Abrasive liners are those that are expected to abrade the tips of the blades. Inversely, abradable liners are those that are expected to be abraded by the blade interaction. Therefore depending on the liner blade combination either one or the other will mostly be abraded. Furthermore it is interesting to note that the tip gap will not necessarily increase constantly along the axial direction. Instead the tip gap could vary in height along the blade tip. For example, the tip gap at the LE and TE could be larger than at the mid-chord.

This also means that the tip gap depends on the operation of the engine and its corresponding deterioration. Engines that are not requested to provide maximum available thrust often are expected to have lower tip gaps that those that do.

2.3.3. Fouling

Fouling is the combined effect of aerofoil thickness increase through particle deposition, and the surface roughness increase connected with it. Increased thickness can reduce the mass flow through a compressor stage, while the surface roughness increase increases friction losses. This can have significant effects on performance. For example, Meher-Homji and Brownley [41] consider an annual loss of 3% of power due to deterioration extremely modest. Fouling accounts for $70 \approx 85\%$ of performance deterioration, as per Diakunchak [10]. The main impact of fouling is the increase of surface roughness as opposed to the increase of aerofoil thickness [47].

Fouling is caused by pollutants adhering to the blade surface through several different mechanisms. These include settling, diffusion, interaction, inertial impaction and electrostatic forces. Common examples of particles that cause fouling are smoke, oil mists, carbon, and sea salts, and are usually roughly $2-10\mu m$ in diameter [32]. High concentration of these particles in the air can be found at airports near the sea and at airports in countries with high levels of air pollution.

Some of the effects of fouling can be prevented, to a degree, and its effects mitigated, again to a degree. Industrial gas turbines commonly have inlet filters installed in order to remove the pollutants from the air before entering the turbo-machinery, and thus prevent much of the fouling. This is impractical in the case of aircraft engines. The effects of fouling can be mitigated by engine washing [33] [41]. Washing can remove the pollutants from the blades and therefore restore the initial performance.

2.3.4. Engine life

The estimates of engine lives are largely based on the behaviour of the engines through their lives. Basically the engines start their lives with some initial performance and then lose it over time due to deterioration effects. After the level of performance drops to the point where the engine can barely still generate the thrust needed to perform a worst case condition take-off (with one engine inoperative) in the worst case ambient conditions (very high temperatures) with a fully loaded aircraft, the engine is removed from service and sent to be refurbished in order to restore its performance. After the overhaul the engine is sent back into service.

The engines are expected to have 3 overhaul services in their entire lives. After that enough time has passed that it is worthwhile to replace them with newer engine types. As an example, an engine that can complete 5000 flights [45] before an overhaul is needed and flies a route lasting 10 hours will typically need an overhaul every 7 years. This is based on the assumption that the airline will fully utilise the aeroplane and make 2 flights per day with the aircraft, with a 2 hour turnaround time per flight. Therefore a total life of an engine can be estimated to be roughly 30 years.

To visualise the performance of an engine fleet over the course of its life a schematic is given in 2.3. This figure shows a life limiting performance parameter plotted on the y-axis against the flight cycles performed on the x-axis. It also shows three overhaul activities as the spikes where the performance parameter is restored to the initial performance level. The intervals between these spikes are called intermediate lives in this chapter.

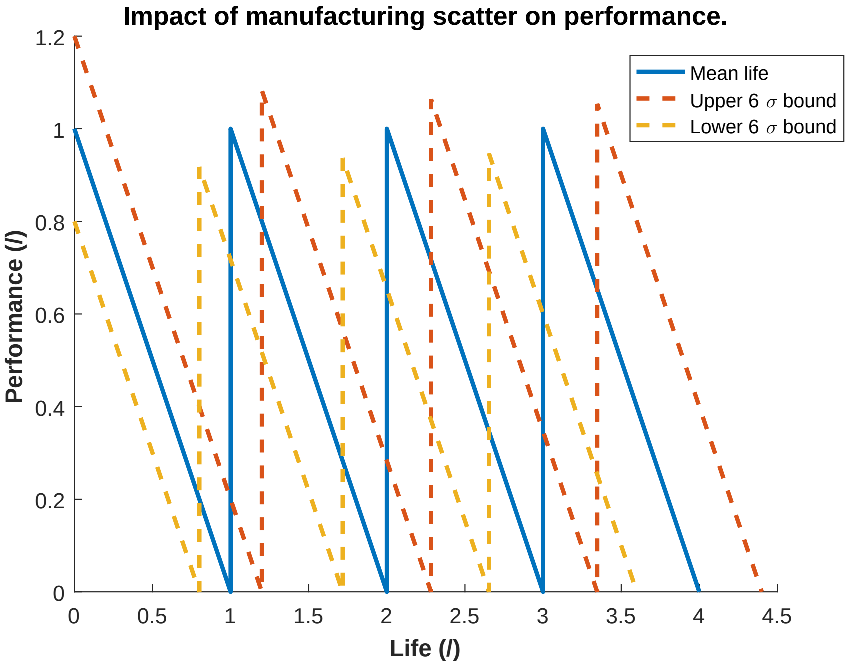

Figure 2.3: Theoretical behaviour of performance bounds of engine fleet performance with operational use.

For simplicity it has been assumed that in each of the intermediate lives the performance scatter due to deterioration that will occur is the same. In other words, no deterioration scatter can be seen in figure 2.3. Furthermore, overhauls will not restore the exact amount of the performance lost as the same manufacturing scatter that caused the initial scatter is also present here. However, an assumption can be made that the mean of the overhauled engine fleet will match the mean of the initial performance. Lastly, it is assumed that the intermediate life values follow a normal distribution.

The initial scatter is combined with deteriorated scatter and translated into the intermediate life scatter at the first overhaul stage of engines lives. The performance of engines is then recovered during overhauls and then depleted once more as the engine returns to operational service. In total the engines experiences the first deterioration process after introduction to service and three subsequent overhaul and deterioration cycles.

As the amount of performance recovered during overhaul is also a random process all of the intermediate lives are also random variables. The total life prediction for an engine fleet is therefore a sum of 4 random variables, each representing a specific intermediate life. As per the assumptions regarding the overhaul effectiveness the total life can therefore be calculated using 2.3:

$$Z = \sum_{i=1}^{n} X_i \quad \text{where} \quad X_i \sim N(\mu_i, \sigma_i^2) \tag{2.2}$$

$$Z \sim N(\sum_{i=1}^{n} \mu_i, \sum_{i=1}^{n} \sigma_i^2) \tag{2.3}$$

It can be seen from this equation that to calculate the estimate of fleet life, the estimates for intermediate engine lives and the scatter of intermediate lives must be given. Based on figures available in literature the following numbers were selected:

2.4. Impact of manufacturing scatter and deterioration on finance

Section 2.3 gave an overview of engine performance behaviour over its life and presented the most important causes of performance deterioration. In the end the impact of this deterioration is financial. The aim of this section is to provide the financial motivation for robust design against deterioration.

The motivation is presented as a quantified example that shows the possible financial opportunity of improving robustness. The results are given as the size of opportunity per % improvement of mean life and % reduction of its scatter. The analysis is supported by sample cost values freely available in internet mass media.

The example is broken into two parts. Subsection 2.4.2 quantifies the opportunity of reducing the scatter of life. Separately, subsection 2.4.3 quantifies the financial impact of reducing the rate of deterioration. Both parts together form the final opportunity available.

2.4.1. Estimation of engine prices and lives

To do the evaluation of the opportunity present within the reduction of performance scatter some initial cost estimates must be made. These are the number of engines of a particular type that are being evaluated, the standard deviation of their lives, the price of an overhaul, the price per hour that the airline pays and the amount of hours flown per every cycle.

The price of small business jet engines is roughly can be roughly estimated from the prices of the aircraft they operate on. The BR700 series operates on business aircraft such as the Bombardier Global 5000 and Gulfstream 550. The various derivates of these airframes cost between 50 and 70 million $ [63]. The costs of the BR715 engine can be estimated from the price of the Boeing 717, which is roughly 33 million $ [2]. Assuming that the engines account for a quarter of the airframe costs, the estimated price for an engine from the BR700 family is between 8.75 and 4.13 million $. Assuming that the airlines opt for the cheapest option the price lower price of 4.13 million dollars is used in this example.

Assuming that the engines are sold under a pay-by-the-hour scheme, this would include all the maintenance work. In this example it is assumed that the maintenance includes 3 overhauls, costing 1.2 million $ each. Assuming that the expected life of the engines is roughly $20,000$ cycles, each lasting 4 hours, the price-per-hour is 41.3$.

In the case of large engines the price for a single engine is roughly 41.7 million $ [50]. Assuming the same expected service life in terms of cycles as for the smaller engines, but with each cycle lasting 8 hours, the price per hour is roughly 260$.

2.4.2. Impact of performance scatter on finance

Under the conditions of Rolls-Royce TotalCare ®the airline only pays a pre-agreed pay by flying hour fee and the manufacturer takes care of the maintenance costs. This gives peace of mind to the airline, and incentivises the manufacturer to ensure that the engines operate reliably and provide the performance agreed upon in the contracts. It also means that the maintenance costs must be known beforehand in order for the contracts to be negotiated correctly.

The largest part of maintenance costs are the engine overhauls, which are expensive. Typically the costs of an engine overhaul are in the range of 1.2 to 2 million $ [1], [51], whereas the amount paid out by the airline for a single large engine over the course of its life is about 41.7 million $ [50]. This means that a single overhaul can account for roughly 5% of the amount received by the engine manufacturer. It is therefore very important that the number of expected overhauls required is projected accurately for a long term. Unexpected engine overhauls represent a loss of profit for the company and in some cases potential overall loss on a contract.

The main idea in setting up the contract is therefore to place its end date such that it corresponds to the expected end of life for the engines. One method for choosing the exact end date could be to ensure that all engine lives fall beyond that point. This can be done by calculating the end of contract as being before the mean of the fleet life by 6 standard deviations of its scatter, which is illustrated in figure 2.4.

Assuming the standard deviation of the intermediate lives as 250 cycles, the end of life scatter, using 2.3, is 500 cycles. The six standard deviations offset of the end of life contract is therefore 3000 cycles. When using the first approach to estimate the losses, the number of unused cycles is the difference between the mean total engine lives and the end of contract point. As all engines fall after this point this difference is assumed to equal the 6 standard deviation range.

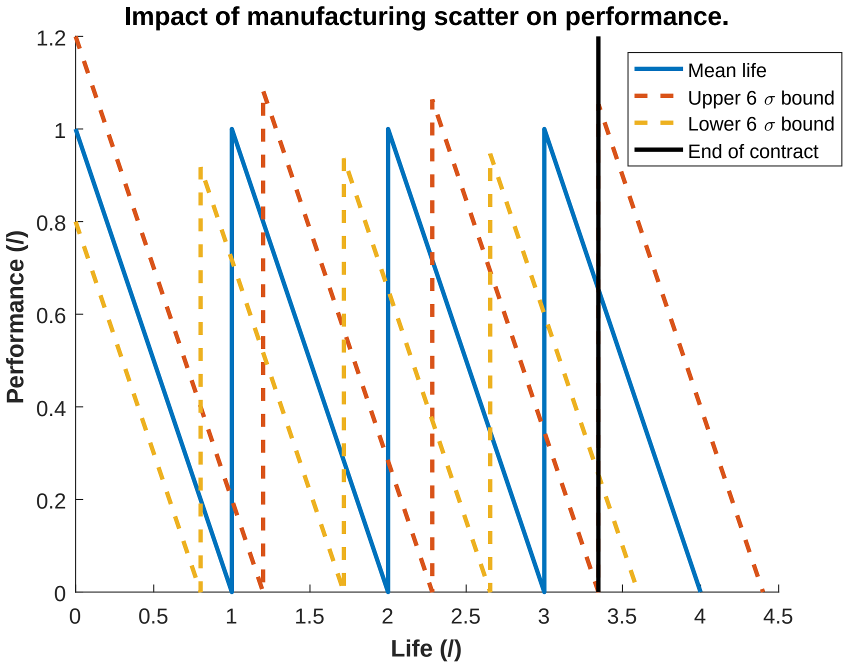

Figure 2.4: Theoretical scheduling of end of fleet support contract.

The cost of these cycles can be calculated using equation 2.4, where V is the value, $n_{engines}$ is the number of engines covered by a particular contract, n_{sigma} is the number of standard deviations used in the setting of the end of contract date, $n_{h.percycle}$ is the number of hours the engines operate for in one cycle, and p is the price-per-hour payable by the airline as per the contract.

$$V = n_{engines} \cdot n_{sigma} \cdot \sigma \cdot n_{h.percycle} \cdot p \tag{2.4}$$

For the case of a small business jet engine fleet consisting of 300 engines, the loss due to the performance scatter, under the assumptions and cost estimations, would amount to 148.68 million % . It can be seen that this cost scales directly with the scatter and therefore a 1% decrease in scatter is a 1% decrease in cost, in this case roughly 1.5 million $.

For the case of a large engine fleet also consisting of 300 engines, the loss due to the performance scatter, under the assumptions and cost estimations, would amount to 1.872 billion % . Again, this cost scales directly with the scatter and therefore a 1% decrease in scatter is a roughly a 18.72 million $ reduction in losses.

2.4.3. Impact of deterioration rate on finance

Another way to improve the profitability of the engines is to increase their service life. This can be done by improving the initial performance margins of the engines, or by improving their deterioration rates. Both of these can be achieved by altering the aerodynamic design of the turbomachinery components. Assuming that the blades are already optimized for maximum initial performance, the improvement of the deterioration rates should be optimised next.

This requires the performance of the deteriorated blades to be increased. As the blade geometries are expected to change with operation the improvement of deterioration rates is likely conflicting with the goal of maximising the initial margins. However, as long as the reduction of life due to smaller margins is less than the improvement of life due to improved deterioration rates, the trade-off is worthwhile.

Since deterioration was assumed to be modelled as a linear function of in service time, the improvement of life due to improved deterioration rates can be expressed using equation 2.4.3. In this equation ΔL is the increase of life gained through the change of the deterioration rate, $k_{initial}$ is the initial deterioration rate, $k_{improved}$ is the improved deterioration rate, n is the ration of the rates, M is the initial performance margin, and $L_{initial}$ is the initial expected life.

$$\Delta L = M \cdot \left(\frac{1}{k_{improved}} - \frac{1}{k_{initial}} \right)$$

$$with \quad k_{improved} = n \cdot k_{initial} \quad and \quad L_{initial} = \frac{M}{k_{initial}}$$

$$\Delta L = L_{initial} \cdot \frac{1-n}{n}$$

Equation 2.4.3 shows that the improvement of life changes with $\frac{1}{n}$. As n is expected to be close to 1, a 1% reduction of the deterioration rates therefore roughly translates into 1$ increase of life. The increased revenue due to this increase of life can be calculated using equation 2.5.

$$V = (1-n) \cdot n_{engines} \cdot L_{initial} \cdot n_{h.percycle} \cdot p \tag{2.5}$$

Applying the price estimates from subsection 2.4.1, the improvement of life due to deteriorated rates being improved by 1% equals roughly 9.9 million $ for business jet engines and 124.8 million $ for large engines.

2.5. Naming convention

To support the discussion the features of the compressor blade design must first be named. This is done for the design that is being evaluated in the subsequent work, namely the BR715 R3 compressor blade. The naming convention is established here in order to improve clarity of the arguments presented regarding the methodology implemented and results obtained.

The BR715 R3 compressor blade is shown in figure 2.5. This figure shows the Computer Aided Design (CAD) design of this particular component and the orientation of the coordinate system within which it is defined.

First, it is important to determine the coordinate system within which the CAD is located and name its axes. A Cartesian coordinate system with the orientation displayed in the figure is used to define the blade. Its x-axis is collinear to the engine axis of revolution. This direction is therefore named the axial direction. The y-axis points tangentially in the direction of rotation under operation and therefore this direction is called the tangential direction. The z-axis points along the span of the blade and is therefore called the spanwise direction.

In general, this CAD can be split into three major features, each with its own purpose. These purposes are:

- To hold the component in place within the assembly
- To form the lower geometry bound of the gas path annulus
- To impart the energy onto the flow

It is therefore convenient to name the features that complete these tasks individually. These are the axial root, platform and the aerofoil, marked as features 1, 2 and 3 in figure 2.5 respectively. To avoid confusion when referring to the “blade”, the entire component is hereafter named the blade, and the aerodynamic “wing” located on the platform is called the aerofoil.

Intuitively, the root positions the blading element in place by use of the dovetail surfaces. The platform bounds the gas path and provides a surface onto which the blade is attached. Near the front of the blade there is a groove on the lower side of the platform.

The aerofoil itself has several features. The location, along the spanwise direction, where the aerofoil is attached to the platform is called the hub. On the opposite end, the location where the aerofoil ends is called the tip.

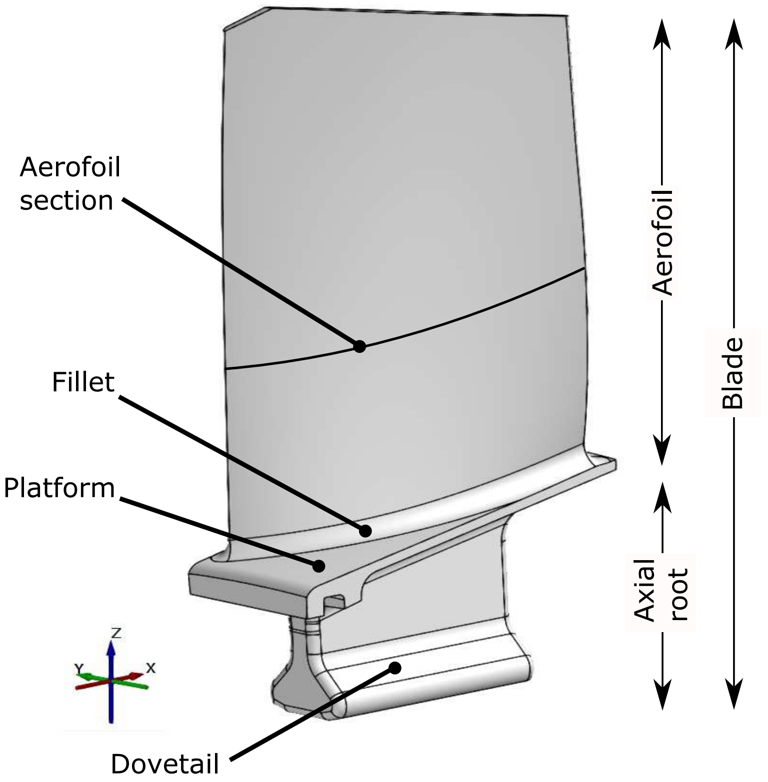

Figure 2.5: The features of the name and their naming convention.

The aerofoil geometry is a complex 3D shape that changes continuously from the hub to the tip. The cross section of the aerofoil at any specific spanwise position is called the aerofoil section.

This CAD shown in figure 2.5 is used to support the manufacturing process and the blades are manufactured according to its dimensions. As the x-axis and the engine axis of revolution are collinear the CAD is already located in the spanwise position that it assumes in the final assembly.

3

Review of probabilistic methods for quality control of compressor and turbine blades

Before attempting to take advantage of the opportunities outlined in chapter 1, their potential benefits need to be quantified. After all, if these possible benefits are not substantial, then taking further steps to take advantage of the available opportunities is not worthwhile. The quantification of the impact of geometrical uncertainty on component performance is therefore the first step to be taken.

Past work performed on performance uncertainty quantification of compressor and turbine blades is presented in section 3.1. The aim of this section is to show that the blades do have geometrical variability and that this causes significant performance variability. Lastly, it provides an overview of methods already developed to characterise compressor and turbine blades, and the methods to quantify their geometric variability in terms of performance.

Provided that the available opportunities are substantial the next step is to take advantage of them by updating the existing designs to robust designs. The past work performed on robust optimisation of compressor and turbine blades is presented in 3.2. This section presents the separate ways blades could be optimised for robustness to geometrical variability.

The conclusions drawn from the literature review are then given in section 1.2. This section summarises the progress made so far and opportunities available for future work. Furthermore, the most frequently used methods are highlighted.

3.1. Uncertainty quantification of compressor and turbine blades

Compressor and turbine blade performance are affected by many sources of uncertainty. Initially, some uncertainty is present during the manufacturing process, which affects the initial performance of the individual blades. Secondly there is some uncertainty in the inlet conditions of the engine that cause different operational performance. Lastly, the blades undergo a deterioration process during the course of their lives, which further changes their performance. These are the uncertainty source that need to be quantified.

In this section an overview of literature on the application of uncertainty quantification methods on aerodynamic performance parameters of compressor and turbine blades is presented. First the applications pertaining manufacturing uncertainty are presented in section 3.1.1, followed by the applications on deteriorated blades in section 3.1.2.

3.1.1. Applications for manufacturing uncertainty

Several approaches for quantifying manufacturing variability of scanned blades have been developed recently. Garzon et al. [16] used PCA to analyse radial station measurements of 150 compressor blades of the same stage taken from 4 different engines. The variables used in the PCA were the measured aerofoil coordinate distances from their corresponding mean aerofoil coordinates. The obtained PCA modes were reported to resemble classical NACA aerofoil parameters. No clear most important mode was identified. A 2D CFD MC assessment has then been performed based on these modes to determine the effect of the geometric variability. Performing an evaluation using five times the geometric variability that was observed in the actual measurements, the authors reported a 1% reduction of compressor efficiency relative to the nominal design to manufacturing variability.

The same procedure has been applied to turbines as well. Duffner [12] used PCA to characterise manufacturing geometrical variability of Coordinate-Measuring Machine (CMM) measured turbine vanes. The measurements were taken from the midspan of the blades only. Based on these profiles a 2D CFD sensitivity analysis was performed. The author reported that there is a high probability that geometrical variability leads to unacceptable losses. The modes were not interpreted in more intuitive geometrical parameters.

Lange et al. [37] have used PCA to analyse scanned 3D compressor blades and performed 3D CFD MC assessment of the effect of manufacturing scatter. They have reported that the most important principal components resemble classical profile parameters which are more intuitive to understand. The total pressure ratio was reported to be highly correlated with stagger angle and maximum camber of the blades.

Lange et al. [35] have developed a method for extracting 14 classical NACA aerofoil profile parameters from 3D scans of compressor blades. Heinze et al. [21] used the method developed in [35] on 500 3D scanned turbine blades and extracted the NACA parameters from the scans at 61 radial stations. This resulted in 854 parameters (61 stations x 14 parameters), which is large. They then demonstrated that a spanwise correlation between the parameters at individual sections exists. Specifically, it was shown that the parameters of individual blades are more or less constantly offset from the nominal values at all stations along the span. This allowed the number of parameters required for the characterisation of the blade to be reduced from 854 to just 14, which is tractable for probabilistic evaluations.

The histogram of one of these parameters, namely the maximum thickness, averaged over the entire blade span is show in figure 3.1. This figure clearly shows a scatter that could have an impact on aerodynamic performance.

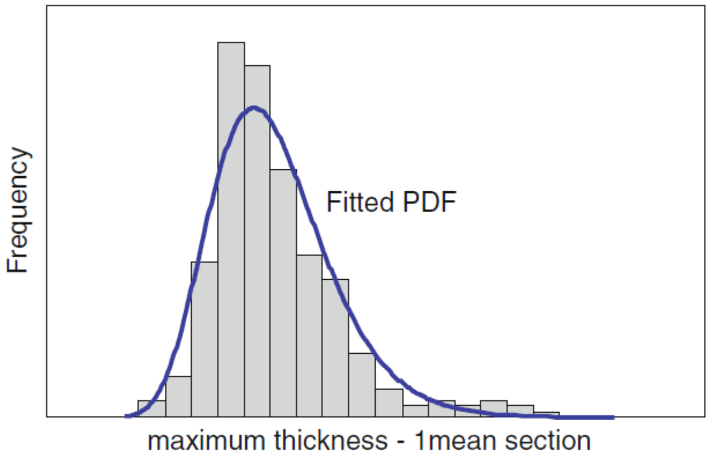

Figure 3.1: A histogram of max thickness of 500 scanned blades and its corresponding PDF [21].

Within the scope of Heinze et al. [21] aerodynamic evaluation of these blades was left as a next step. To address this step, Lange et al. [36] scanned 150 high pressure compressor blades. Using methodology from Lange et al. [35] they extracted the NACA parameters from the scans. The parameters were then characterised by a multivariate PDF. A sample of 500 new parameter sets was extracted from this PDF and used as a basis to build randomized realisations of the original blades. Based on these blade realisations 3D CFD was performed in order to assess the performance sensitivities in respect to the NACA geometrical parameters.

Figure 3.2 shows the pressure coefficient plots based on the results of 3D CFD of realised blades for a particular span-wise section. The pressure coefficient distribution for individual realised blades are plotted on top of each other in grey to illustrate the overall scatter in the results. Additionally, two realisations are plotted in red and black to highlight the differences between individual realisations. The difference between these two highlighted blades is significant in the first quarter of the chord. This different behaviour and its corresponding different performance is a direct consequence of geometrical variability.

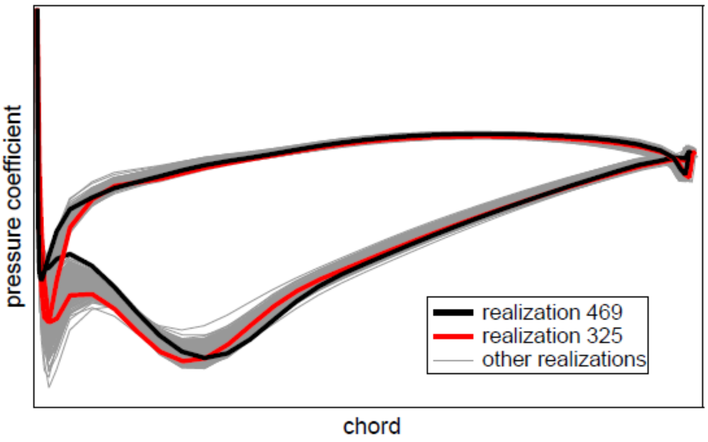

Figure 3.2: A comparison of the coefficient of pressure distribution of uncertain aerofoil geometries [36].

Lastly, Lange et al. [36] computed the sensitivities of performance parameters to the geometrical parameter variabilities. The isentropic efficiency was most sensitive to thickness related parameters, with the Leading Edge (LE), maximum and Trailing Edge (TE) thickness the three most important parameters. The loss coefficient was most sensitive to these same parameters, in the same order. The last performance parameter evaluated was the blade turning which was most sensitive to the stagger angle and the maximum camber of the blades. Lastly, it was reported that these sensitivities are in agreement with the results of Garzon et al. [16]. The similarity of both studies suggests that both the PCA and parametrisation approaches are equivalent in terms of identifying performance sensitivities with regards to geometrical parameters.

In Lange et al. [36] the authors assumed that the impact of geometric uncertainty on performance of a single blade are representative of the entire rotor. Obviously this is not truly representative of reality, as the basis for uncertainty quantification of blades is that all the blades have slightly different geometries and therefore slightly different performances. To remove this assumption Lange et al. [38] simulated up to eight passages composed of blades with individually generated geometries. A part of the simulation results presented in Lange et al. [38] is shown in figure 3.3.

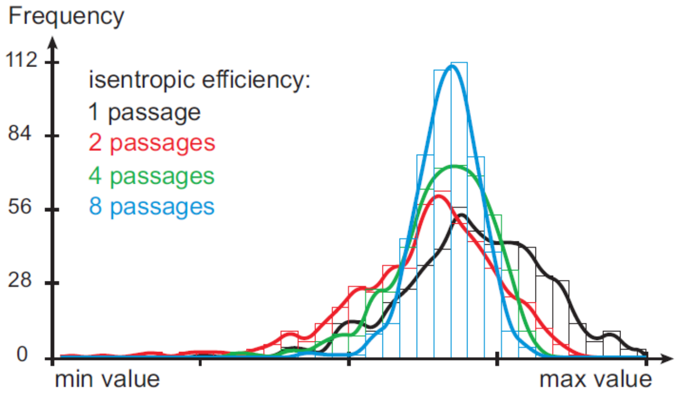

Figure 3.3: Isentropic efficiency histograms and their corresponding PDFs for CFD models containing several different numbers of passages. [38].

Figure 3.3 shows the histograms and their PDFs for simulations using 1, 2, 4, and 8 passages. Note that a comparison of the mean performance of the different multi-passage results can not be done on the basis of figure 3.3 as the values of the histograms are scaled between their respective minimum and maximum values.

From figure 3.3 it can be seen that the performance scatter reduces with the increase of the number of passages simulated. In contrast to the performance scatter decrease, Lange et al. [38] report that the performance offset converges to a value other than the results of the single passage model. As the blades have geometry that is completely independent from one another it is more likely that models using few passages will be comprised of only "good" or "bad" blades, which increases the scatter. Models with several passages will most have a mix and therefore the scatter would be lower.

Lange et al. [38] also report that the results of this work, where the blade geometries representing an entire stage were independent of one another, can be contrasted with the results of Lange et al. [36], where they were not. These results and methods can therefore be used to determine the lower and upper bounds of the performance scatter that could be observed in reality.

Expanding on that work, Lange et al. [39] investigated the impact of manufacturing variability on multiple stages at once. They performed 3D MC CFD on a 10 stage compressor represented by a single rotor stator pair for each of the stages. The manufacturing scatter was characterised for three distinct stages. For the rest of the stages the scatter was interpolated between the values of the evaluated stages. The final geometry of the generated compressor realisations was then governed by 140 PDFs (14 parameters x 10 stages) that were independent from stage to stage. The realisations were built based on samples obtained from the PDF's using a Latin-Hypercube sampling and used in a 3D MC simulation. To represent the speed line seven operating points were used in the CFD assessment.

The results were the sensitivities of global parameters such as mass flow and isentropic efficiency to geometric parameters and the sensitivities of individual stage performance uncertainties to the overall compressor performance. They suggest which of the geometric parameters in a full compressor and which stage uncertainty in a full compressor are the most important respectively.

The results showed that front stages had larger impact on the overall performance than rear stages. Furthermore, with more throttled operating points the front stages became increasingly important. For the initial operating point the most influential parameter was the leading edge thickness. For more highly throttled operating points this changes however and the most important parameter became the maximum camber.

The MC requires a high number of CFD evaluations to be performed. In order to reduce the required computational time Schmidt et al. [56] have compared the MC approach against a less computationally intensive adjoint approach. This was done by comparing the sensitivities of performance parameters to geometric variability of compressor blades obtained by both methods. The MC method relies on evaluating CFD for several blades based on randomly generated parameters. Based on these evaluations the sensitivities of performance to geometric parameters can be determined. This is done by a polynomial regression of the results. The blades were characterised with the NACA like parameters.

Alternately an adjoint CFD evaluation of a blade can be performed to obtain the partial derivatives of performance parameters to the geometrical parameters. These partial derivatives can then be seen as a linear model of performance. Combining this model with the actual variability of geometrical parameters their impact on the actual performance can be determined. The main advantage of this approach is the reduced computational time as only one CFD evaluation needs to be performed.

However, for this approach to be valid the assumption that the performance varies linearly with the geometrical parameters must be valid. The initial adjoint CFD from which the partial derivatives are obtained is based on a mean blade.

The sensitivities of isentropic efficiency η_{is}, pressure ratio π_t and turning angle $\Delta\beta$ obtained using these methods are shown in figure 3.4. The x-axis of the plots features the sensitivity parameter which is calculated by multiplying the standard deviation of the geometrical parameter with the partial derivative of the specific performance with respect to the geometrical parameter.

Figure 3.4 shows that both methods roughly agree with each other. The agreement is best for the case of pressure ratio π_t where both the first and second most important parameters of both methods are the same even though the actual sensitivity values are not the same. In the case of turning angle $\Delta\beta$ both methods identify stagger angle γ as the most important parameter, but do not agree exactly how important it is. Furthermore,

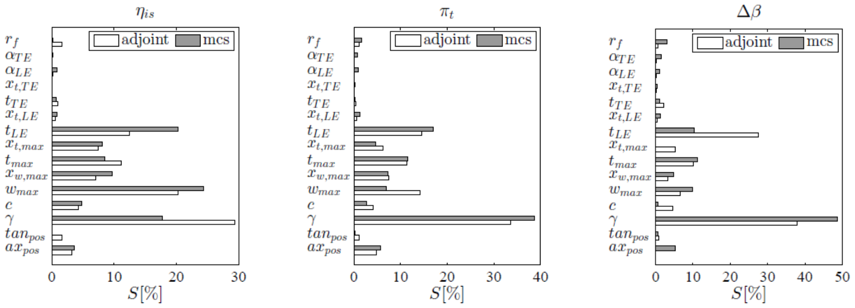

Figure 3.4: Sensitivities of performance to geometrical parameters calculated using the adjoint and MC-type methods. [56].

no agreement is reached for the second most important parameters. Lastly, for the case of isentropic efficiency η_{is} none of the three most important parameters are in agreement.

Due to the observed differences in the calculated sensitivities it can be concluded that the methods are not interchangeable. As the MC method relies on less assumptions it stands to reason that its results are more credible.

The approach of using NACA parameters for characterisation of compressor blades has also been applied to turbine blades. Scharfenstein et al. [55] analysed scans of 500 new and used turbine blades to obtain their respective NACA aerofoil parameters. Based on PDFs of these parameters they performed a 3D MC CFD evaluation. A plot of the results can be seen in figure 3.5.

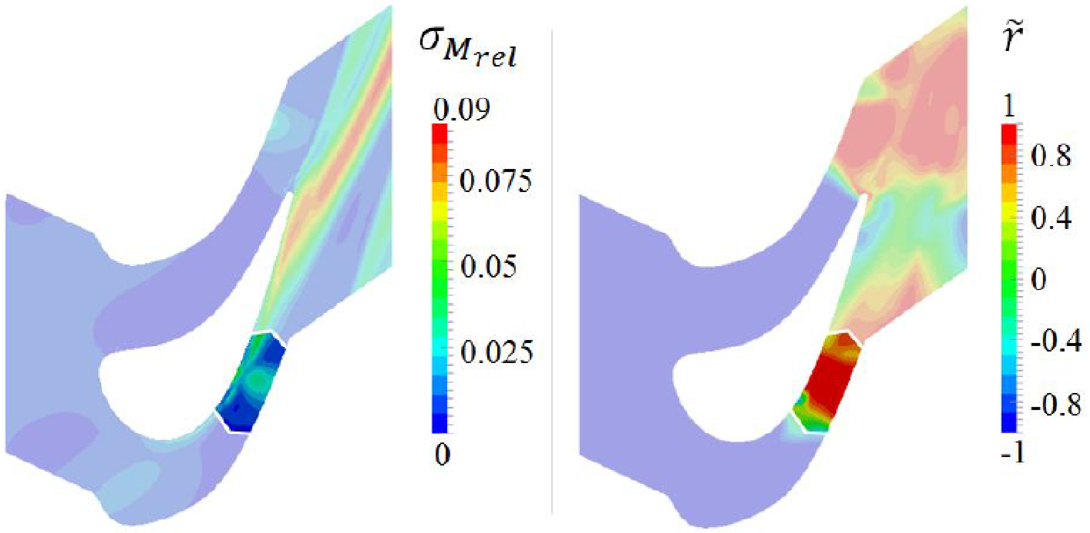

Figure 3.5: Statistics on Passage of the relative Mach number and its correlation with stagger angle γ. [55].

Figure 3.5 shows a plot using a visualisation technique called Statistics on Passage. In this case this technique plots the statistics associated to all points of a 2D grid section of a 3D CFD solution. The statistics are calculated among the solutions of the different blades. On the left side of figure 3.5 the standard deviation of the Mach number is presented. On the right side the correlation of the local relative Mach number with the blades stagger angle γ is shown. The blades are visualised in white in both of the sub-plots. The figure also highlights a region on the suction side of the blade at approximately half chord location.

On the left sub-plot this region shows moderate standard deviations which are a consequence of geometrical variability. On the right sub-plot it is shown that these are strongly correlated with the blades stagger angle γ. Furthermore, Scharfenstein et al. [55] reported that the turbine rotor could choke due to the impact of geometric variations. Note that the geometrical variability consists of both manufacturing variability and deterioration. The impact of these individual sources of uncertainties was not separated within this work.

Similar work has been performed by Kolmakova et al. [27] who performed a 3D quantification of turbine blade performance with uncertain geometrical and inlet and outlet air parameters. The blade geometry description was limited to stagger angle, chord length, blade height, and tip radius. The uncertain air parameters were the inlet total pressure and temperature and the outlet static pressure. Performance was characterised through loss coefficient, mass flow and the outlet flow angle. The assessment showed that the uncertainties in the air parameters had no significant effects on performance. On the other hand, geometry variations had a significant impact on performance of turbine blades. Stagger angle and displacement of the trailing edge caused the largest impact on the performance parameters, up to 4% and 20%, respectively.

Turbine blades have a different geometry than compressor blades as they typically exhibit cooling holes and slots in them. It is important to take these into account when characterising the blades with the NACA parameters and subsequently analysing them. Högner et al. [22] analysed scans of 20 turbine NGVs that featured cooling holes and a trailing edge slot. A special method was presented that removed the cooling holes from the geometry before the calculation of the aerofoil parameters. The trailing edge slot was characterised by its own specific parameters such as the trailing slot position, thickness and height. Again, these were combined in a multivariate PDF from which randomised samples could be drawn.

The turbine blades used in this work were cast as a pair, which needs to be addressed in the simulation. Therefore two individually created blades were joined in a pair. This was done to keep the dimensionality of the PDF to a minimum and therefore the number of samples and CFD calculations needed to quantify the performance. The 3D CFD was then performed on the generated blade pair. To expand the results of a single blade pair to the entire stage a so-called synthesis approach was implemented. This approach linearly superimposes individual blade pair results to give the full stage results. The use of this method has been supported due to the linear behaviour of the results and the absence of cross-effects in the solutions.

The cooling holes and the trailing edge slot on turbine blades introduce cooling air to the gas path. Within the CFD setup this was modelled as a line source on the surface of the blade. The output parameter of interest was the relative capacity change. As in the cases of Lange et. al. [38] and Scharfenstein et. al. [55] the size of the performance scatter was decreased for the full stage results compared to the single blade pair results. Furthermore, it was noted that a possible explanation for the relatively low geometrical scatter observed is the low number of scanned blades used in the evaluation.

Lastly, Högner et al. [22] reported that parameters that somehow affected the passage cross sectional area between turbine blades had the largest impact on performance. These were the profile offset at the suction side and the trailing edge cooling slot position.

The assessment of the impact of uncertainty of film cooling has been performed by D'Ammaro et al. [9]. In this work the authors quantified the impact of uncertain cooling passage inlet conditions on the cooling adiabatic effectiveness. The results showed that a 40% variation in the inlet conditions can reduce the turbine blade life by 5 times. Even though the work focuses on uncertain inlet conditions it can be seen that variability of the cooling geometry could have a potentially large impact on the blade life as well.

Quantification of turbine tip gap and fillet radius geometrical variability has been performed by Montomoli et al. [46] using 3D CFD. The tip gap and fillet radius were modelled using a two dimensional Gaussian distribution. The parameters of the distribution were selected arbitrarily and did not originate from measured data. The results were reported as response surfaces and showed that relatively small geometrical variations can cause significant changes in the flow field and therefore performance. Both the tip gap and the fillet radius were found equally important.

The geometrical variability has impact not only on the performance of the individual engine components, but also on the entire engine. Gorelik et al. [19] have created a probabilistic whole engine model that is capable of taking into account the variability of real components. This model allows to evaluate the impact of component geometrical variability on whole engine parameters such as efficiency. An example application to 3D centrifugal compressors has been shown. In this assessment only the blade angles at the compressor inlet and thickness along the shroud line were varied. Therefore only global changes in geometry could be quantified, whereas the local effects could not be captured.

3.1.2. Applications for deterioration uncertainty

First assessments of deterioration focused on a holistic approach backed by operational data. Lakshminarasimha et al. [34] describe a stage stacking approach to whole engine deterioration modelling based on empirical operational data. Zaita et al. [64] developed a holistic deterioration model that relied on aircraft mission profiles, environmental conditions and physical and design characteristics based on the stage stacking approach. A current example of holistic deterioration prediction method has been developed by Spieler et al. [59]. This model is capable of taking into account specific airline mission profiles and using correlations derived from operational data to predict deterioration.

As opposed to holistic modelling, current trends attempt to model deterioration on a component level. Kumar et al. [28] simulated erosion on a 2D compressor aerofoil to show the impact it has on performance. The erosion was simulated as a series of depressions modelled using Hicks-Henne functions. The parameters of these functions were selected by the user and were not obtained from real blades. These erosion patterns were then applied to the pressure side of the blades, as shown in figure 3.6.

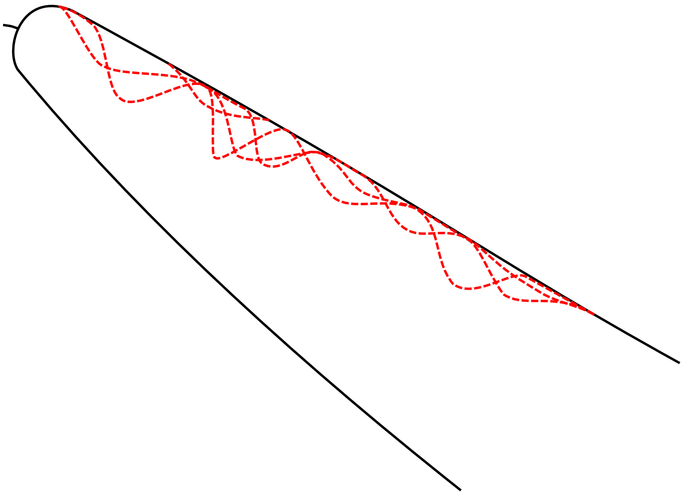

Figure 3.6: Erosion patterns constructed on the 2D aerofoil pressure side. [28].

To expedite the calculations a surrogate model was constructed and validated to approximate the flow equations based on the results from the HYDRA flow solver. The analysis of blade erosion showed that it could cause an upto 5% increase in pressure loss for the blade in question.

Goodhand et al. [18] showed that compressor blades susceptible to switches between flow mechanisms, such as separation, are most sensitive to geometry variations in 2D. Here the Hicks-Henne functions were again used to simulate blade erosion. The function parameters were also selected by the user and were not obtained from real blades. The results showed that blades with sharp leading edges are susceptible to leading edge separation as a consequence of geometrical variability.

A quantification of fouling in 2D has been performed by Qui et al. [52]. The Hicks-Henne functions were once again used to simulate the dirt on a 2D compressor aerofoil. Once again, the parameters governing the functions were not based on real blade data. The results showed that performance losses increased monotonically with the amount and size of the fouling.

In 3D, Marx et al. [40] have scanned 1400 variously deteriorated and 300 newly manufactured blades from 10 different stages and compared their geometries. The blades were then parametrised similarly to the process described in Lange et. al. [35]. The parameters of newly manufactured and deteriorated blades were then compared to each other. The authors reported that deterioration increased the geometrical scatter of the blades. An example plot for the leading edge thickness can be seen in figure 3.7.

Figure 3.7 shows 10 color coded vertical bands that each represent blades of a particular stage. The bands are scaled to have their heights in proportion to the real blades. Furthermore the passage contours are drawn in with white as well as some grid lines to support visual comparison. The colours of the band represent the scatter of the particular geometrical parameter for the particular stage and location on the blade. Note that the

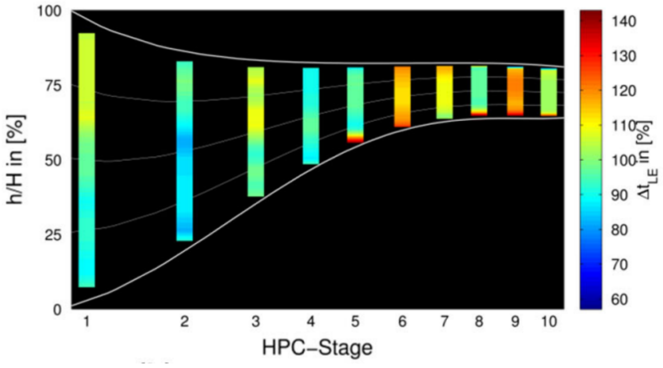

Figure 3.7: Change of the leading edge thickness due to deterioration [40].

color scale is given in percent of the reference blade values and ranges from 60% to 140%. Therefore any value below 100% is a reduction in that parameter and vice versa.

From figure 3.7 it can be seen that deterioration sharpened the front stage leading edges while it blunted the rear stage leading edges. Furthermore, the authors report that over all the parameters the increase of scatter was small for front stages and larger for the rear stages. As the analysis also included blades that have been previously repaired these effects were also included in the results.

Reitz et al. [53] have developed a process for evaluation of performance of deteriorated compressor blades using 3D CFD and showed its use on an example blade design. This work followed the approach described in Lange et al. [35] to parametrise the blades and to rebuild blades based on an arbitrary set of geometrical parameters. The deterioration variability came from a set of 120 scanned blades. The rebuilt blade geometries were then evaluated using the CFD tool TRACE. The difference of performance between maximum and minimum Mach numbers at a 2D section at 95% of the blade height can be seen in figure 3.8.

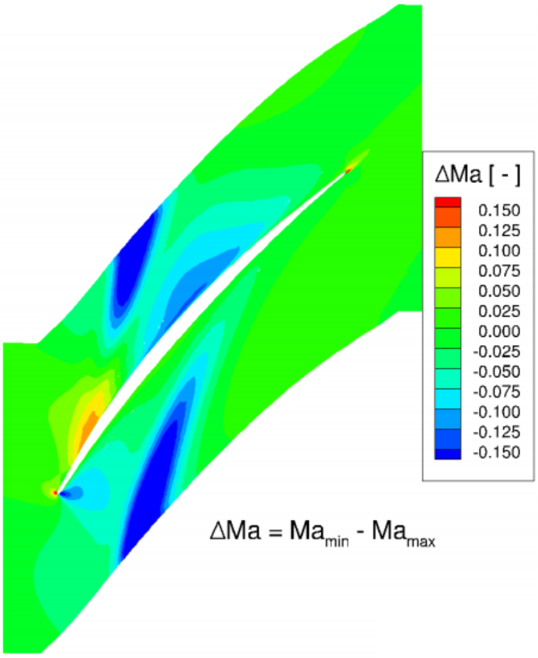

Figure 3.8: Difference between the maximum and minimum Mach number of realised deteriorated blades for a 2D section at 95% of the blade height[53].

Figure 3.8 shows that for the geometry uncertainty and the selected set of boundary conditions the variation in Mach number was up to $\Delta M = 0.15$. One of the conclusions reported was that deterioration changes the stage performance noticeably. Note that this work did not take into account blade roughness and deteriorating tip height. Furthermore it did not attempt to decouple the effects of manufacturing variability from deterioration, therefore the reported effects were a combination of manufacturing variability and deterioration.

Work done so far on quantifying compressor deterioration on a blade level relied either on arbitrary deterioration modelling or included the effects of both manufacturing variability and deterioration. In the case of turbines no previous work that would address purely the quantification of the effect of deterioration has been found.

3.2. Robust optimisation of compressor and turbine blades

Uncertainty quantification of performance as given in section 3.1 reveals what the impact of the uncertainty is. Essentially the magnitude of the impact directly translates into the magnitude of the opportunity for using robust optimisation.

Provided that the impact of uncertainty is large, the natural next step is to use the information of the impact in the design phase. In the current scope the uncertainty relates to two processes, manufacturing and deterioration. For the purposes of optimisation these effects must therefore first be separated as they might have different impacts.

After the effects of manufacturing variability and deterioration have been properly separated, modelled and evaluated in terms of performance, the blades can be optimised by including these effects as well. This optimisation can aim to achieve various goals. By optimising the tolerancing of blade surfaces the performance deviations of a particular nominal design can be improved. This is done by recognising the most important blade sections and reducing their tolerances. This however increases manufacturing costs, which can possibly be offset by reducing the tolerances in less important sections. Alternately the blade shape can be changed to provide improved results over the original nominal design. In this work it was assumed that the manufacturing was already optimised economically, therefore the optimisation of the tolerancing was not considered.

This section gives an overview of literature that deals with robust optimisation of compressor and turbine blades. It is split into two parts that each deal with one of the uncertainty processes. Both sections first discuss the applications on compressor blades and follow with the discussion of turbine blades. For both of these parts, the work performed in 2D is presented first. Based on this, the work in 3D is presented next.

3.2.1. Optimisation of blades for uncertainty in manufacturing

Kumar et al. [29] have designed a 2D compressor aerofoil robust to manufacturing variability. The Hicks-Henne functions were used to parametrise the entire aerofoil, with 5 function describing the suction and pressure sides each. The amplitudes of the Hicks-Henne functions are then used as the design parameters which results in the aerofoil being described by 10 parameters. The bounds for these parameters were specified by the user.

Based on the geometries obtained with this process an optimisation of the aerofoil was performed. The aim of the optimisation was to minimise the standard deviation of the normalised pressure increase over the blade while maximising its mean. The normalized pressure is the difference of the total pressures at inlet and outlet normalised by the inlet total pressure and multiplied by 100. The results of the deterministic optimal design and the robust design are shown in figure 3.9.

Figure 3.9 shows two histograms overlapped that represent the performance scatter of two designs. Furthermore two vertical lines are drawn which denote the deterministic performance of the two designs. This figure shows that a robustly designed aerofoil would be significantly less sensitive to manufacturing scatter. Furthermore the robust blade had better mean performance, lower performance scatter and better worst case performance than the deterministic design, at the expense of worse nominal design performance. Because the actual performance of the blades is the performance after manufacturing is taken into account, the worsening of nominal performance is not significant in this case. Note that mean performance does not correspond to the peaks in the histograms in this case as both have a strong right tail.

In subsequent work Kumar et al. [30] have extended the method developed in Kumar et al. [29] to include the information on the process capability into the assessment. The process capability data is used to effectively

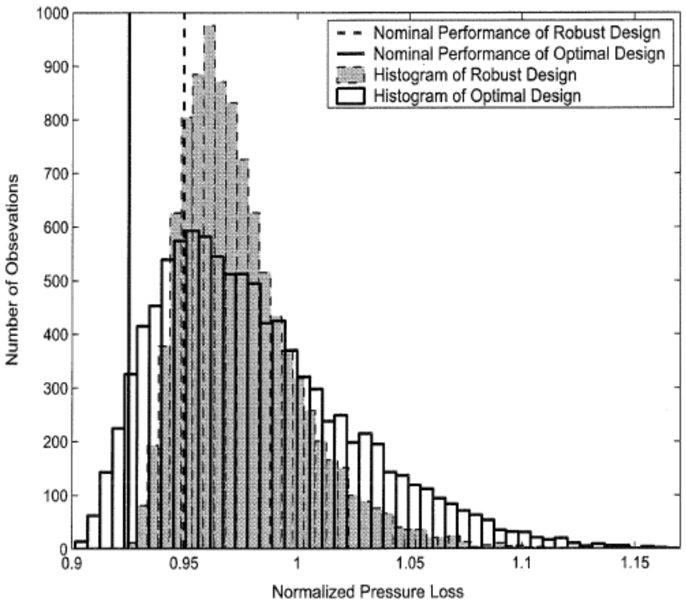

Figure 3.9: Comparison of normalized pressure loss coefficients of deterministically and robustly designed blades [29].

determine the bounds of the parameters of the Hicks-Henne functions. As this work is just an adaptation of the inputs of Kumar et al. [29], the results of both studies are in direct agreement with one another.

Similarly, Dow et al. [11] have performed an optimisation of 2D compressor aerofoils taking into account manufacturing variability, and simultaneously optimised the manufacturing tolerances of the blade. The blade was simulated by a vector of normal differences between the manufactured and nominal geometries, which was then modelled using a multivariate normal distribution. Using this setup the tolerancing scheme of the blade can be altered by changing the covariance of the multivariate normal distribution. Parameters of the multivariate distributions were not based on measurement data. The geometries generated with this process were evaluated by a 2D solver called MISES. The results showed that the leading edge geometry was the most important and the tolerances in this region were reduced in the optimal design.

For the case of turbines, Moeckel et al. [42] presented a frame work to calculate cost savings associated to various tolerancing schemes for turbines in quasi-3D. This work focused solely on the thermal analysis of the blades. The uncertain parameters were the operational conditions, such as the inflow temperatures and mass flows, the film cooling hole diameters and heat transfer coefficients. The results showed that the film-cooling hole diameter is one of the most important parameters to be correctly toleranced.

Javed et al. [24] have performed robust optimisation of 3D centrifugal impellers against manufacturing uncertainty using manufacturing capability assumptions. The design variables were several geometrical parameters. Based on the tolerances of these parameters the uncertainties of impeller performance could be quantified. The tolerances and the level of robustness were arbitrarily selected by the authors. The objective of the optimisation was to maximise the means of pressure rise and isentropic efficiency while minimising their standard deviations.

The realisations of the impeller were evaluated based on a response surface model, itself based on 3D ANSYS CFX calculations. The results showed that the impellers could be optimised to several levels of robustness as desired by the designer.

Focusing on uncertainty of operational conditions, Vinogradov et al. [61] have performed a 3D robust optimisation of a axial high pressure turbine blade. In this case the inlet conditions were subject to variability, and the dimensions of the LE film cooling rows and the tip clearance were design parameters of the optimisation.

Arsenyev et. al. [4] have developed a methodology for optimisation of 3D low pressure turbine stator vanes against geometrical variability. The method is multi-disciplinary and involves aerodynamic and thermal evaluations as well as an evaluation of mechanical stresses. An example application of a robust aerodynamic design of turbine blades was presented. In this example the blade thickness at the LE and TE, and the blade axial and tangential positions were the uncertain parameters. The optimisation focused primarily on blade reliability as opposed to aerodynamic design.

Bestle et. al. [6] have proposed a quasy-3D approach to robust design of compressor blades. The quasy-3D approach means that the blade is represented by several 2D aerofoil sections, which can be stacked together to give the final 3D geometry. The 2D aerofoil sections were parameterised using NACA like parameters, as seen in Lange et. al. [35]. These parameters are subject to uncertainty. The CFD evaluations needed for the robust optimisation were performed on the 2D sections independently to save time.

3.2.2. Optimisation of blades for uncertainty of deterioration

Kumar et al. [31] have performed a robust optimisation of 2D compressor fan aerofoils against erosion using methodology similar to Kumar et al. [29] and Kumar et al. [30]. Erosion was simulated as depressions in the blade surface modelled by Hicks-Henne functions. The parameters of these functions were selected arbitrarily. The aerofoil shape was then optimised to have its performance robust against this erosion. The objective was to maximise the mean of the normalised pressure loss while minimising its standard deviation. The results showed significant improvements in terms of performance of the robustly designed aerofoils.

Yaping et. al. [25] have performed a robust optimisation of centrifugal compressors against fouling. Fouling has been simulated using a 3D surface roughness map. The optimised results showed improved robustness to surface roughness uncertainty while retaining good design point performance.

No other work with regards to robust optimisation of compressor or turbine blades against deterioration has been found.

4

Research question formulation

Chapter 1 gave a general introduction to the need of quality control for compressor and turbine blade performance. This has been further explored in chapter 3, which showed that the blades indeed have geometrical variability that translates into performance variability.

Based on the literature review given in section 3, the next step that should be performed is the uncertainty quantification of the impact of deterioration on aerodynamic performance variability based on blade measurements. This step has been selected as it is a prerequisite for robust optimisation activities also identified in section 3.

The aim of this chapter is to demarcate the problem framework within which to address the selected next step and to formulate an appropriate research question. The research steps required in order to answer the research questions are identified and the methodology and software to do so is selected as well.

First, the measured data and its features are represented in section 4.1. The features of the data determine what kind of knowledge can be obtained by analysing it.

Several sources of uncertainty, as identified in chapter 3.1, can be included in the subsequent performance analysis. These include the variability of the inlet air conditions. The assumptions regarding the uncertainty of the inlet conditions are given in section 4.2.

The blade measurements inherently include the effects of all sources of geometrical uncertainty. In order to improve the understanding of blade deterioration the impact of these sources must be separated. This is discussed in section 4.3.

Aircraft engines are expected to operate in a range of atmospheric and operational conditions. Some of these conditions are also required to be selected as the boundary conditions for the CFD evaluation of the blades. The conditions used in performance quantification must therefore be carefully selected in order to arrive at meaningful results. After all, assuming that the performance uncertainty is small, it makes little sense to quantify it under operational conditions that the engine will rarely or never experience in its life. The selection of these conditions is done in section 4.4.

Based on the demarcation of sources of uncertainty, the approach to separate its effects, and the demarcation of the operational conditions, the research question can be formulated. This is done in section 1.3.

Lastly, based on the best practices found in the chapter 3 the methodology used to complete parts of the answer of the research question is selected.

4.1. Overview of the dataset

This chapter gives an overview of the dataset that was used to analyse the impact of deterioration on performance. Measuring deteriorated blades is essentially measuring an outcome of an experiment. The parameters of this experiment, such as the origin of the blades, the size of the sample, and the design of the experiment, are presented in section 4.1.1.

The important features and the disadvantages of this dataset are further discussed in section 4.1.2.

The actual measurement campaign that produced the 3D point cloud data is presented in section 4.1.3. This section gives an overview of the process used and highlights the specific preprocessing that is needed before the data could be properly analysed.

4.1.1. Experiment setup

The blades used in the assessment were 36 BR715 stage 3 compressor rotor blades. They came from 3 separate engines operated by a single operator and were manufactured to the same specifications by three separate manufacturers. The manufacturing process involves first forging the blades and subsequently machining parts of the blade to the end geometry requirements.

This particular type of blade was selected as they are fitted into engines operated on commercially used Boeing 717 airframes. As these airframes fly relatively short routes it typically makes several flights per day. This means that the blades are returned into refurbishment in a shorter period of time relative to engines on larger aircraft that typically fly longer routes.

The aim of quantifying the impact of deterioration on performance is ultimately to support robust design. From the perspective of robust design it is interesting to understand how deterioration impacts different blade geometries. The ideal Design of Experiments (DOE) would therefore send several batches blades based on separate nominal design into service and measure them upon return. This would give an overview of the impact of deterioration on performance as a function of the nominal blade design. The ideal DOE however cannot be realistically be performed as, assuming that the current blades are deterministically optimal, installing blades that are suboptimal would be bad for business and therefore counter-productive.

Another important aspect of deterioration are the operational conditions. Therefore another ideal DOE could have several blade batches of the same design go into service under different operational conditions. This is mostly not practical as there are many operational conditions to consider and Boeing 717, which is powered by the BR715, has too few operators to support an extensive DOE. Therefore to isolate operational conditions from the assessment, blades from a single operator have been selected for analysis.

4.1.2. Features

The main feature of the dataset is that the blades have been measured at two time points in their life. They were first measured before being installed for the first time – these blade scans are referred to as the “manufactured blades”. This allows for an assessment of the impact of deterioration with respect to operational use.

The engines they were fitted in flew within a commercial fleet and performed the same routes in similar conditions as the other engines of the fleet. Therefore the blades experienced deterioration and wear representative of the particular fleet flying within the particular operating conditions.

The blades were then scanned again when the engines returned from service for a refurbishment – these blade scans are referred to as “deteriorated blades”. The measurements of the two states together allow the modelling of geometry changes with operational use of individual blades.

The blades are manufactured with their heights higher than the prescribed values. This is done as the tip heights are prescribed for the entire stage assembly as opposed to individual blades. Therefore, before being placed into an engine, the blades are assembled into their stage and have their tips rubbed to conform to the geometrical requirements. The manufactured blade scans were made before this process took place. Because of this, no information on the blade heights of manufactured blades can be extracted from the dataset.

4.1.3. Measurement campaign

The blades were measured using an optical 3D measuring system IVB Kolibri Flex 100 at TU Dresden. The resolution of the scanner is approximately $30\mu m$ [23]. The scanner produces an unstructured point cloud representation of the blade. Note that the scanner resolution determines the scale of the features that can be captured. For example, small holes due to blade pitting and surface roughness cannot be sufficiently captured using this resolution.

For scanning the blades are clamped into place within the measurement system manually. The resulting point clouds are therefore oriented within the same coordinate system, but occupy different positions in it. Therefore, in order to allow the comparison of individual blades they must first be individually aligned to their appropriate position within the compressor coordinate system before analysis.

Choosing the common coordinate system as the compressor coordinate system further allows the discussion of parameters to be made in meaningful intuitive values.

Within the measuring campaign the blades were clamped into place using their root geometry. A part of the blade geometry covered by the clamping geometry was therefore not captured within the measurements. The clamping geometry itself has already been removed from the measurements during the measurement campaign.

4.2. Uncertainty of engine inlet conditions

Inlet air conditions play a major role in engine performance. Its uncertainty can be separated into two contributions: the flight to flight uncertainty of the atmospheric conditions, and the flight to flight selection and execution of the flight profile. Essentially, the inlet conditions depend on the daily weather, the flight profile selected by the pilot and how well the pilot can adhere to this flight profile. This uncertainty could hypothetically be used in the design phase to optimise engine performance for the most commonly encountered operating conditions.

4.2.1. Atmospheric conditions

In aviation the daily atmospheric conditions are expressed as offsets from International Standard Atmosphere (ISA) conditions at the sea level. ISA is an idealised atmospheric model based on thermodynamic equations. Daily atmospheric conditions can therefore be split into the daily variation of the conditions at the sea level and the variation of the actual altitude profiles about the ISA model. Additional factors that are not included in the ISA model are wind, turbulence and water vapour which can impact performance drastically.

The atmospheric conditions at sea level have a more drastic impact on airline operations than their altitude variation. In certain circumstances these conditions can reduce engine performance to a point where taking off from a certain airport is not possible. Therefore in the design phase a set of worst case atmospheric conditions under which the engine still produces adequate thrust is selected. This limits airline operations to the airports that reliably experience atmospheric conditions less severe than the worst case conditions designed for.

For a particular engine design performance variability dictates how severe these conditions can be. However, as the worst case conditions are deterministically selected it is the geometrical uncertainties that govern the severity of these conditions.

The uncertainty of the conditions about the ISA model with altitude are secondary in magnitude of performance impact, but have a larger economic impact, provided that missed business opportunities due to engine performance limitations are not included. Profit oriented airlines typically do not schedule flights to airports they cannot take off from. The economic impact therefore originates from in flight fuel consumption variability due to atmospheric conditions uncertainty.

The aircraft spend most of their operating time in cruise, which is typically where most of the fuel is consumed. Cruise altitudes are typically in the tropopause where the temperature is approximately constant, therefore temperature variations are expected to be negligible.

The exact altitude is selected such that the fuel efficiency is optimised. The fuel efficiency is governed by the operating point efficiency of the engine and therefore an altitude with appropriate pressure to keep the engines running at maximum efficiency design point. Assuming that the airlines fly at the optimum altitude the cruise pressure is approximately constant as well.

Furthermore, the engines are designed to operate in a wide range of atmospheric conditions and therefore their efficiency is not expected to be sensitive to very small deviations in atmospheric conditions. As the uncertainty of atmospheric conditions in cruise is expected to be very small it can be assumed that it will have no major impact on performance and therefore airline economy.

4.2.2. Airline operations

Different airlines will have separate preferences regarding flight profiles. One airline might prefer to climb with maximum allowable power to reach maximum fuel efficiency cruise altitude as soon as possible. Another airline might prefer to climb with reduced power to trade the extra fuel burned to preserve engine life and minimise maintenance costs.

Furthermore, different airlines operate in different regions of the world and have different route structures and business models, which all impact the decisions regarding flight profile selection. For example, airlines operating in predominantly cold environments are interested in cold weather performance and vice versa.

Lastly, during flight the pilot might be asked to deviate from the selected flight profile due to ATC requirements. To take this into account the world aircraft fleet would have to be modelled to determine the occurrence of such requests.

To take the uncertainty of airline operations into account an accurate estimate of the future engine customers behaviour as well as the world flight routes would have to be made. This is practically impossible and therefore this source of uncertainty is excluded from the assessment.

4.3. Separating manufacturing from deterioration geometrical uncertainty

The real issue in determining the impact of deterioration is that the deteriorated geometry also invariantly includes the manufacturing variability. The challenge therefore lies in separating these. This can theoretically be done in two ways:

- A set of blades can be scanned when being assembled into an engine and than scanned again when being removed from the engine during an overhaul activity.
- A large set of blades can be scanned when being assembled into an engine and another, unrelated set of blades can be scanned when being removed from the engine during an overhaul activity.

The first approach allows to directly determine the impact of deterioration as the differences between both scans of the same blade. The penalty for this is the time it takes for the engine to return from service to perform an overhaul.

The second approach allows the data to be gathered quickly as there is no need to wait for the blades to return from service. However this mean that the impact of deterioration can only be presented as a change of a particular performance statistic, such as its mean or standard deviation. Moreover, it is not possible to completely separate manufacturing effects from deterioration effects, as the manufacturing variability also changes with different blade batches.

The most important reason for this would be the fact that the tooling used in manufacturing the blades itself deteriorates with use. Therefore the blades manufactured when the tooling was new could be significantly different from the blades manufactured when the tooling is about to be changed out. This effect can introduce statistically significant differences between batches of blades of a single manufacturer.

Secondly, the same type of blades could be manufactured by several different companies that use different tooling. This could result in statistically significant differences between the manufacturers as well. This would then further skew the differences between the scanned manufactured and deteriorated blades.

In order to achieve greater relevance of the final conclusions the first approach should be used to obtain the blade scans for further analysis. The differences in geometry and performance due to deterioration should be evaluated for individual blades as well as for the entire blade set statistics.

4.4. Operating point

The engines will realistically be operated in a variety of atmospheric conditions using various power settings. The engines are therefore designed to produce a certain level of performance in all expected situations. This performance can be summarised in terms of performance maps. However, computing these performance maps of engines is a large computational task in itself as it requires CFD simulations for several operating points. In order to reduce the complexity of the task a suitable operational point for the analysis needs to be selected.

Following the argument from section 4.2.1 the engines spend most of their operational time in cruise where most of the fuel is being consumed as well. This is therefore the most important operational point for the engine. Therefore the optimal design cruise point should be used for the evaluation. Specifically, this is the maximum efficiency point on the 100% speed line.

4.5. Research project overview

In order to answer the research question given in section 1.3, several intermediate objectives need to be completed. The aim of this section is to provide an overview of these objectives, the research sub-questions that they answer, and the selection of the methodologies used to complete these objectives.

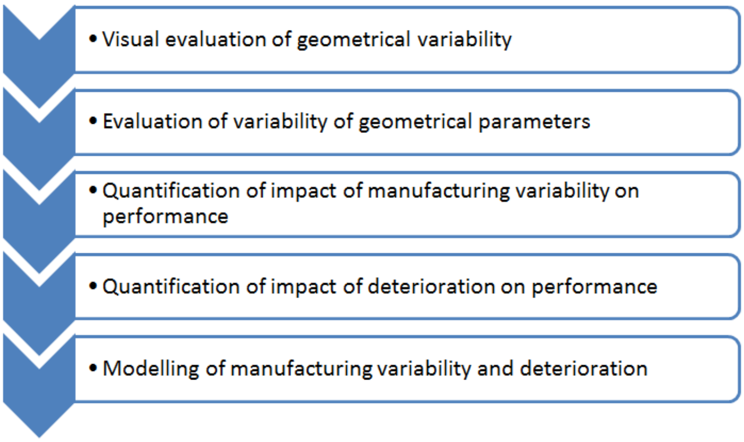

Figure 4.1: Intermediate objectives to be performed in order to answer the research question.

The last step in the flowchart is essentially a prerequisite for any subsequent robust optimisation activities. Robust optimisation relies on the ability to model manufacturing variability and deterioration. Therefore the modelling of the effects of manufacturing and deterioration itself is interesting.

4.5.1. Research sub-questions

The sub-questions that need to be answered in order to answer the research question are presented below:

1. What is the geometric impact of deterioration?
    (a) How does the geometry change from early to late life of the blades?
    (b) How do we separate the effects of different types of deterioration?
    (c) How do we model the impact of these effects on geometry?
2. What is the performance impact of deterioration?
    (a) How does the performance change from early to late life of the blades?
    (b) How do the combined deterioration effects compare with manufacturing variability in terms of performance impact?
    (c) How do the separate deterioration effects compare in terms of performance?

4.5.2. Methodology selection

In order to answer sub-question 1, a blade characterisation approach needs to be selected. Based on the findings of section 3, the NACA like parameters were selected to characterise the blades. The alternate approach identified is the PCA approach. The NACA like approach was chosen as it was used more often in literature, gave results comparable to the PCA approach, and is intuitively easy to understand and interpret. The parametrisation of the blades into NACA like parameters has been implemented within a software package developed by TU Dresden named B2P [62], which was selected for the current work.

The performance impact is calculated based on results of 3D CFD evaluation of blades. The Rolls-Royce proprietary CFD software suite Hydra [43] was used for this. To create the mesh for Hydra the in-house meshing software PADRAM [57] was used. The visualisation of the flow results was performed using an open source software application ParaView.

An additional aim of this thesis was also to produce a generalised uncertainty quantification workflow that could support future investigations into importance of manufacturing variability and in service deterioration of other available measured blade data sets. This workflow was implemented in iSight and used to obtain the uncertainty quantification results presented a subsequent chapter.

5

Geometrical analysis of manufacturing and deterioration variability

The aim of this chapter is to give an overview of the geometrical changes of the blades due to the deterioration processes. The changes are obtained by comparing the measured manufactured and deteriorated geometries of the same blades. The comparison is done visually, by calculating the point-wise surface deviations in the GOM Inspect software package, and by comparing the blade parameters calculating the B2P software. The comparisons aim to show the scatter of the deviations and parameters, and any offset from the intended geometry.

First, sections 5.1 and 5.2 present the process of calculating the GOM surface deviation plots, and the B2P calculation of the geometrical parameters used in the analysis respectively. The geometrical parameters are the NACA parameters and the blade heights.

Before any comparison of the data can be made, the blades need to be aligned to the appropriate common coordinate system. Obviously deviations would also be observed if the nominal CAD blade were compared to the same blade that was offset, rotated or a combination of both. Section 5.3 gives an overview of the alignment methodology, the importance of selecting the appropriate alignment and the actual alignment used to orient the measured blade scans.

The geometry deviations that remain after the scans have been properly aligned include the effects of dovetail and aerofoil scatter combined. Section 5.4 gives the motivation for separation of these deviations as well as an overview of the methods to achieve this.

Section 5.5 presents the visual analysis of the scanned blades. The visual analysis is performed on the basis of surface deviation plots made in GOM Inspect. The effect of deterioration is extracted by comparing the deteriorated to the manufactured blades. In order to show how these compare to the original manufacturing deviations, the results of comparison of manufactured blades to the nominal CAD design are also given.

Lastly, section 5.6 presents the analysis of the NACA like aerofoil parameters obtained using B2P. The analysis focuses on spanwise profiles on the parameters as well as the blade-to-blade scatter of the parameters.

5.1. Calculation of geometry deviations

The manufactured blade geometry can be compared to the design intent by comparing the scanned point cloud representation of individual blades with the CAD drawing of the design. This is done in the “GOM Inspect” software package. In GOM Inspect the scanned blade is represented as a mesh with patches filling the areas between the scanned points. Furthermore two geometries can be compared one to another in the software by calculating the difference between them and displaying this difference in a colour range. This allows the user to visualise the regions where and how the manufactured blades deviate from the design intent.

5.2. Calculation of geometry parameters

This section gives an overview how the geometrical parameters that are used to characterise the blades are calculated. Subsection 5.2.1 gives a brief overview of the calculation of NACA like aerofoil parameters based on Lange et. al. [35] and the B2P user manual [62].

The NACA like aerofoil parameters give the shape of the blade. Another important feature for compressors are also the tip gaps, for which the blade height is required. Subsection 5.2.2 gives an overview of the blade height calculation based on the B2P user manual [62].

The tip gaps depend on elements other than the blade, such as the drum and the casing. Subsection 5.2.3 presents how to combine the calculated blade heights with the information regarding other elements to calculate the tip gaps.

5.2.1. Calculation of NACA like aerofoil parameters

The reduction of the data for analysis can be done using NACA style geometric parameters as argued for and described in [35]. The use of these parameters provides clarity and control to both the analyst that is interpreting the data and the designer who subsequently aims to produce an improved design. The calculation of the parameters has been performed using a tool called B2P developed by TU Dresden on the basis of [35] [62]. The parameters that are calculated with this tool are illustrated in figure 5.2 for a specific blade section.

The parametrisation approach aims to calculate the parameters at several user defined span-wise positions along the blade. First, the section points at each of the positions are identified and converted into streamline coordinates. Then a camber line is fitted to the section outline. Based on the camber line the chord line is fitted as well. These lines are then converted from sets of 2D points into distributions and a thickness distribution is calculated as well. The NACA like aerofoil parameters are calculated based on these distributions.

The section positions are defined by the user as the percentage location between the hub and the tip of the blade. The parameters calculated for each of the stations together give a 3D characterisation of the blade. Note that the parameters are only calculated for the selected sections, and therefore any information of the blade geometry between them is lost in the parametrised characterisation of the blade. Furthermore, any effects with a frequency higher than the spanwise section frequency are lost. Lastly, surface roughness cannot be calculated from the measurements, as their resolutions is too large.

An example is shown in figure 5.1. It shows the determination of these coefficients for an example blade. In this figure three major features are plotted. These are the tip and hub contour, plotted in blue and brown respectively, and the parametrisation section outlines that lie between them in cyan.

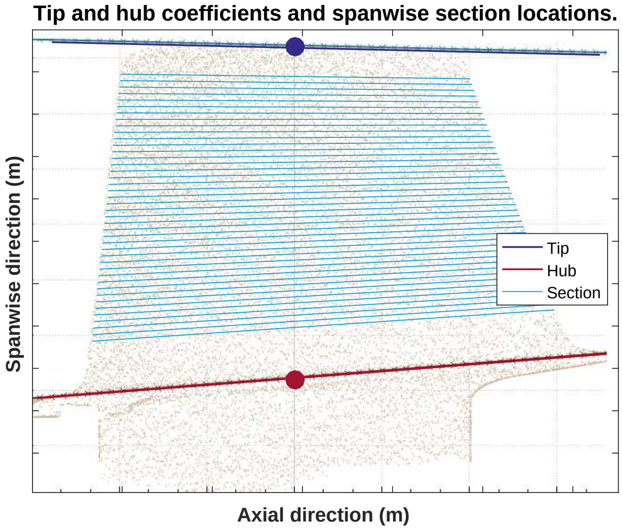

Figure 5.1: Example illustrating the blade parametrisation.

To allow the calculation of absolute span-wise positions of the individual sections and their slopes, the information regarding blade height and its position within its coordinate system must be provided by the user. This is done by inputting the so called blade tip and hub coefficients, which contain this information.

These coefficients are defined as the spanwise position of two points that together give the blade height and location, and the slopes of the hub and the tip contour at these points. Only the spanwise coordinates of the points are given as coefficients as the axial location is hardcoded within B2P. In total there are therefore four coefficients: two spanwise positions and two contour slopes. They need to be determined manually.

After data points of individual sections have been identified, the camber line and camber distribution can be calculated. The camber line is determined using the Rolling Ball Method (RBM). The RBM works by "rolling" balls of various radii inside the outline of the aerofoil section. The centres of the ball will meet twice for each ball radius, once near the LE and once near the TE. The intersections are points on the camber line.

To improve the quality of the calculated camber line a subsequent step is then performed. This step aims to find a camber point for each of the points that make up the section outline. It calculates the camber point as the mid point between a point on the suction side and its corresponding point on the pressure side of the blade. However, as the data is unstructured these points do not exist and must therefore be interpolated between the points closest to the correct location.

To identify these points the camber points calculated with the RBM are interpolated by a polynomial. A line is then drawn that runs from the selected point orthogonally through the camber line and pierces the section outline at the other side. The points closest to this line where it pierces through the section outline are used to calculate the location of the point corresponding to the chosen point. With both the corresponding points available the camber point is now calculated.

After the camber line points have been calculated, the corresponding chord points can be calculated as the projection of the camber points onto a straight line connecting the first and last points of the camber line. The chord can then be expressed as a chord distribution by calculating the distance of each of the chord points from the LE point and normalising with the total length of the chord. This converts a vector of 2D points into a vector of values between 0 and 1.

As the camber points are calculated for each of the section outline points the blade thickness values can also be calculated for every point of the outline using 5.1 [35].

$$t_k = 2\left|\vec{X}_k^{profile} - \vec{X}_k^{camber}\right| \tag{5.1}$$

Furthermore, the camber can now be translated from a vector of 2D points to a 1D vector of values by calculating the distance between each of the points and its corresponding chord points using equation 5.2 [35].

$$w_k = \left|\vec{X}_k^{camber} - \vec{X}_k^{chord}\right| \tag{5.2}$$

This procedure works well in regions where the section outline curvature is low. If the curvature of the outline is high, the procedure does not produce good results due to an assumption made when calculating the camber line, specifically in the step where the selected points corresponding point was calculated as a linear interpolation between the nearest points. This assumes that the interpolation is very close to where the real outline would be. In highly curved regions this is not true. For these locations the camber line points are calculated based on extrapolation from the known points. This extrapolation is done for a fraction of the entire chord length however and varies from application to application. It can be approximated as the ratio of the summed leading and trailing edge radii and the chord length.

Finally, the blade geometry is described using a chord, thickness and camber distributions. The chord distribution gives the non-dimensionalised locations along the chord line and thickness and camber distributions together give the shape of the fitted aerofoil section. Note that these distributions only give the general shape of the aerofoil section, but not its position, orientation or size. For this, additional parameters are required.

The parameters that give the position are the axial and tangential positions. They are obtained as the coordinates of the first point in the chord line. The orientation of the blade is given by the stagger angle and is calculated as the angle between the tangential axis, as shown in figure 5.2 [62] and the chord line, after the

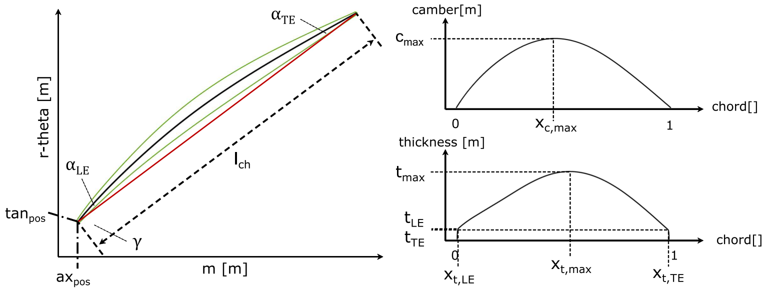

(a) Profile section parameters (b) Thickness and camber related parameters

Figure 5.2: Parametrisation of a blade section into thickness, camber and chord distributions and NACA like parameters.

symbol	profile parameters
ax, tan	axial and tangential position of section outline
λ	stagger angle
c	chord length
t_{LE}, t_{TE}	thickness of leading and trailing edge
$pos_{t_{LE}}$, $pos_{t_{TE}}$	assigned position on chord
t_{max}, w_{max}	maximum thickness and camber of profile
$pos_{t_{max}}$, $pos_{w_{max}}$	assigned position on chord
α_{LE}, α_{TE}	angle at leading and trailing edge

Table 5.1: NACA like parameters calculated by B2P for a particular spanwise section.

offsets from the origin have been removed. The size of the aerofoil section is given by the chord length, which is calculated as the distance between the first and the last points of the chord line.

The rest of the parameters are calculated directly from the camber and thickness distributions. The angles at the LE and TE are calculated as the angle of the camber distribution at the LE and TE respectively. This is done by fitting a 2^{nd} order polynomial through the camber distribution within the first and last 1.5% of the chord length respectively. The slopes of the fitted polynomial at the LE and TE are then the LE and TE angles.

The thicknessess of the LE and TE are calculated at a position where the slope of the thickness distribution reaches certain hard-coded values. The thickness is calculated as the distance between the suction and pressure sides of the blade.

The positions where the LE and TE thickness are calculated are also stored as individual parameters. Essentially, these positions are a measure of how elongated the LE and TE are. The nominal blades will have a constant LE thickness along the span, as the minimal thickness is desired. As the chord decreases along the span, the relative position of the LE therefore increases. As such, the LE becomes a longer part of the chord. At a single section, the LE position can also increase if the blades become elongated. Therefore by comparing the LE positions of several blades, and keeping their chord lengths in mind, the elongation of the LE shape could be assessed.

Additionally the maximum thickness and camber and their respective positions are calculated from the appropriate distributions. These help compare the thickness and camber distributions without the need to actually plot them.

The entire calculation therefore results in 14 individual parameters and 3 distributions for each of the sections evaluated. These parameters are summarised in table 5.1 and illustrated in figure 5.2 [62].

5.2.2. Blade height calculation

The blade heights are calculated to determine the blade height variability. Because the heights are variable and unknown the section corresponding to the blade tip cannot be calculated. Furthermore, the tip has a 3D shape which could not be captured with a 2D slice. A separate method is therefore required to determine the tip contour.

The deteriorated blade heights were calculated from the scanned data by using the B2P tip contour extraction functionality. The application and methodology of this method on an example case can be seen in figure 5.3.

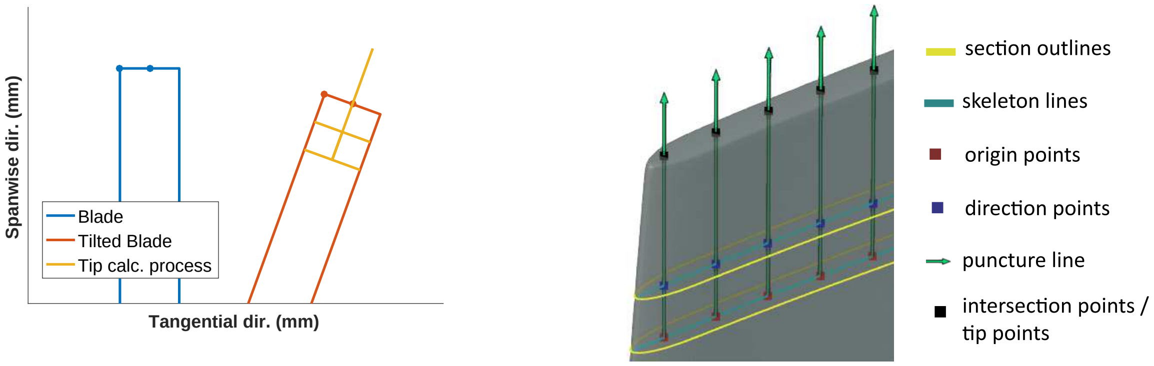

Figure 5.3: Calculation of tip contour. Left subfigure illustrates supports the reasoning behind the method implemented in B2P. Right subfigure shows the application of this method on a blade [62].

Figure 5.3 contains two subpanels. The left panel shows two blade geometries simplified as rectangles in 2D. The blue blade represents the ideal geometry, whereas the brown blade represents an example non-ideal realisation. Both blades have two points plotted on their blade tips. The points desired for the construction of the tip contour are the ones that lie at the midpoint of the tip. Lastly, the brown blade also has a construction used to arrive at this point drawn in. This panel supports the discussion of the reasoning of the method implemented within B2P. The right panel shows the visualisation of this method in 3D from the B2P user manual [62]. This panel supports a more detailed description of the method to calculate the tip contours.

The blade tip shown in figure 5.3 is flat, which is an approximation. The tips of the real blades are rounded by the manufacturing and deterioration processes. However, because the radius of the tip is very large and the thickness of the blade is very small the difference between a flat approximation and the expected height is lower than the resolution of the scan. The approximation can therefore be made.

An approach to calculate the tip contours could be to simply select the highest points at each axial position. The major issue with this approach is that it is very sensitive to the orientation of the blade. If the blade is perfectly upright, as the blade drawn in blue in figure 5.3, then the approach is sound.

However, the blade geometry is variable and therefore it is very likely that the all the blades will be tilted with respect to the tangential direction by some degree, as shown with the brown blade in figure 5.3. The highest point of the brown blade is located on the left edge of its tip. In this case therefore the approach to select the highest point does not give the desired result. This problem remains even if the approach selects the highest point in radial coordinates. In this case, the edges of the tip are always farthest away from the axis of revolution and will therefore always be selected by this method.

To avoid these problems the tip contour extraction implemented in B2P uses a different approach, which is drawn in yellow on the brown blade. The two lines that are parallel to the tip represent two sections of the blade. A line perpendicular to them can be drawn through two corresponding points of these sections. In this case these points are the midpoints of the yellow lines. The location where this perpendicular line punctures the blade tip is a point of the tip contour.

The selection of the sections and points used for the calculation of these positions in 3D can be seen in the right panel of figure 5.3. The sections selected to support the calculation are drawn in yellow and their skeleton lines in blue. In this case, the points of the skeleton lines at the same axial locations are used to construct the lines that point the tip.

The skeleton lines lie in the middle of the aerofoil sections in the tangential direction, and it is assumed that the lines going through them also go through where the skeleton point of the tip would be. For this assumption to be valid the blade geometry should not change significantly from the location of the yellow sections to the tip. For usual blade designs this can be ensured by selecting these sections close to the tip.

The lines that determine the points of the tip contour are drawn as green arrows on figure 5.3. The actual points of the tip contour are plotted using black markers. Note that this method can only determine the tip contour points for the axial positions that are within the domain of both skeleton lines. In other words, the two skeleton lines span different ranges of axial positions. The tip contour points can only be calculated for axial positions that lie in the intersect of these ranges. If the blade chord length would increase drastically above the selected sections the tip contour would be determined only partially. Similarly, if the blade is tilted heavily in the axial direction, the procedure will return only a part of the tip heights. Usual compressor designs have a reducing chord length with span and are not tilted in the axial direction.

5.2.3. Tip gap calculation

The tip gap can be calculated as the distance between the casing and the blade tip height. The determination of the actual values is very complicated as it depends on multiple factors, which range from geometrical dimensions to the mechanical and thermal stresses when operating. A typical arrangement of a compressor blade assembled within the compressor module can be seen in figure 5.4. This figure shows 4 principle components which together make up the tip gap. These are the compressor drum, blade, casing liner and casing. Note that the tip gap does not depend on the dimensions of any single component. Instead it depends on multiple components at once.

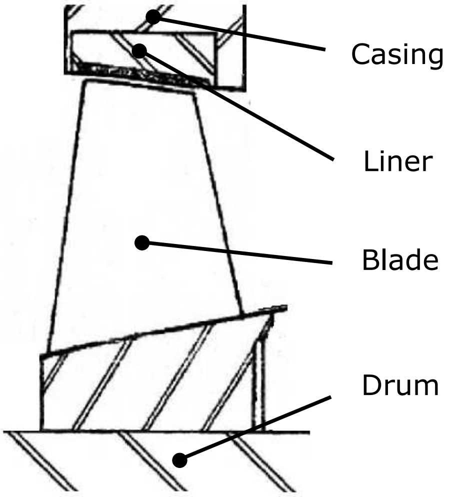

Figure 5.4: General arrangement of the blade within a compressor assembly, and the resulting tip gap..

The drum and the blade together make the lower bound of the gap and the casing and liner make the upper bound. The determination of the tip gaps therefore depends on the dimensions and tolerances of all of these components. These dimensions are not available for the drum, casing liner and the casing as these elements have not been measured along with the blades. However, within the scope of this work the interest lies solely with the blade performance. This can be isolated from the rest of the assembly by making appropriate assumptions regarding all other components, which are summarised in table 5.2.

The manufactured blades were scanned before their final machining process, as mentioned in section 4.1. Therefore for the manufactured blades no data is available for any of the compressor components and the tip gaps therefore cannot be calculated.

Part	Manufactured tip gap	Deteriorated tip gap
Drum	Assumed nominal	Assumed nominal
Blade	N/A	From scans
Liner	Assumed nominal	N/A
Casing	Assumed nominal	Assumed nominal

Table 5.2: Source of dimensions required for the calculation of the tip gap.

For deteriorated blades the blade heights can be calculated. However, the stage 3 rotor blades that are being evaluated have an abradable liner. As explained in section 2.3.2 the blades are therefore not expected to have their heights reduced and the majority of the tip gap increase is expected to come from liner abrasion. The assumption of the liner dimensions made in the case of the manufactured state therefore cannot be made for the deteriorated case. The tip gaps cannot be calculated for either the manufactured state or the deteriorated state.

5.3. Aligning the scanned blades

Subsection 5.3.1 presents a method to align the blade measurement scans to the nominal design. The alignment is required because the blades are secured into their position manually in the blade scanning process. This introduces differences in the orientation of the blade coordinate systems and the location of the blades within their respective coordinate systems. Furthermore the nominal design will have its own coordinate system and location within it as well. In order to compare the scanned blades to one another and to the design intent, the scanned blades need to be aligned to a common reference – the design intent blade.

Subsection 5.3.2 shows the impact different alignment techniques might have on the results of the subsequent analyses. Different alignment techniques include different selections of alignment surfaces as well as different methods of calculation of the alignment.

Lastly, subsection 5.3.3 presents the alignment selected to align the blade measurement scans for the subsequent geometrical and performance analyses.

5.3.1. Alignment procedure

GOM Inspect offers the user several methods for aligning the blades and allows for selection of specific areas of the scanned blade for comparison with the nominal blade. As the exact parameters for the alignment are unknown, the “Local Best-Fit” method was selected.

The "Local Best-Fit" method relies on a user selected subset of the geometry to align it to a reference object. In other words, a user selects a part of the measured blade mesh that will be used for alignment. The selected mesh geometry is then aligned to the corresponding surfaces on the CAD model. The operation that is required to align these point is then performed on the rest of the geometry, thus aligning the entire measured blade mesh.

For this, the GOM software must have the information on which CAD surfaces to use to align the selected mesh geometry. This is not done explicitly, by specifying the exact surface to be used, but rather implicitly, as the surfaces nearest to the selected mesh geometry. Therefore the blades must be pre-aligned such that these surfaces are reasonably close to each other, before the final alignment is made.

After the geometries have been pre-aligned and the alignment surfaces selected, GOM inspect will find the optimal position that minimises the distances between the selected alignment surface. This is done by minimizing the sum of the squared distances between each of the points on the selected mesh geometry and the nominal blade [17]. Specifically, the distances are the lengths of the normal vectors from the CAD alignment surface to the individual points of the selected geometry mesh point.

The alignment therefore depends on the areas selected for the alignment. If very small areas are considered the end results are sensitive to local geometric uncertainties of these areas. Furthermore selecting different areas for alignment will produce different results. For example, the alignment could be chosen such that the fit between the blade geometries is optimised while the root is not considered. Such a fit would show only minimal variation of the blade stagger angle, even if in reality the variation is large.

Lastly, this also means that any points selected far from the rest of the selected mesh geometries will have a relatively larger impact on the final alignment. Therefore care should be taken that the points are selected in roughly equal quantities from all the important alignment geometry features.

5.3.2. Importance of alignment technique

The blade parameters are extracted from the scanned geometry at various positions. As the blade shape changes along its span only the parameters at the same span locations can be compared. This means that the analysis of NACA like parameters is sensitive to the blades being orientated correctly in the parametrisation process.

To illustrate the importance of the alignment method on the results an analysis of the stage 3 blades has been performed using two different alignments. First of these alignments uses all the surfaces identified in green in Figure 5.5 and the second uses only some.

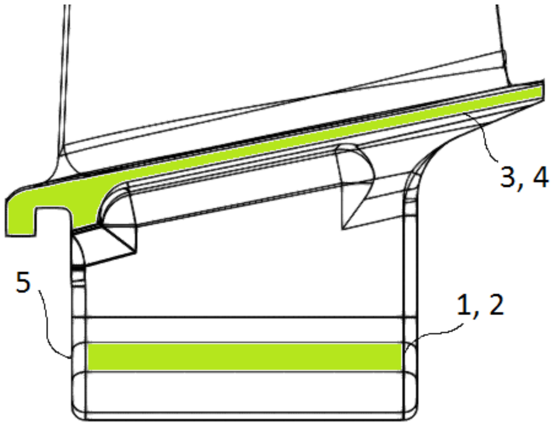

Figure 5.5: Surfaces used in alignment of scanned stage 3 compressor rotor blades.

Figure 5.5 shows a CAD drawing of a stage 3 compressor rotor from the side and from the front. Surface pairs 1, 2 and 3, 4 correspond to surfaces on both sides of the blade.

Alignment 1 uses surfaces 1, 2 and 5 for alignment. Alignment 2 additionally uses surfaces 3 and 4. The competing alignments result in different deviations from the nominal design blade. The difference in deviations for a sample blade is shown in Figure 5.6.

Figure 5.6 shows the normalised deviations from the design intent if alignment 1 and 2 are used, in the left and right sides of the figure respectively. The colour-bars are selected such that they show the same range of values. Note that the histograms of the deviations next to the colour-bars change somewhat with the change of alignment as well.

Several differences can be observed by comparing the left and right sides of figure 5.6. The most obvious is the reddish patch in the upper half of the blade that is present on the left side. This patch is missing from the right side. Both blades show high deviations from the nominal geometry in the region of the fillet. The region of high deviations is larger for the left blade. Furthermore, comparing the vertical faces just below the platform it can be seen that the deviations of this face are also larger for the right side. Lastly, at the tip near the TE the right side shows a region of negative deviations that is missing from the left side. Together these features indicate that there is a offset between the blades in the tangential direction. Specifically, the blade aligned using alignment 1 is offset to the suction side compared to the blade aligned using alignment 2.

However, comparing the deviations of the root slanted surfaces it can be seen that the deviations do not change. Furthermore, this is also true for the geometry of the platform. This indicates that the blade aligned using alignment 1 is not offset in the tangential direction, but rather slanted along it.

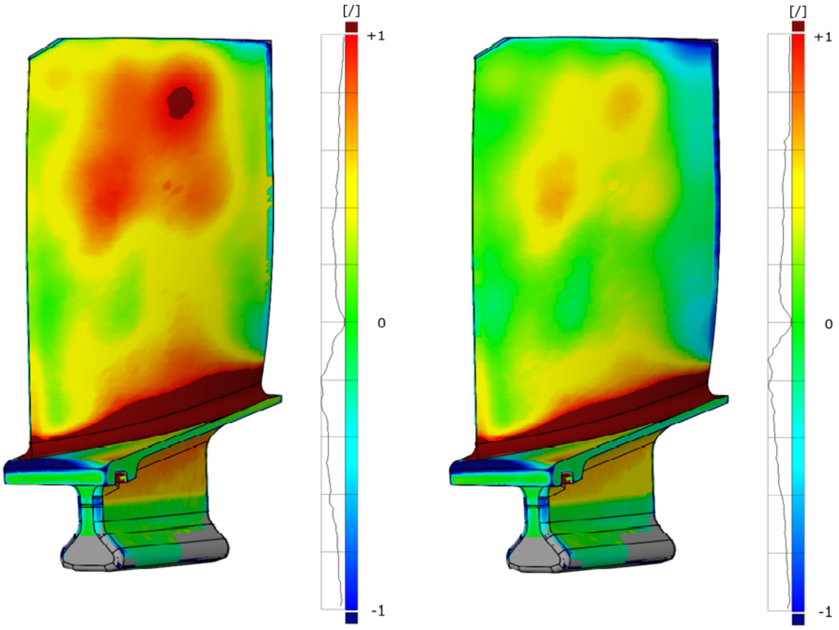

Figure 5.6: Impact of alignment technique selection on the geometrical deviations from the nominal design for a sample scanned stage 3 compressor rotor blades. The image aspect ratio has been distorted.

The NACA parameters obtained from the scanned blades also depend on the alignment. It is these parameters that are used for the generation of the blades used for the aerodynamic evaluation of the impact of geometric uncertainty. An analysis on the effect the alignment change has on the NACA parameters has also been performed.

Figure 5.7 shows the impact of the alignment on the calculated axial position of the blade. Note that axial position is determined relative to the B2P internal coordinate system centre, which is near the centre of the blade. The values along the axial direction are increasing from the LE to the TE, therefore all the axial positions will be negative. This means that negative differences mean that the blades have moved towards the front of the engine and vice versa.

The left panel shows the normalised differences due to alignment of individual blades. These normalised differences were calculated by subtracting the parameters of blades aligned using alignment 1 from the blades aligned using alignment 2, and normalising them with an appropriate arbitrary length scale. Mean spanwise normalised differences were also calculated and are plotted as a bold blue line.

The right panel shows the histogram of the spanwise averaged normalised differences. These are calculated by averaging all the normalised spanwise values of an individual blade. In this case this statistic is a meaningful measure of the scatter as all the blades exhibit a particular behaviour.

The left panel of figure 5.7 shows that there are offsets between the blades that are differently aligned. Furthermore, these offsets are not constant along the span, and their absolute value increases with the span. This suggests that the LE of blades aligned using alignment 1 is more vertical than the LE of blades aligned using alignment 2, which means that, relative to the latter, the former are rotated along the tangential axis towards the LE.

It is also interesting to note that there seems to be some correlation between the axial positions at the root and the slope of the spanwise differences, e.g. the line that has the highest value at the first spanwise position also has the highest increase of the differences with the span. This further supports the idea that the blades are misaligned by some rotation about the tangential axis.

This is surprising, since the difference between the alignments is the addition of two parallel surfaces that are expected to determine the position about the vertical and not the tangential axis. However, as these surfaces position the blade about the vertical axis they also change the position of the surfaces that position the blade

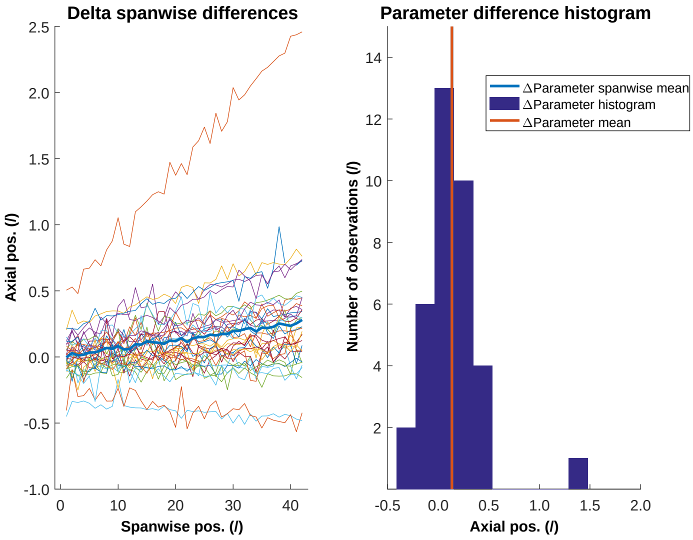

Figure 5.7: Comparison of axial position values calculated from scanned stage 3 manufactured compressor blades aligned using alignments 1 and 2.

about the tangential axis. This change of position then manifests in a rotation about the tangential axis.

The right panel of figure 5.7 shows that the mean normalised difference in axial position due to alignment is about 0.10 which is well below the scanner resolution of $30\mu m$. In most cases therefore this offset is negligible. However, there are cases where this is not so, and one of the blades was offset by approximately 1.50.

Similarly, figure 5.8 shows the impact of the alignment on the calculated tangential position of the blade. Note that tangential position is determined relative to the B2P internal coordinate system centre, which is near the centre of the blade. The values along the tangential direction are increasing from the Suction Side (SS) to the Pressure Side (PS), therefore all the tangential positions will be negative. This means that negative differences have moved the blade towards the PS and vice versa.

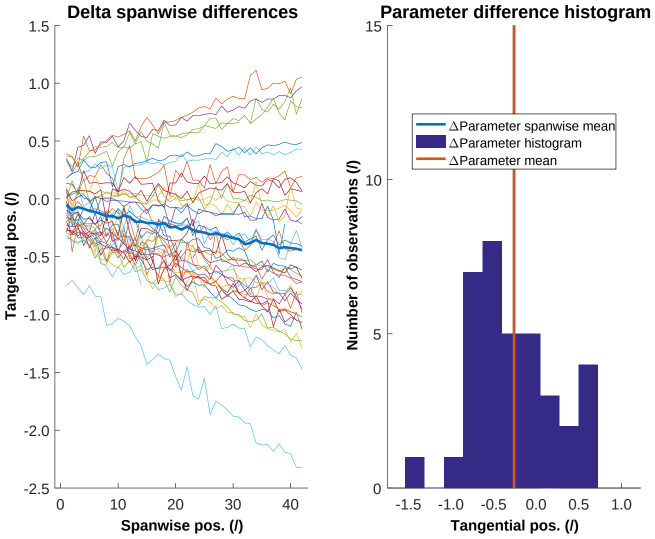

Figure 5.8: Comparison of tangential position values calculated from scanned stage 3 manufactured compressor blades aligned using alignments 1 and 2.

The left panel of figure 5.8 shows behaviour similar to that seen in figure 5.7. The differences exhibit a linear behaviour and their slope seems to be correlated to the value at the first section. However, the differences are larger in this case. The alignment had a relatively large impact on the value of the tangential position of the blade. The linearly increasing tangential positions with the span also mean that the blade is rotated about the axial axis towards the suction side, which is in agreement with the conclusions of the visual observations of a sample blade. Similarly as before, this is a consequence of rotating the blade along the vertical axis.

The extent of the rotation along the vertical axis can be seen by looking at the impact of the two alignments on stagger angle. This is done in figure 5.9.

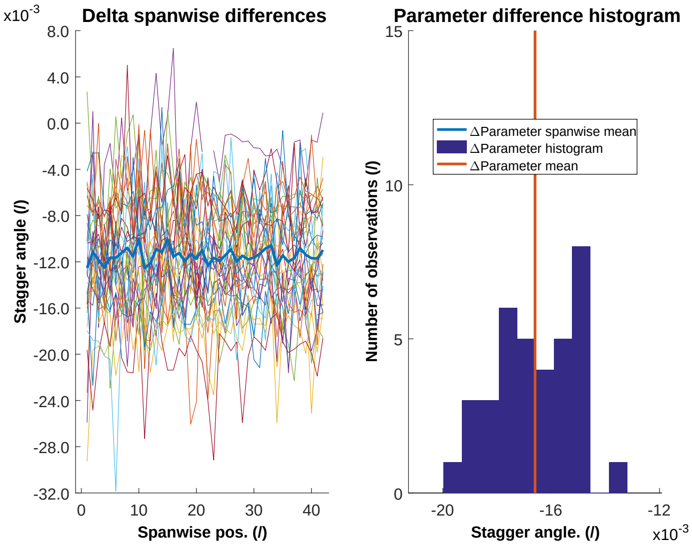

Figure 5.9: Comparison of stagger angle values calculated from scanned stage 3 manufactured compressor blades aligned using alignments 1 and 2.

In the case of stagger angle, the deltas due to the alignments are quite noisy. However, from left panel of figure 5.9 it can be seen that almost all the data lies below 0. This shows that all of the blades have opened to some degree, i.e. rotated such that the distance between the LE and TE in the tangential direction has decreased. As expected, no real trend with the span can be observed, however this is not so clear due to the magnitude of the noise.

The last parameters where significant changes were observed due to the alignments are the LE and TE angles. The results for the TE angle showed a larger dependency on the alignment and are shown in figure 5.10.

In the right panel of figure 5.10 it can be seen that the largest average normalised differences are somewhat higher than -0.08. The left panel of this figure then shows that the noise in the normalised differences due to alignments is significantly larger than this. Specifically, all the blades have values centred about 0 with the largest differences at approximately 1.4 and -1.4, which is more than 15 times the value of the maximum average differences. It is therefore likely that the average differences are simply an artefact of the noise in the LE and TE difference calculation. The same can be seen for the case of LE angles as well, therefore these results are omitted.

Furthermore, for parameters c, ω_{max}, $x_{\omega max}$, t_{max}, x_{tmax}, ℓ_{LE}, t_{LE}, x_{LE}, ℓ_{TE}, t_{TE}, x_{TE} the differences due to the change of alignment are much smaller than the scanner resolution, and are therefore omitted. Note that these parameters, along with the LE and TE angles depend solely on the shape of the aerofoil section and not on its orientation. Therefore it is expected that these parameters would not change due to the proposed alignments.

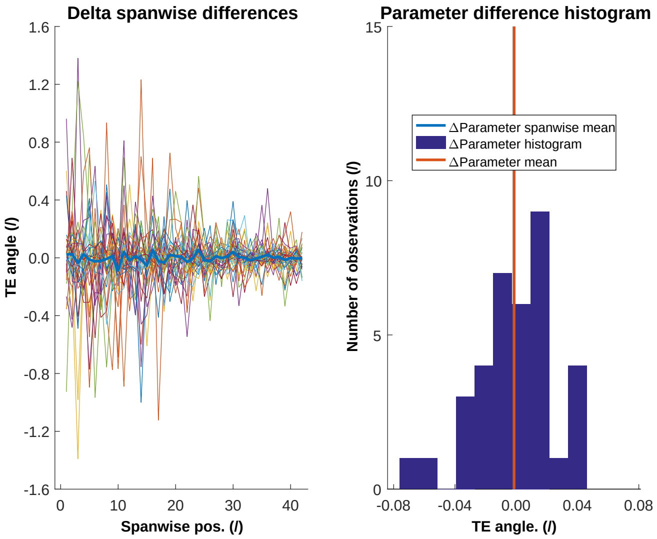

Figure 5.10: Comparison of TE angle values calculated from scanned stage 3 manufactured compressor blades aligned using alignments 1 and 2.

5.3.3. Actual alignment

Sections 5.3.1 and 5.3.2 illustrate the how the alignment is performed and its importance. This subsection presents the geometry selected to be used in the alignment of the scanned blades.

The alignment method for the manufactured blades is therefore based on how the blade is installed in the final assembly. Figure 5.11 shows a schematic of how blades with an axial root are installed in their respective position with the surfaces used in the alignment highlighted in green:

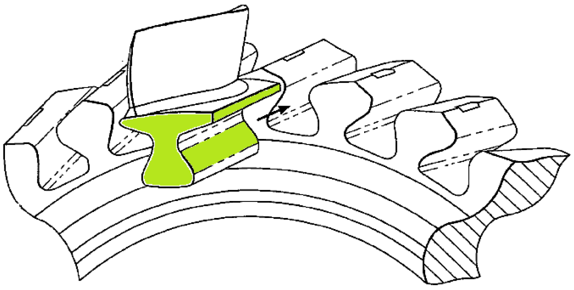

Figure 5.11: The assembly of a blade with an axial root type, analogous to the one of the stage 3 compressor blades that were scanned, into its final position.

Figure 5.5 showed the surfaces used in the analysis of the impact the alignment has on the parameters and subsequently on the performance evaluation. These are also the alignment surfaces used in aligning the manufactured blades.

The surfaces are on the blade root as well as the side surfaces which touch adjacent blades. The axial position is given by the front vertical surface. The front vertical surface has a step feature on it. The faces that actually touch the drum are the ones within the groove that can be seen in figure 5.12. The centrifugal forces are borne by the dovetail slanted surfaces. The vertical surfaces on the side of the blade platform touch the neighbouring blades and help determine the shape of the gas path passage. These surfaces are the ones used in the alignment of the blades for inspection after they are manufactured.

Note that the scans have missing geometry in the location of the groove that is in contact with the drum, and therefore another surface had to be chosen to determine the axial position. The flat surface below the step

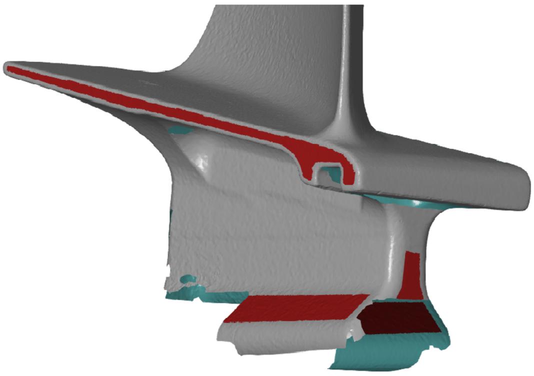

Figure 5.12: An example selection of the scanned data to be used in the alignment procedure of stage 3 manufactured compressor blades. The image aspect ratio has been distorted.

feature has therefore been used for alignment purposes. Note that this is not a high priority surface in the machining of the dovetail and therefore this selection will add some uncertainty to the geometry.

Figure 5.12 also shows the geometry that is missing due to the measurement process. At the back of the dovetail there is a large area where the geometry of the slanted surfaces is unavailable. This is where the blade had been clamped during the scanning process. Furthermore, the bottom side of the axial root is missing as well.

5.4. Separating the impacts of a and blade deviations

Blade geometry deviations due to manufacturing variability and deterioration occur due to the effects of these processes on its individual features. Knowing which features contribute most to deterioration geometry deviations enables the subsequent efforts to focus on improving these specific features as opposed to the entire design.

Subsection 5.4.1 gives an overview of the specific features that contribute to the geometry deviations and the advantages of focusing on individual features. Essentially, this subsection explains why this separation is desired, provides motivation for it and presents the levels at which it can be performed. Furthermore, the implications of using the separation at either level are then discussed, and the decision to use the separation at the alignment level is made.

Based on subsection 5.4.1, subsection 5.4.2 discusses the alignment techniques that can be used to determine which feature of the blade is responsible for most of the gas path deviations. Two techniques and their advantages are explained.

5.4.1. Possible levels for the separation of impacts

The blades are being analysed because they have some geometric uncertainty associated with them. Ultimately the interest does not lie in the geometric uncertainty of the blades, but in the uncertainty of the gas path passages between them. The shape of the passages depends mostly on the shape of the blade, the axial root, and the drum that holds the blades in place. Of these three, the axial root and drum geometries have a particularly large impact, as together they define how the blade will be positioned within the gas path. For example, the aerofoils can be manufactured perfectly, but if they are not oriented properly within the gas path due to poor axial root quality, the gas path will still be geometrically uncertain.

For the purpose of the current assessment it can be assumed that the drum is perfectly manufactured, which allows the nominal CAD blade to be used as the reference for alignment. By performing CFD on blades that are aligned to the nominal CAD design the performance scatter that exists due to their geometrical deviations can be calculated. Provided that the scatter is sufficiently large this analysis would warrant a robust optimisation of the blading element. However, the analysis would not tell which part of the blading element should be optimised for improved results.

This information is particularly worthwhile as the geometry and the manufacturing process of the axial root are relatively straight forward, whereas the ones of the aerofoil are not. Due to the large impact of the dovetail on the gas path shape and its relatively simple optimisation, it is therefore worthwhile to try and optimise the axial root first. However, if the actual geometrical deviations due to the axial root shape are small, then the aerofoil needs to be optimised. To summarise, improving the robustness of performance should be performed on the feature from which most of the deviations originate from.

The separation of the effects must occur somewhere in the performance evaluation process. Figure 5.13 shows the two identified methods of separation and the levels of their evaluation process. There are five levels all together. The levels are: the raw data, alignment, parameter, CFD evaluation and post-processing levels. The results in the postprocessing level are presented as pie charts, where the separate colors represent the degree of impact of the individual features on performance. Specifically, blue and brown are used to denote the aerofoil and axial root uncertainty impact respectively. Note that this figure also visualises how the performance impact can be traced back from the CFD evaluation to the original data.

As there are two features for which the impact on performance should be quantified individually, at least two separate performance evaluations are needed. Therefore the separation has to happen between the raw data level and the CFD evaluation level. This leaves two levels at which the separation can be performed, namely the alignment and the parameter levels. The two methods in figure 5.13 therefore show the processes of separation at these levels. These processes also show how the resu

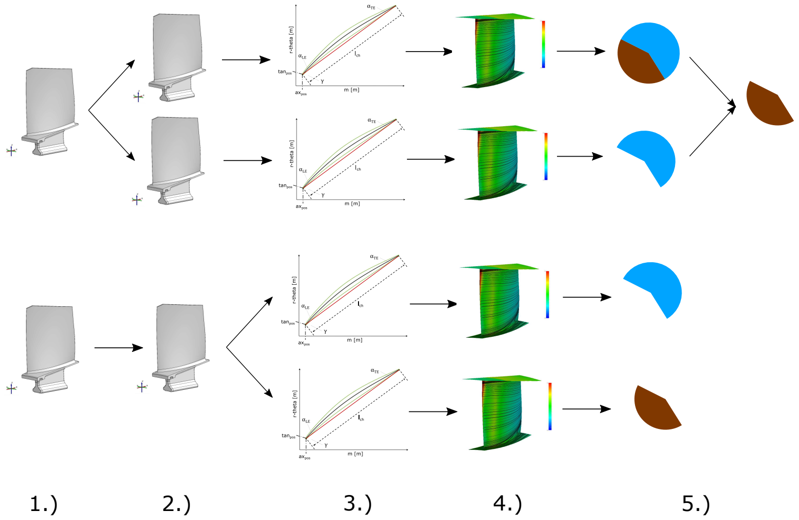

Figure 5.13: Visualisation of the separation of the impact of dovetail and blade deterioration at the alignment and parameters levels. 1.) Raw data level 2.) Alignment level 3.) Parameters level 4.) CFD level 5.) Postprocessing. The blade image aspect ratios have been distorted.

The upper process demonstrates the separation at the alignment level. From the raw data two separate blading element data are created by using separate alignments. The idea is that one of these geometries will include the effects of both dovetail and blade uncertainty, whereas the other will include only one of these. These geometries are then parametrised and evaluated with CFD. Lastly, a post-processing step is required to obtain the performance impact of the dovetail uncertainty.

Alternately, the separation of the impact of axial root and aerofoil deviations can be done at the parameter level. Here only one alignment is used, and the separation of the feature impacts is performed by separating the geometrical parameters into two groups that that define the position of the blade and its shape separately. The positioning parameters can be seen as part of the axial root geometrical deviations and therefore the impact of these parameters on performance would be the impact of the axial root deviations on performance.

These groups can then be complemented with the parameters of the nominal blade to perform the CFD analysis. Specifically, the group of position parameters can be completed with the nominal blade parameters, and the group of blade parameters can be completed with the nominal position parameters. The evaluation of these geometries therefore directly yields the impacts of the separated groups.

However, the position parameters would position the aerofoil along the axial and tangential axes, but would not account for any deviations along the spanwise axis. Furthermore, any rotation of the blade as a consequence of the axial root deviations would be included in the impact of the aerofoil, not the axial root. This is because the rotation of the blade is determined by the stagger angle, which cannot be decomposed into the two feature uncertainties using this approach. Stagger angle can therefore only be classified either as a position parameter or as an aerofoil parameter. The former case assumes that blade manufacturing results in perfect stagger angle, and the latter assumes that the axial root will not induce any stagger angle.

On the other hand, the process involving different alignment procedures quantifies the impact of the feature uncertainties without splitting the geometry. This avoids the issues highlighted for the parameter separation process.

5.4.2. Methods for separation

The question at hand here is how to separate the geometry deviations of the axial root and aerofoil for both the manufactured and deteriorated blades at the alignment level.

Note that the methods of using alignment to isolate the deviations of the root and the blade hinge on there being two different references available to align the scans to. Specifically, these need to be references from an earlier stage of the blade's life, as the intent is to capture the effect of the process that occurred in between. As discussed below, in the case of deteriorated blade scans these are their manufactured counterparts and the nominal CAD design. However, for the case of manufactured blades only the nominal CAD design is available as a reference. Therefore, in this case an assumption is required.

During the manufacturing process the axial roots are already machined to high precision using a manufacturing technique called broaching [49]. It could therefore be assumed that all the deviations that occur for manufactured blades are deviations of the aerofoil. Robust optimisation taking into account the manufacturing variability should therefore focus on the aerofoil.

Something similar clearly cannot be claimed for deteriorated blades, as in that case the dovetail will also deteriorate. A prominent example of dovetail deterioration is the abrasion of the slanted surfaces of the dovetail, which take the centrifugal load. The abrasion of these surfaces means that the entire blade will be offset from the nominal design in the radial direction under operating conditions which will have a direct impact on performance as well. This effect needs to be separated from blade effects such as blade abrasion.

To separate the effects of dovetail and aerofoil deterioration two approaches were identified:

1. The deterioration of the blade and the dovetail can be separated by comparing the scans aligned to the nominal blade with the scans aligned to the manufactured scans
2. When comparing the deteriorated blade to the manufactured blade, geometry that is considered to have not deteriorated is used (e.g. platform geometry).

Approach 1 can be used to separate the effects of deterioration of the aerofoil and dovetail. As mentioned in the above paragraph, due to the abrasion of the faces that take the centrifugal load the blades can be offset in the radial direction. That would mean that if we want to compare the manufactured parameters of some span-wise location H_{manuf}, the deterioration parameters would be taken from a section that is at $H_{deter} = H_{manuf} - H_{offset}$. Note here that since these surfaces are slanted at an angle α degrees the offset will be $H_{offset} = H_{abrasion} \cdot cos(\alpha)$.

While optimising the aerofoils to minimise the effect of the blades being offset in the radial direction is possible (more uniform blades with less twist) this should be done through optimising the axial root design instead (by using less slanted surfaces). This is not possible using the current methodology as no parameters of the axial root are captured. Nevertheless the optimisation of the aerofoil for deterioration would still be useful for engines currently in production. For these engines the design of the axial root should remain the same to allow for interchangeability with the current blades as this would avoid the need to re-design the drum. These results would then include an implicit assumed relation between blade deterioration and dovetail surface abrasion.

To sum up, the blades can be optimised using this approach, but there would be no control over the axial root, and no way to exclude it from the analysis. Some assumptions on the dovetail and its deterioration would therefore be baked into the result. Given the aerofoil shape resulting from the optimisation, the assumptions might no longer hold, as the new design could feature lower centrifugal forces, different directions of aerodynamic loading,...

Approach 2 would give more distinction between dovetail and aerofoil deterioration, as the alignment surfaces are assumed to have deteriorated insignificantly. The selected surfaces would be either on the gas path platform surface or the surfaces underneath the platform. These surfaces are not heavily loaded and are not expected to erode or accumulate dirt. The question here is how much dirt accumulation could affect the alignment. This question could be avoided by cleaning the blades prior to scanning them.

To sum up, approach 2 would not produce the issues in comparing the parameters and would at the same time allow for separate evaluation of the aerofoil and dovetail deterioration. In this approach the dovetail design and its deterioration would have to be represented by a model.

5.5. Visual comparison of scanned blades

This section provides a first glimpse of the geometrical deviations of the scanned blades. The deviations are presented visually, as surface deviation maps of entire blade surfaces, with the deviations of each of the geometrical points being expressed with colors. Secondly, blade aerofoil deviations for the hub, mid-span and tip sections are plotted.

Subsection 5.5.1 gives an overview of the calculation procedures that are required to obtain the desired visualisations. These are the calculation of the surface deviations as performed in GOM Inspect, and the calculation of the section geometries as implemented in B2P.

Subsection 5.5.2 gives a visual presentation of the manufacturing deviations. Lastly, subsection 5.5.3 gives a visual presentation of the deterioration deviations.

5.5.1. Methods for visualisation of measured blades

The calculation of surface deviations is performed in the GOM Inspect software. The surface deviations are calculated on the basis of a measured point cloud and a 3D CAD design. These geometries must have been aligned before the surface deviations have been calculated, and therefore the alignment will have an impact on the magnitude of the deviations.

Calculation of surface deviations

GOM Inspect offers two main ways to calculate the surface deviations. One is to calculate the deviations on the nominal CAD design, and the other is to calculate them on the measured mesh. This means that the methods also differ on which geometry the results are visualised. Specifically, the deviations are plotted on whichever mesh they are calculated on. The two approaches result in opposite results, as where the deviation from the nominal to the measured mesh is positive for one of the methods, it will be negative for the other.

In order to be able to compare the surface deviations of different scanned blades to one another, as well as to compare them for both manufacturing and deterioration deviations, they need to be calculated for comparable geometrical points. Therefore all of the surface deviation maps have been calculated based on the nominal CAD design as that allows a point-to-point comparison of different surface deviation maps.

The calculation process of the individual point deviations is quite simple. The deviation is the magnitude of the normal vector from the surface of the base geometry to the other geometry. In this case the base geometry is the nominal CAD design and the deviation at an individual point of the CAD design is the distance at which a normal vector of the nominal CAD design surface pierces the measured geometry mesh.

Lastly, the fact that the deviations of all the measured points are calculated at the same locations on the nominal CAD design allows to visualise statistics of the point-wise surface deviations. Specifically, the mean blade deviations can be calculated as the point-wise mean deviations among all of the deviation maps. Similarly, a standard deviation of the point-wise geometrical deviations can be made. This is done as a post-processing step outside of GOM Inspect.

Calculation of section geometries

An additional output of the B2P processing of measured blades are the coordinates of the individual spanwise sections for which the parameters and their corresponding distributions have been calculated. These are not the coordinates of the actual points measured by the 3D scanner, but the coordinates of the section outline fitted to them. These coordinates at given at predetermined percentages of the chord.

5.5.2. Visual inspection of the geometrical impact of manufacturing

To get a first impression of the impact of manufacturing on the blade geometry the surface deviation maps can be calculated. Figure 5.14 shows the mean geometrical deviations of the scanned blades. The left panel of this figure shows the deviations on the pressure side, and the right panel shows the deviations on the suction side. It can be seen in both of these panels that the deviations are predominantly positive. This means that the manufactured blades are generally thicker than the nominal CAD design.

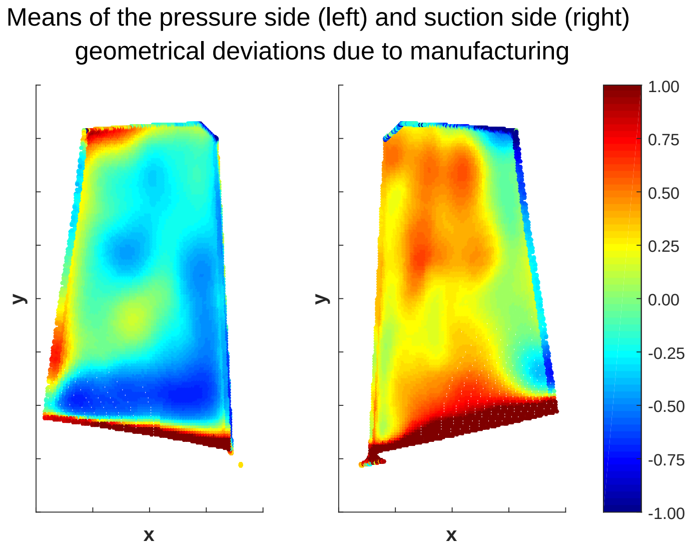

Figure 5.14: The mean geometrical deviations due to manufacturing variability of the pressure side (left) and the suction side (right).

Furthermore, regions with high deviations can be seen at the hub in both panels. Deviations in these regions are due to the fillet extending to high up the blade itself. On the pressure side it can be seen that the highest deviations are towards the TE. One such region is at the tip and another is at about quarter span from the hub towards the tip. Another clearly visible region of positive deviations of the pressure side is a circular feature located just behind the mid-chord position, slightly below the mid span section.

The mean deviations of the suction side are less pronounced. The majority of the deviations occurs behind the LE area and in front of the 75% of the chord. The deviations are not as high as those of the pressure side. This shows that the mean suction side is better manufactured than the mean pressure side.

The mean deviations only tell a part of the story however. For example, it could be that in reality most of the blades have high deviations, and some of the blades very low deviations on the suction side. The mean blade could therefore still have low surface deviations. To address this, figure 5.15 shows a plot of the standard deviations of the geometrical deviations at each of the points on the blade.

This figure shows the standard deviations color coded from blue to red, where blue means that the blades are all very similar in that region, whereas red means that the deviations are quite different from blade to blade. The left panel of figure 5.15 shows that the deviations are relatively small for most of the blade. The region with the lowest deviations is at the hub at the LE and the region with the highest deviations is at the LE of the tip. Note that here the deviations at the TE are not considered. This shows that the geometries at the hub are very similar to each other, while the geometries at the tip can vary much more.

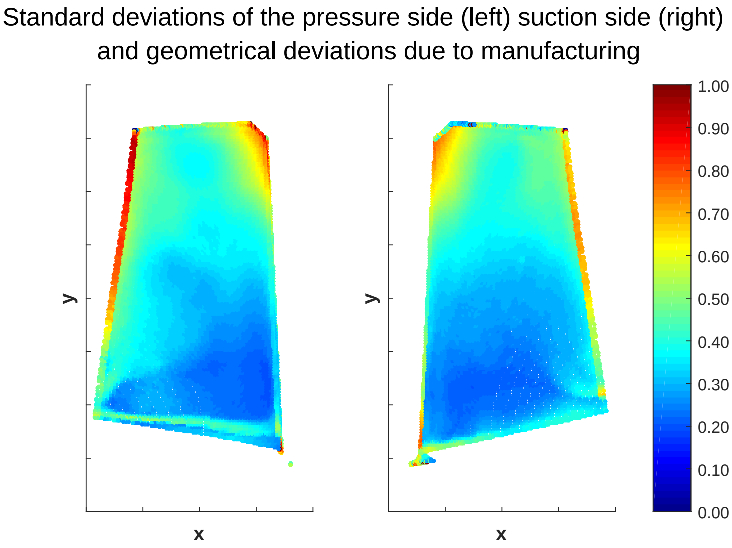

Figure 5.15: The standard deviations of the geometrical deviations due to manufacturing variability of the pressure side (left) and the suction side (right).

The same is true for the suction side, which is shown in the right panel of figure 5.15. In this case the deviations are higher than on the pressure side. This means that the blades are more different to one another at the suction side than on the pressure side.

This could be due to the blade suction surface changing independently from the pressure surface, or that both surfaces change together and the blades have different amounts of twist in that region. The first can be seen as standard deviation regions that appear on only either the pressure or the suction sides. The latter can be seen as regions that appear on both simultaneously.

At the tip at the LE it can be seen that the pattern of standard deviations visible on the suction side also appears on the pressure side. This indicates that this region is likely twisted by various amounts in different blades. Furthermore, it can be seen that in this region, the standard deviations on the suction side are larger. This means that in addition to the twist, the thickness in this region is variable, with most of the variability towards the suction side. Lastly, it can be seen that the entire suction side above the mid-span section has large standard deviations, relative to the region below the mid-span section. This further supports the idea that the blades are more variable towards the tip.

In order to visualise the effects of manufacturing better, the aerofoil sections of all the blades are plotted in the figures below. The comparison of the different aerofoil sections allows to compare their shape, position and orientation visually. The aerofoil sections are plotted for sections 1, 21 and 41, corresponding to the hub, mid-span and tip sections respectively. Note that the aerofoil sections have been rotated such that the LE and TE lie roughly on the same horizontal position, and that the aerofoil section shape itself is distorted by different scale axis for ease of comparison. Additionally to the individual aerofoil section plots their LE and TE have been plotted individually for ease of comparison as well.

Figure 5.16 shows the aerofoil section variability at the hub. The variability of the aerofoil sections at can be best seen at $x = 0$ as the contours of the manufactured aerofoil sections are distributed above the nominal CAD design contour. Towards the TE and the LE it can be seen that the geometries are quite different as well. This is better illustrated in figure 5.17, which shows these regions close up.

Figure 5.17 shows the LE contours on the left panel and the TE contours on the right panel. It can be seen that the manufactured blades generally have their LE contour behind the nominal design contour. Furthermore, the manufactured blades are thicker at the LE than they are intended to be. This can be seen by comparing the scatters at the lower and upper contours. At the lower contour the scatter is centred about the nominal contour, whereas on the upper contour the scatter is completely above it. This means that at least the blades that have their lower contour below the nominal contour are thicker than the nominal CAD design. Lastly,

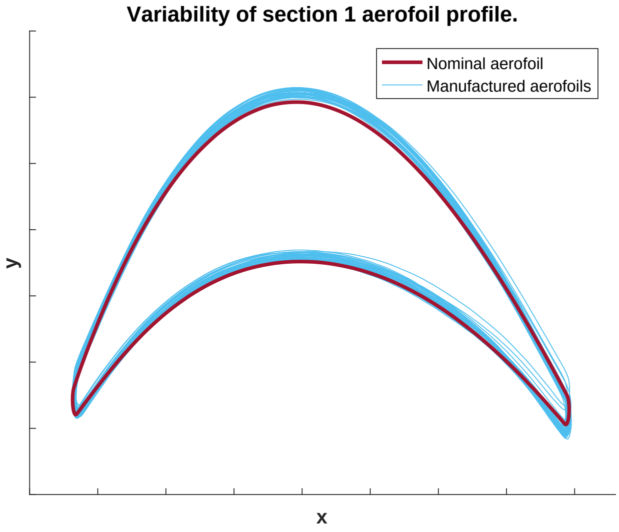

Figure 5.16: The variability of manufactured aerofoil section profiles for spanwise section 1. Note that the scaling of the axis distorts the actual aerofoil section shape.

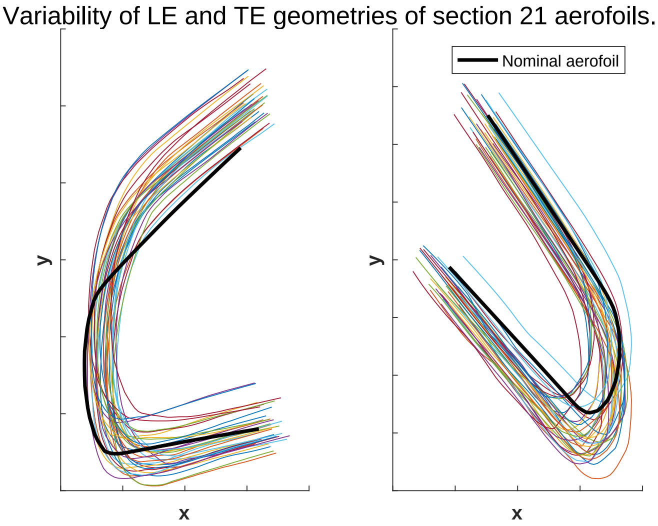

Figure 5.17: The variability of manufactured aerofoil section profile LE and TE geometries for spanwise section 1. Note that the scaling of the axis distorts the actual LE and TE shapes.

it can be seen that for the upper contour the manufactured blades follow a different slope than the nominal contour.

At the TE, the positions of the contours are more scattered than at the LE. The manufactured TE contours now lie both behind and in front of the nominal TE contour. Unlike for the case of the LE, the TE thicknesses appear to be more in line with the design intent. This can be seen as the scatter of manufactured contour lies beneath both the upper and lower nominal contours.

Figure 5.18 shows the scatter of the aerofoil sections for the mid-span region. Similarly to figure 5.16, figure 5.18 shows that $x = 0$ most of the manufactured aerofoil sections lie above the nominal contour, with a scatter comparable to the scatter seen for the case of the hub section. However, in the mid-span case the scatter increases towards the TE, where it is much larger than that of the case of the hub, seen in figure 5.16. Towards the LE the scatter has increased as well, but not as dramatically as towards the TE. Together this shows that the mid of the blades are manufactured well, with the scatter mostly increasing for the LE and TE.

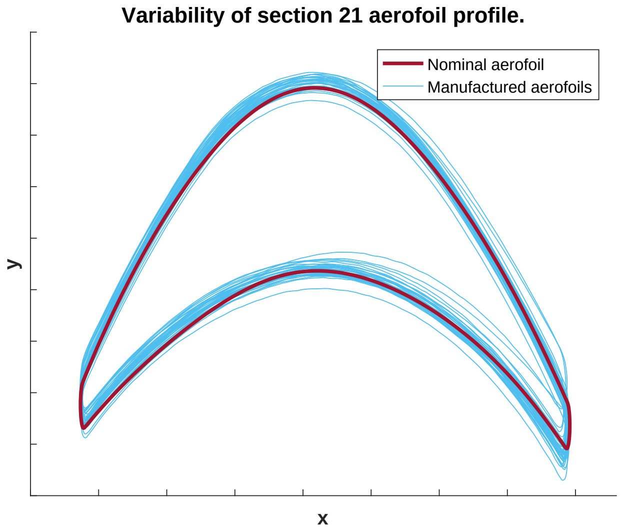

Figure 5.18: The variability of manufactured aerofoil section profiles for spanwise section 21. Note that the scaling of the axis distorts the actual aerofoil section shape.

The LE and TE regions can be seen better in figure 5.19. This figure further supports the statements made in the discussion of figure 5.18. Furthermore, it can be seen that the LE and TE contours appear to have a shape comparable to the nominal contour. In other words, the shapes of the manufactured contours are not the main driver of the scatter seen in figure 5.18. Instead it is the positioning of the contours that causes the scatter. Lastly, in the right panel of figure 5.19 the manufactured TE contours all fall in front of the nominal contour. The left panel shows that the manufactured contours fall slightly in front or behind of the nominal contour. This suggests that at the mid-span the chord lengths of manufactured blades are slightly shorter than intended.

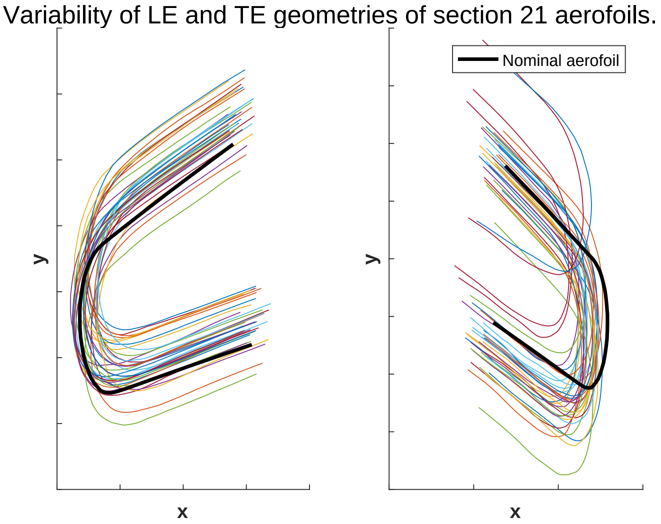

Figure 5.19: The variability of manufactured aerofoil section profile LE and TE geometries for spanwise section 21. Note that the scaling of the axis distorts the actual LE and TE shapes.

To see if this trend continues towards the tip, figure 5.20 shows the aerofoil sections of the tip section. As expected the scatter for this section is larger than both for the hub and the mid-span sections, as seen in figures

5.16 and 5.18 respectively. Furthermore, note that in the case of the tip contour scatter, the scatter at the TE is larger than the scatter at the TE. This further supports the discussion of the surface deviation maps as given in section 5.5.

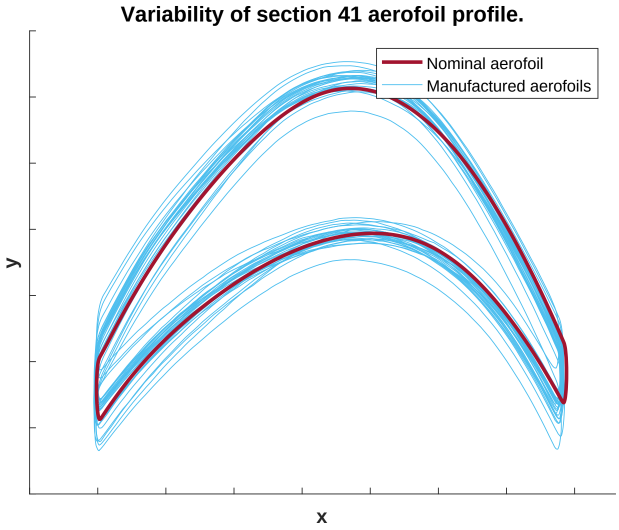

Figure 5.20: The variability of manufactured aerofoil section profiles for spanwise section 41. Note that the scaling of the axis distorts the actual aerofoil section shape.

The LE and TE contours of the tip section have been plotted in figure 5.21. This figure shows that, unlike the figures of the mid-span and hub contours, the LE contours now lie in front of the nominal contour. This means that along the span the LE of manufactured blades has leaned forward at some point above the mid-span position.

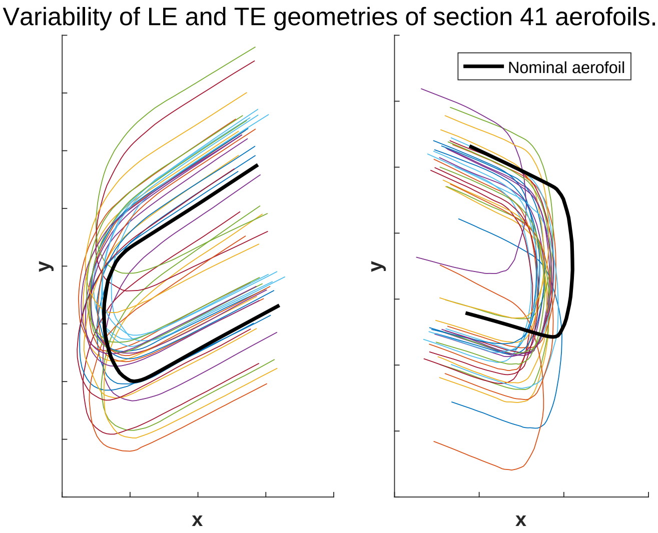

Figure 5.21: The variability of manufactured aerofoil section profile LE and TE geometries for spanwise section 41. Note that the scaling of the axis distorts the actual LE and TE shapes.

5.5.3. Visual inspection of the geometrical impact of deterioration

To get a first impression of the impact of deterioration on the blade geometry the surface deviation maps were also calculated. To extract just the impact of deterioration, the surface deviation maps of both the manufactured and deteriorated blades were created in GOM Inspect. The maps were all performed on the basis of the nominal CAD design, therefore the geometrical deviations of all the measured blades were calculated for the same set of points of the nominal design. The differences between the manufactured and deterioration geometrical deviations of individual blades then give the impact of deterioration.

Figure 5.22 shows the mean geometrical deviations of the impact of deterioration. The left panel of this figure shows the deviations on the pressure side, and the right panel shows the deviations on the suction side.

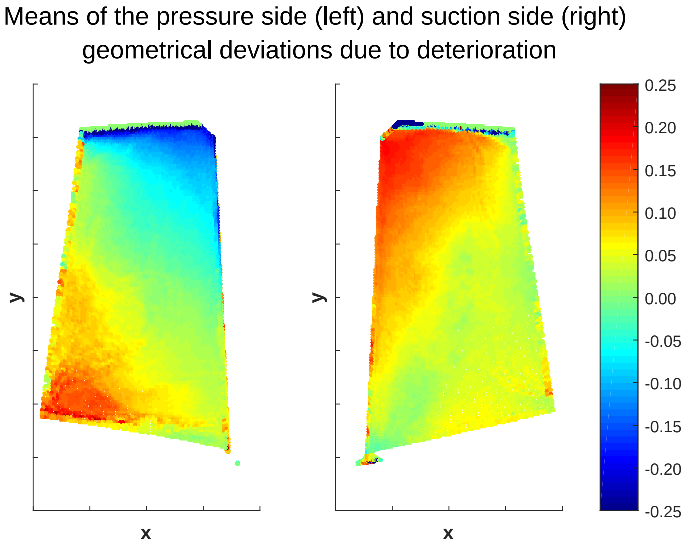

Figure 5.22: The mean geometrical deviations due to deterioration of the pressure side (left) and the suction side (right).

It can be seen in both of these panels that the deviations are predominantly positive. This means that the deteriorated blades are generally thicker than the nominal CAD design. However, on the pressure side the deviations are negative in a triangular region extending form the TE of the tip , LE of the tip and roughly half span along the LE. On the pressure side this region has positive deviations.

There are several possible explanations for this behaviour. It could be explained by aggregation of dirt in the location of positive deviations and blade erosion in the location of negative deviations. Furthermore, at the tip the blade might have been bent towards the suction side, or had its stagger angle decreased. Furthermore, the blade might have been leaned towards the suction side across the diagonal stretching from the tip TE and the hub LE.

To show how consistent this behaviour is, figure 5.23 shows the standard deviations of the geometry deviations due to deterioration. Note that the standard deviation value range of figure 5.23 is more than four times smaller than that of figure 5.15.

It can be seen that in general the standard deviations increase from the hub to the tip, which means that the deterioration at the tip is more variable than the deterioration at the hub. Furthermore, as the standard deviations do not appear in the same triangular fashion as the mean deviations observed in figure 5.22, it could be concluded that the pattern of the mean geometrical deviations at the tip LE occur to some extent for all blades.

Furthermore, patterns can be observed on both the pressure and suction sides. These patterns are a consequence of the blade geometries moving in the axial and radial direction. The patterns actually occur at the edges of deviation regions. Specifically, these are the edges across which the gradient of the geometrical deviations is large. In other words, these edges separate regions of very different deviations. These regions will largely appear on both the manufactured and deteriorated blades. As the blades move in the axial and

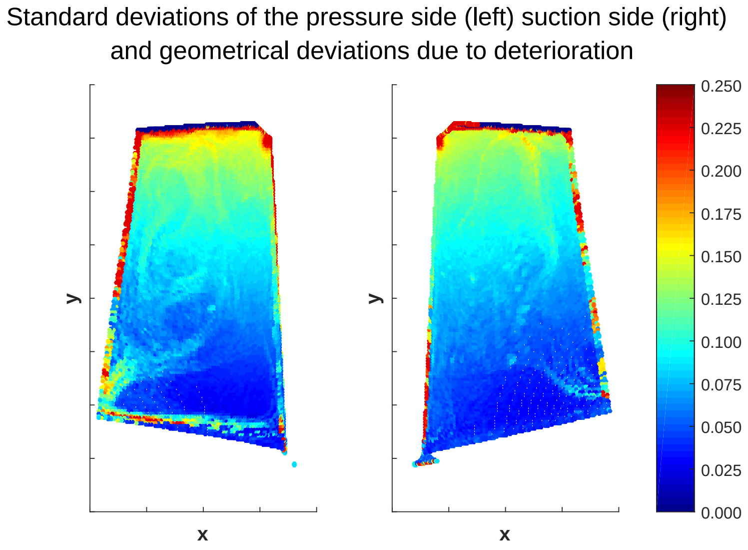

Figure 5.23: The standard deviations of the geometrical deviations due to deterioration of the pressure side (left) and the suction side (right).

radial directions, the edges of the corresponding regions will misalign, causing a line of higher deviations at the location of the misalign. All the lines due to this misalignments form a specific pattern on an individual blade. These are not visible on the mean deviation plots as they are averaged out. However, the most prominent do appear in the standard deviation plots as standard deviation is more susceptible to outliers than the mean.

As the geometrical deviations due to deterioration are relatively small, the plots of the entire deteriorated aerofoil sections gives no interesting information. Therefore they are omitted here. However, the LE and TE are expected to change. The LE of the deteriorated blades for section 1 can be seen in figure 5.24.

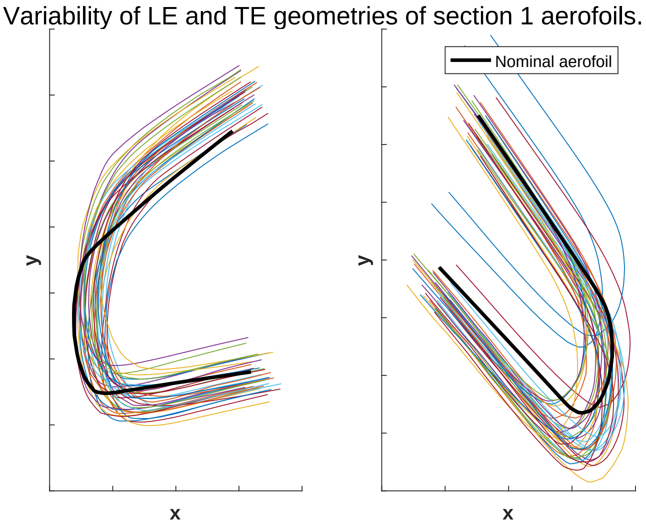

Figure 5.24: The variability of deteriorated aerofoil section profile LE and TE geometries for spanwise section 1. Note that the scaling of the axis distorts the actual LE and TE shapes.

Some subtle differences can be observed between the left panels of figures 5.17 and 5.24. In figure 5.17 notice the highes line of the bottom contour. In this figure this line represents two separate blades as their contours lie on top of each other. In figure 5.24 they are split, meaning that the blade LE have deteriorated differently. Furthermore, note that in figure 5.17 no scanned upper contour lies beneath the nominal upper contour. In figure 5.24 a line now lies beneath the nominal upper contour.

However, the major differences can be seen in the right panel. Most notably, figure 5.24 shows two scanned blades that have their bottom contours directly in the middle of the nominal blade, which is not the case in figure 5.17. Furthermore, the deteriorated blades are now thicker. This can be seen as the increased gap between the scanned bottom countours and the scanned nominal contour.

The comparisons of LE and TE contours of sections 21 (mid-span) and 41 (tip) show similar behaviour, therefore they have been omitted.

5.6. Parameter comparison of scanned blades

The aim of this section is to present the analysis of the deviations of the NACA like geometrical parameters obtained using B2P. The most interesting geometrical deviations are those due to deterioration. However, these must be seen in the perspective relative to the manufacturing deviations. Therefore these deviations are first presented in subsection 5.6.1. Afterwards, the deviations due to deterioration are presented in subsection 5.6.2.

In the case of blade heights a comparison between the manufacturing and deterioration deviations is not possible, as discussed in section 5.2.3. Instead, subsection 5.6.3 gives a comparison between deteriorated heights to the design tolerances to establish how large the entire deviations are.

5.6.1. Geometrical deviations due to manufacturing scatter

The aim of this subsection is to quantitatively describe the manufacturing scatter of the scanned blades in order to establish a reference when quantifying the deviations due to deterioration.

The positioning parameters such as the axial and tangential positions, and the stagger angle have the largest impact on the geometrical deviations of the gas path, along with the sizing parameter chord length. These are therefore analysed first.

Secondly the blade camber and thickness distributions are analysed. First the actual distributions are compared. Subsequently, the parameters that characterise these distributions are analysed as well.

Of these, the maximum thickness, and maximum camber, have the largest impact on the size of the blade, and are therefore inspected first. The LE and TE shapes typically have a large impact on compressor performance, and their parameters are analysed next. Lastly, the parameters with the lowest expected impact, such as LE and TE positions, and the positions of the maximum camber and thickness are analysed.

The comparisons of the parameters are done by first plotting the differences between the parameters of the manufactured blade and the nominal CAD design. This results in a single spanwise differences line per blade. Additionally, the mean, median, and the 6σ bounds of the differences at each spanwise section are plotted as well. This plot shows the scatter of the measured blade parameters about the nominal CAD design. These values in these plots have been normalised by an appropriate length scale. However, different plots that show parameters with similar units have been normalised by the same length scale in order to retain proportionality.

Secondly, the blades are expected to follow some systematic behaviour about the nominal CAD design. In order to determine the scatter of the individual lines about this systematic behaviour, the mean parameters at each spanwise section were subtracted from all the lines and plotted. This plot shows the scatter of the blades when compared to one another.

After the scatter about the mean spanwise differences has been determined, it is interesting to determine their origin. It could be due to some behaviour of the individual blades, due to offsets from the mean, or simply due to noise. In order to see which of these effects has a larger impact on the overall deviations a second plot, which shows the spanwise differences lines with their means subtracted, is shown.

Axial position

The manufacturing deviations of spanwise axial positions are shown in figure 5.25. The left panel of this figure shows that at the hub the differences of axial positions are between roughly −1 and 2. This means that most of the manufactured blades have their LE in front of the nominal CAD design. Furthermore it can be seen that the axial position differences decrease with the span. This can most clearly be seen by the mean and the median lines, which have negative slopes along the span. Therefore, at the tip most of the manufactured blades have LE in front of the intended design position.

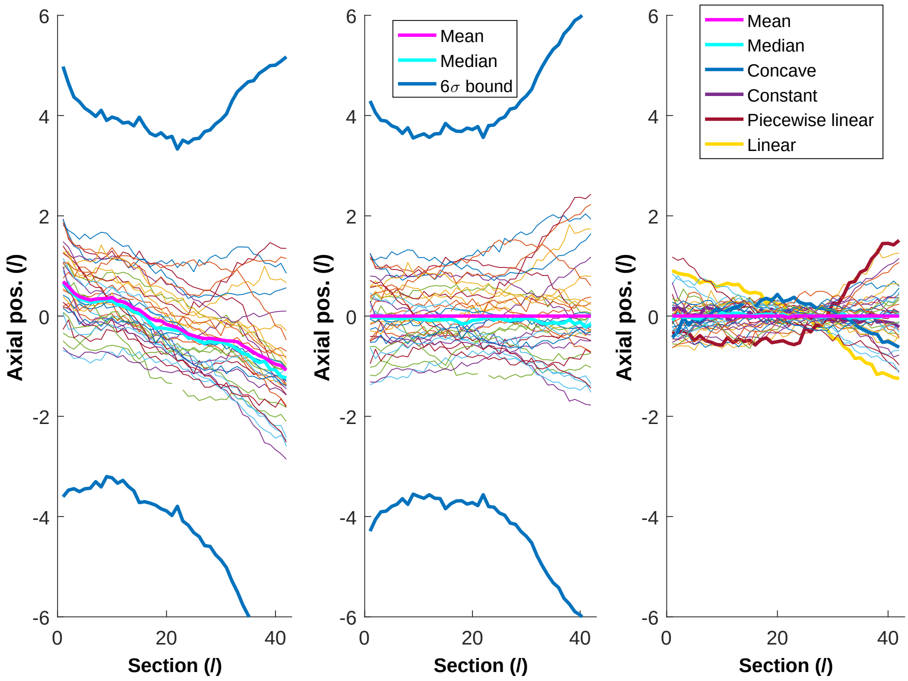

Figure 5.25: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design axial position.

It can also be seen that the scatter at the tip is larger than at the hub. This, along with the spanwise change of scatter, can be better seen in the centre panel. This panel shows that the manufacturing scatter about the mean differences is roughly within ±1. Furthermore, it can be seen that the scatter increases from about midspan towards the tip. The spanwise differences between the mean and the median lines can also be seen in this panel. Specifically, it can be seen that the median line lies slightly below the mean line for the entire span, however the difference appears negligible.

The right panel shows that after the means are subtracted from the centre panel, the scatter in the differences is considerably smaller. It can also be seen that from the hub the scatter due to the behaviour of the lines decreases along the span up to roughly section 25. After section 25 it increases rapidly towards the tip. This shows that the blades have different slopes of the LE in the axial direction to one another. This is caused either by the positioning of the blade by the root or through imperfections in the manufacturing of the blade.

The right panel of figure 5.25 also has several blade difference profile lines highlighted. These lines show the different observed profiles. Note that most of the line showing constant behaviour is directly underneath the mean and the median lines, and therefore cannot be seen. The lines that feature constant or linear behaviour could possibly attribute it to the root deviations. On the other hand the difference profiles such as the highlighted piecewise linear or concave profiles can only owe their shape to manufacturing imperfections of the blade itself. Under the assumption that the blade roots are manufactured perfectly, made in subsection 5.4.2, the absolute differences shown in the left panel of figure 5.25 are due to the manufacturing process of the blade itself.

Tangential position

The manufacturing deviations of spanwise tangential positions are shown in figure 5.25. The left panel of this figure shows that at the hub the differences of tangential positions are between 0 and roughly 2.5. This means that all of the manufactured blades have their LE offset in the direction of the suction side relative to the the nominal CAD design. Furthermore it can be seen that the tangential position differences lower with the span. This can be seen most clearly by the mean and the median lines, which have negative slopes along the span.

Furthermore, it can be seen that the scatter approximately doubles along the blade span and is within roughly −2 and 3 at the tip. This can best be seen in the centre panel. In fact, the scatter increases especially in the last 25% of the span, in the tip region. Lastly, it can be seen that the median line lies almost directly on top of the mean line. The differences between these two lines are very near to zero, and remain so along the entire span.

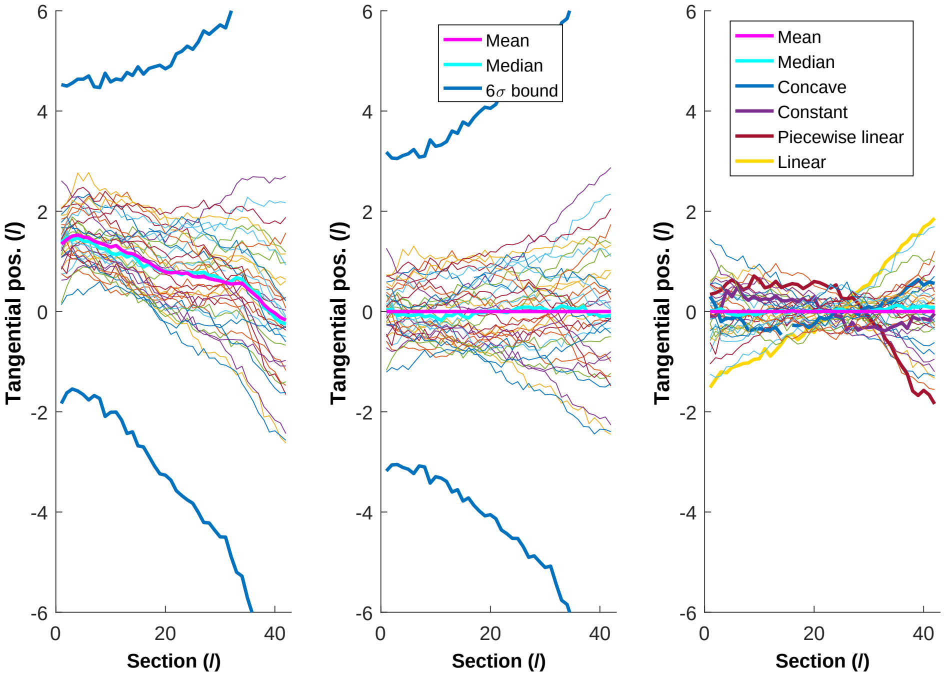

Figure 5.26: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design tangential position.

Similarly to the discussion of figure 5.25, the right panel of figure 5.26 shows that a large part of the variation is due to the profiles of the blades. The highlighted lines correspond to that ones highlighted in figure 5.25. These show that the blade-to-blade differences of axial and tangential position have the same spanwise behaviour, with the exception of the line named "constant". This line now appears to have a negative linear behaviour with the span.

The blade manufacturing imperfections therefore had a 3D impact on the LE contour, as the behaviours in the axial and tangential directions are the same. In the case of linear and constant behaviours this impact is merely an offset, or lean of the LE contour in some direction. In the case of piecewise linear and concave profiles the shape of the blade itself is imperfect.

Like in the case of the axial position, the blade-to-blade differences shown in the right panel of figure 5.26 are definitely due to aerofoil section manufacturing imperfections. Furthermore, under the assumption of dovetail manufacturing, made in subsection 5.4.2, all the differences, seen in the left panel of figure 5.26, are due to aerofoil section manufacturing imperfections.

Stagger angle

Unlike the axial and tangential positions, where the general behaviour of the differences along the span was linear, the stagger angle shows a clear dip in the first half of the span, near the hub. Furthermore, it can be seen that most of the stagger angle parameters are below 0. This can be seen more specifically by noting that the mean and the median lines both fall well beneath 0 for all spanwise positions. It can also be seen that the scatter increases along the span.

This is best illustrated in the centre panel. It shows that the scatter roughly doubles along the span. The blade-to-blade differences of stagger angle fall within ±0.1. Furthermore, it can be seen that the scatter increases roughly linearly from the hub towards the tip. The differences between the mean and the median lines are again roughly 0 all along the span, with the median line lying slightly beneath the mean one.

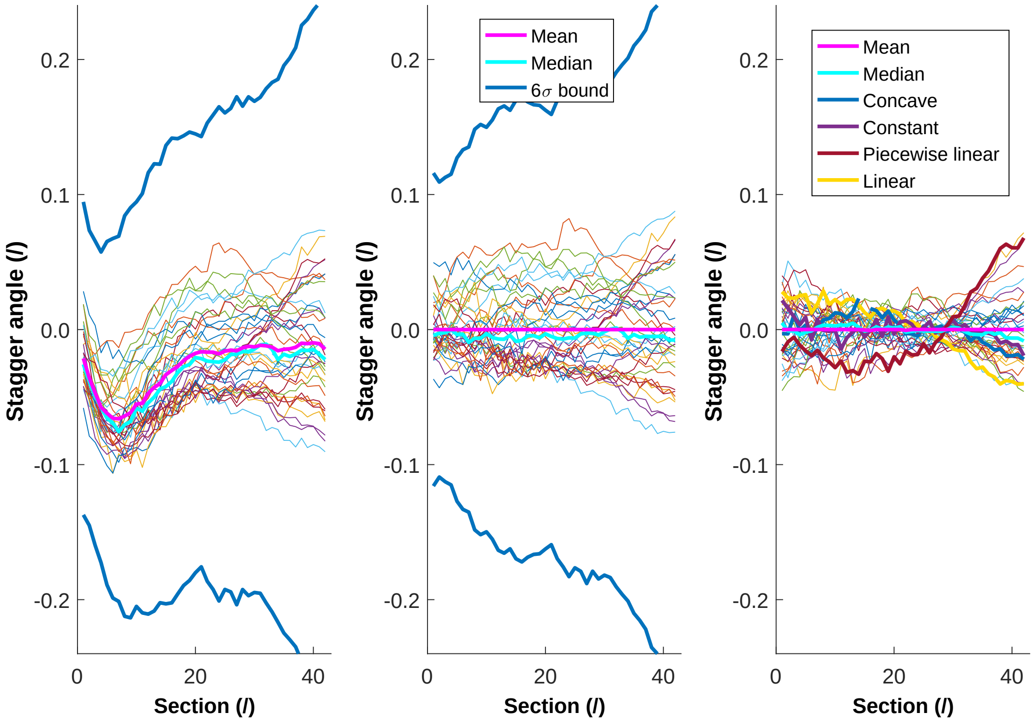

Figure 5.27: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design stagger angle.

The right panel of figure 5.27 shows that a large part of the scatter can be attributed to offsets between the lines. However, it can also be seen that in the middle third of the blade span, between sections 15 and 30, the scatter due to the behaviour of individual lines is at its minimum. Outside of this region it increases significantly toward the hub and the tip. This shows that the behaviours also contribute a significant amount to the scatter.

The behaviours of individual lines are again highlighted to show the differences between them, and to compare to the findings of figures 5.25 and 5.26. The highlighted lines correspond to the blades whose lines have been highlighted in those figures.

Note that a direct comparison between these lines is not possible, as the mean behaviour that was subtracted in the case of the right panel of figure 5.27 is not approximately linear. In this case the subtraction of the mean behaviour imposes some general behaviour on the blade-to-blade differences as well. Despite this, the major features of the individual highlighted lines were maintained. It can be seen that they have either a concave or convex feature, most of them near the hub. These behaviours are due to the manufacturing imperfections of the blade itself.

Chord length

The axial and tangential positions, and stagger angle are all positioning parameters and do not have an impact on the size of the actual blade. Among the considered NACA like parameters it is the chord length and maximum thickness that control the size of the blade. If the blade is imagined as a simple cube, the chord line would be its "depth" and the maximum thickness would be its "width".

The chord length deviations due to the manufacturing scatter can be seen in figure 5.28. The left panel of this image shows that virtually all of the chord lengths differences fall below 0. This means that the manufactured blades are shorter than the design intent. Furthermore, it can be seen that the chord length of the manufactured blades drops sharply within the first few sections. This is likely an impact of the fillet, even though the sections used in the parametrisation were selected to lie above it.

After the initial drop the chord lengths appear to stabilise and remain constant up until the tip. Furthermore, the scatter reduces slightly towards the tip. This is illustrated by the 6σ bounds both in the left and centre panels. The centre panel further shows that the mean and median lines are again very close together, and the differences between them can be neglected.

It is interesting to note that visually most lines lie above the mean and median lines in the centre panel. Most of the lines indeed lie near or above the line. However this is offset by the few lines that lie roughly at -2.

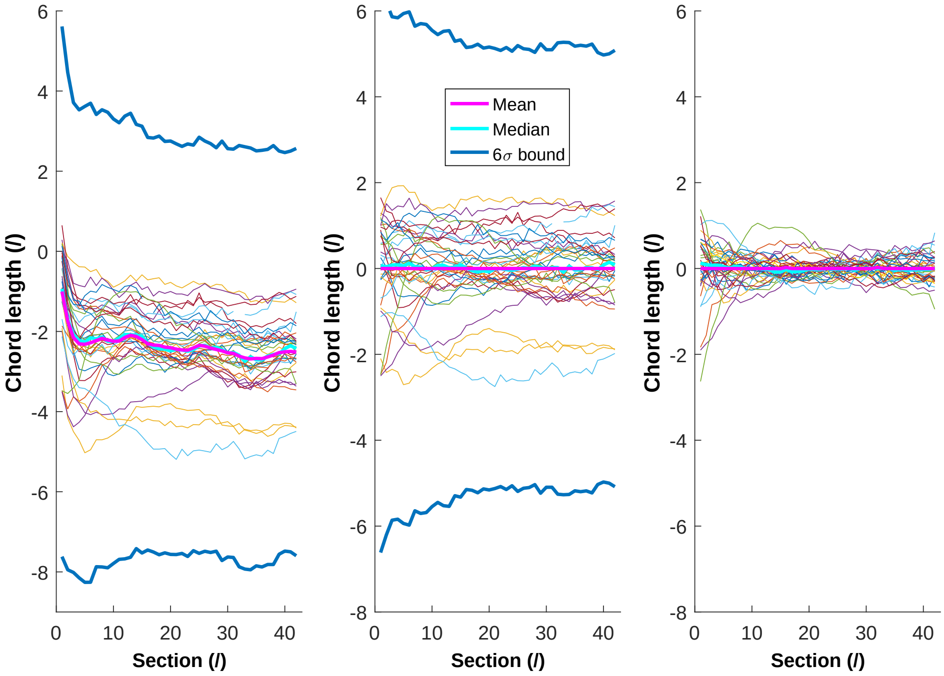

Figure 5.28: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design chord length.

The right panel of figure 5.28 shows why the scatter of chord length is increased near the hub, as seen in the left panel. Most lines have constant behaviour along the span and therefore remain within the scatter seen in the upper half of the blades also near the hub. A few blades have different behaviours near the hub however, which increases the scatter. As mentioned above, this behaviour is due to the transition between the fillet and blade LE and TE geometries. Some blades have already performed the transition, or most of it, before the evaluated sections (lines that originate from -2), whereas some are still doing the transition (lines that originate from 1).

Furthermore, the scatter seen in the right panel is significantly smaller than the one seen in the left panel. Therefore the behaviours of the differences profiles have relatively small contributions to the absolute differences. Lastly, the lines highlighted in figures 5.25, 5.26, and 5.27 are highlighted in the right panel as well. In the case of the chord length the behaviour seen in the previous figures is not retained.

Thickness distributions

The thickness distribution holds the information of the thickness of the blade along its chord. As such, this distribution contains all the extracted information on the LE and TE thickness, maximum thickness, etc. The comparison of the measured thickness distributions and the nominal thickness distributions can be seen in figure 5.29. In this figure the thickness distributions for all the spanwise sections are shown at once.

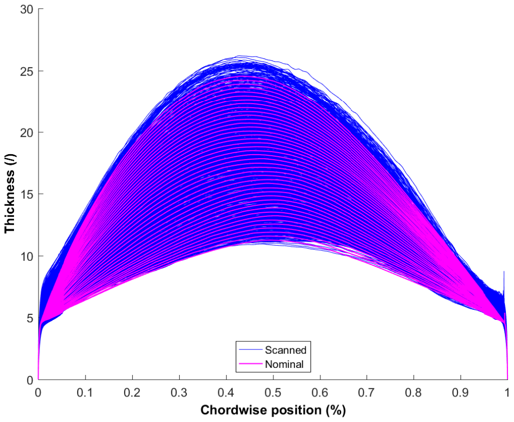

Figure 5.29: The comparison of measured manufactured and nominal thickness distributions for all the spanwise sections.

Figure 5.29 shows the measured distributions in blue and the nominal distributions in magenta. It can be seen that as a whole, the blue lines generally lie on top of the magenta lines, which is most obvious at the peak of the distributions. This means that the manufactured blades are thicker than the nominal blade. Furthermore, it can be seen that the behaviour of the blue and magenta lines is different in the LE and TE regions. This means that the manufactured blades generally have thicker LE and TE geometries.

To better see the differences between the measured and the nominal distributions, they are shown in figure 5.30. This figure shows the differences of the spanwise sections near the hub plotted in blue colours, and the sections near the tip in red colours. It also shows that all the distribution differences follow a wavy behaviour along the chord, which is due to imperfections during manufacturing. These forging imperfections could be due to cast imperfections or due to different cooling rates along the chord.

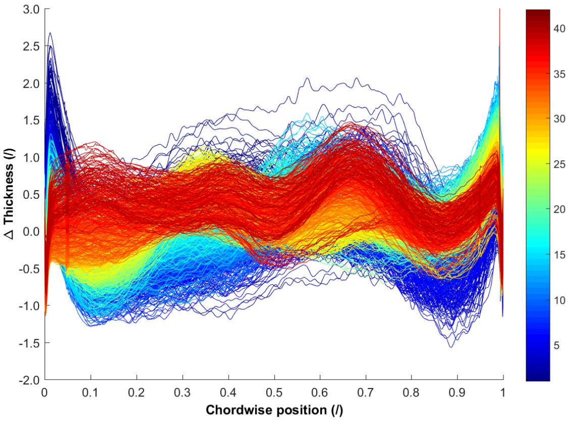

Figure 5.30: Differences between the measured manufactured and the nominal thickness distributions for all spanwise sections.

It can be seen that the thickness distribution deltas are roughly centred about 0, when all the distributions are considered together. Focusing on the distributions near the hub it can be seen that many of them lie beneath 0 for the majority of the chord. The opposite is true for the distributions near the tip. Furthermore, it can be seen that the manufactured thickness distribution differences near the hub are more scattered than those at the tip.

Camber distributions

The camber distribution holds the information of the blade curvature along its chord. As such, this distribution contains all the extracted information on the LE and TE angles, maximum camber, etc. The comparison of the measured camber distributions and the nominal thickness distributions can be seen in figure 5.31. In this figure the camber distributions for all the spanwise sections are shown at once.

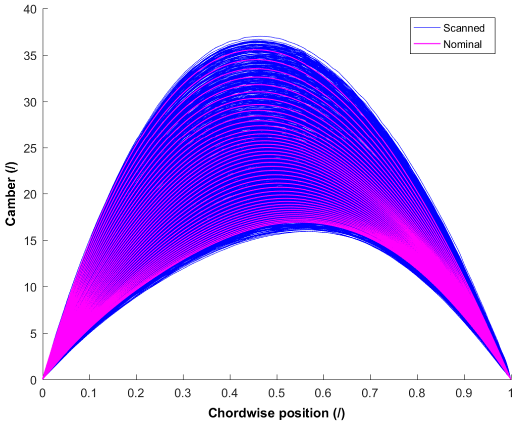

Figure 5.31: The comparison of measured manufactured and nominal camber distributions for all the spanwise sections.

Figure 5.31 shows that the manufactured camber distributions lie both beneath and above the nominal camber distributions. At the hub the scanned camber distributions are below the nominal ones, and at the tip they are above. This means that some of the manufactured blades are curved more at the tip, and some lass at the hub, than intended.

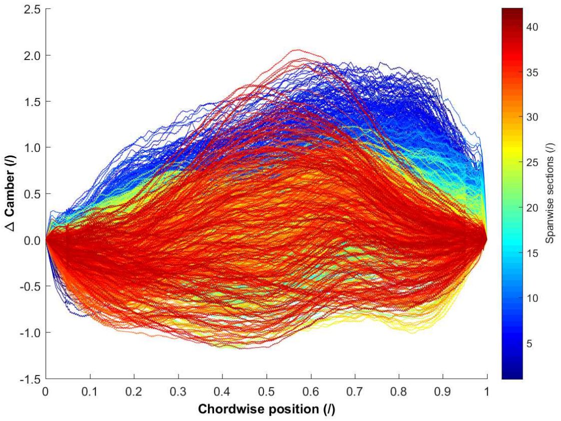

Figure 5.32: Differences between the measured manufactured and the nominal thickness distributions for all spanwise sections.

The actual differences between the measured manufactured and nominal camber distributions can be seen in figure 5.32. This figure shows the differences of the scanned and their corresponding nominal camber distributions. The scatter of the camber distribution differences is lowest near the LE of the blades, and highest near roughly 70% of the chord lengths. Furthermore, the positions of the maximum camber distribution differences are moving from near the TE for the the hub sections, towards the mid-chord for the tip sections. At the mid-chord location the tip sections have the highest scatter, as they span from -1 to $+2$.

Maximum thickness

Another important parameter is the blade maximum thickness, as it gives an indication on the compressor throat area. Figure 5.33 shows the differences of maximum thickness due to manufacturing imperfections. The left panel of this figure shows a wavy behaviour of maximum thickness along the span. Furthermore, a large dip can be seen in roughly the first third of the span. Furthermore, the differences can be seen to be roughly within ±1.

The central panel shows that the blade-to-blade scatter is high at the hub, decreases towards the midspan and subsequently increases again towards the tip. However, it is interesting to note that the blade-to-blade scatter at the hub is high largely due to two lines originating at around -2. Comparing the shapes of these lines to the shape of the mean line from the left panel it can be seen that they are very similar, but inverted. Therefore the lines that are causing the high scatter in the centre panel actually have more or less perfect maximum thickness behaviour along the span.

Furthermore, excluding these lines and the two lines with a that originate at roughly $+0.8$, the lines lie roughly within -0.2 and $+0.3$ at section 1. Again excluding the aforementioned lines, it can also be seen that the scatter of blade-to-blade differences increases with the span. The difference between the mean and the median lines is more or less 0 along the entire span and can be neglected.

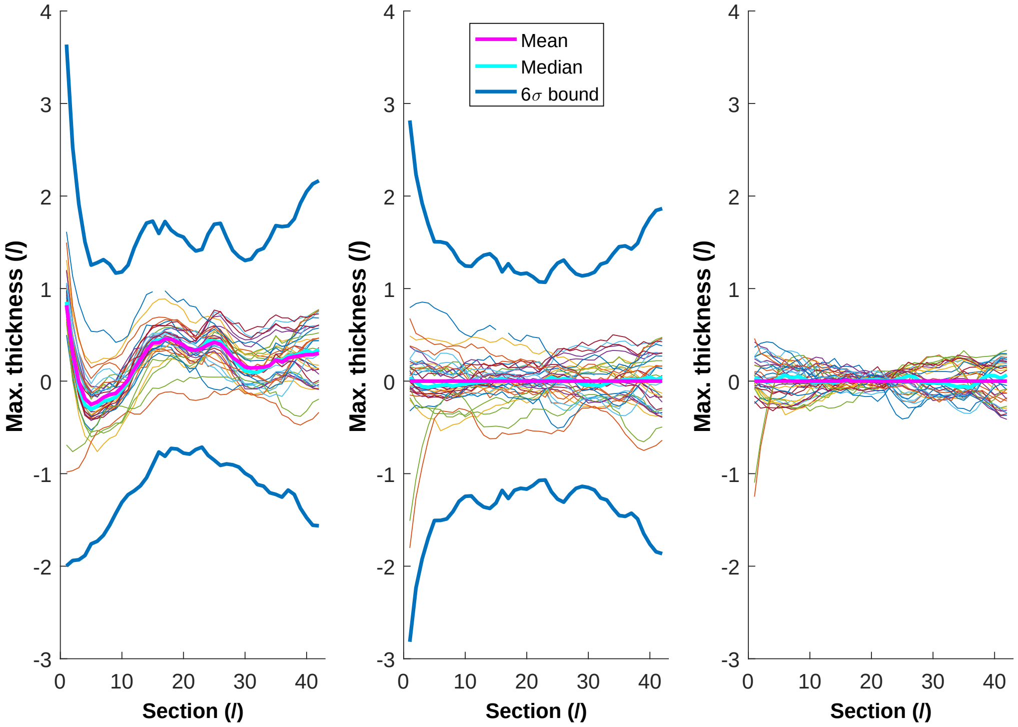

Figure 5.33: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design maximum thickness.

The right panel of figure 5.33 shows the behaviours of the lines without the mean spanwise offsets. It can be seen that the impact of behaviours has a large impact on the blade-to-blade differences. Furthermore, in this panel, at around section 30, it can be seen that two groups have formed. One of the groups is above the mean and the median lines in this region, whereas the other group is just below these lines. In order to see if there are other groups, the lines with similar behaviour have been plotted with similar colors in figure 5.34.

Figure 5.34 shows that five major groups could be identified. There are three groups with relatively many lines (≈ 10), and two groups with relatively few. The groups show different spanwise behaviours, with some groups decreasing with the span and others increasing or constant along the span.

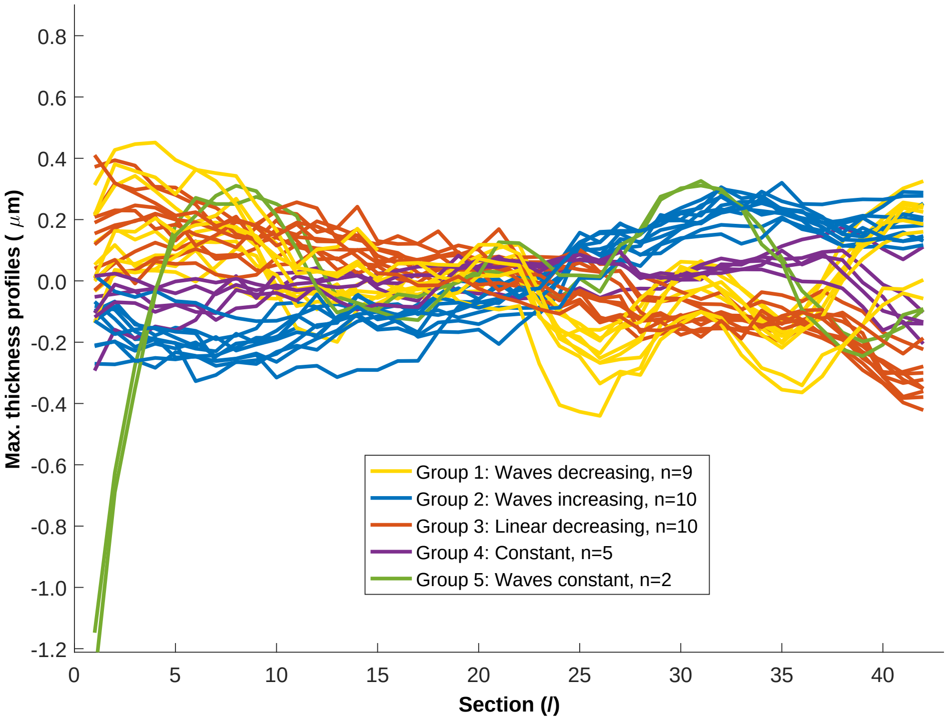

Figure 5.34: Groups of spanwise max thickness difference behaviour.

Due to the consistency of the difference behaviours within individual groups it is likely that the groups are a consequence of manufacturing imperfections. Specifically, in the case of forged blades, the waviness could be due to imperfections of the cast and any variability of the cooling rates of different regions of the blade. Most of the groups have all their values within roughly -0.2 to $+0.4$, which is the impact the different behaviours can have on the absolute differences.

Maximum camber

The maximum camber is a measure of the maximal asymmetry of the blade sections suction and pressure sides, and has a direct impact on the maximum lift coefficient of the section. Its deviations due to the manufacturing imperfections can be seen in figure 5.35.

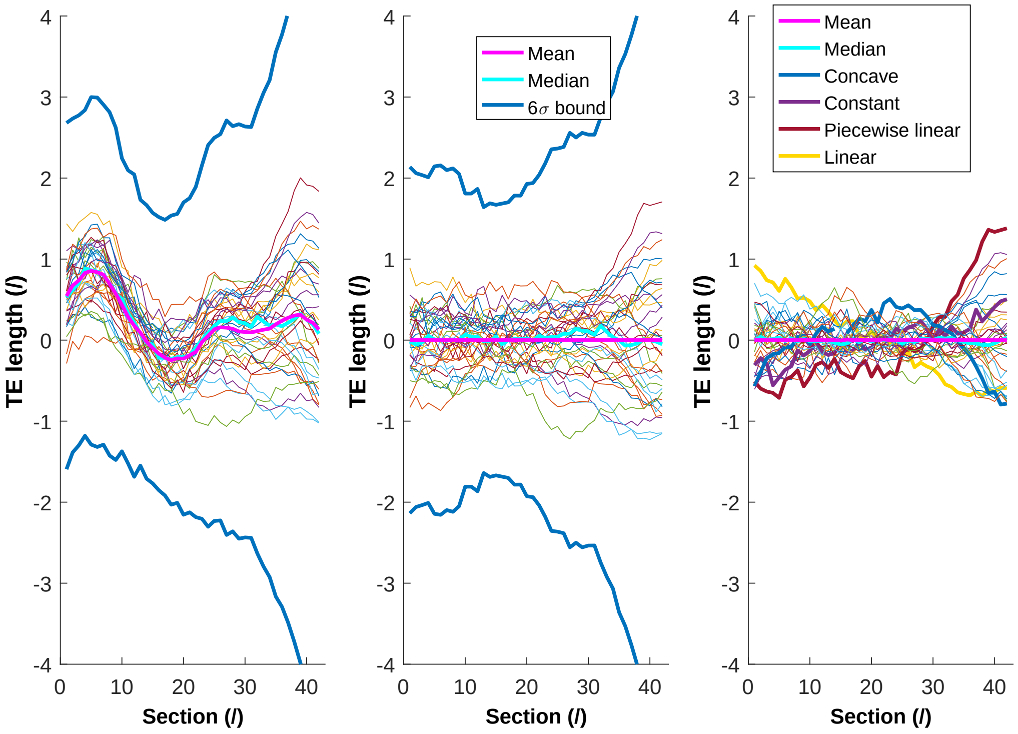

Figure 5.35: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design maximum camber.

Left panel of figure 5.35 show that the maximum camber differences also have a wavy behaviour, and a peak and a valley can be identified in the first half of the span. Furthermore it can be seen that the scatter increases significantly from the hub to the tip. This is again best seen in the centre panel.

The centre panel shows that the differences are between ±1 for the majority of the span and increase significantly in the last quarter of the span towards the tip. It can be seen the the differences are centred around 0 and the differences between the mean and median lines are nearly 0 along their entire span.

The right panel shows that removing the spanwise means decreases the scatter at the tip. However, a significant part of the scatter remains and is due to the behaviours of the individual lines. In order to analyse this behaviour several lines were highlighted as well. These are the same lines as those in figures 5.25, 5.26, and 5.27. It can be seen that the behaviour observed in these figures is also present here for the same blades.

LE angle

The leading edge angle deviations due to the manufacturing imperfections can be seen in figure 5.36. First of all the left panel shows that the LE angle lines are very noisy which is expected to be an artefact of the calculation of the angles as opposed to the measured geometries being noisy. Despite the noise, it can be seen that some sort of spanwise behaviour is present. This can be seen by observing the mean and median lines.

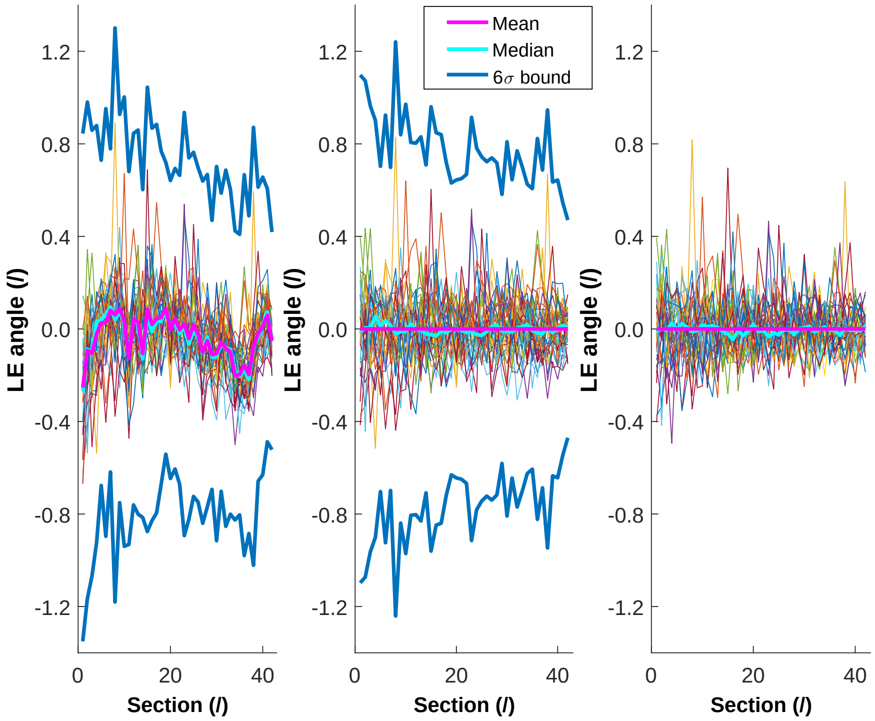

Figure 5.36: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design LE angle.

The left panel shows that the deviations are roughly contained within a range of 0.8 along the entire span. This is more clearly seen in the centre panel, where virtually all the blade-to-blade differences fall within ±0.4. Furthermore, it can be seen that the scatter reduces from the hub to the tip.

The right panel shows that the noise has the largest impact on the blade-to-blade differences seen in the centre panel. The plot in this panel is virtually identical to the one in the left panel, suggesting that any signal hidden in the data is centred about 0 as well. However, any signal is simply lost in the noise.

To try and reduce the impact of the noise the LE angle data was smoothed using a Savitzky-Golay filter [54]. The results can be seen in figure 5.37. Compared to figure 5.36 the left panel of figure 5.37 now shows the spanwise behaviour more smoothly. This behaviour shows that the manufactured LE angles are mostly smaller than the design intent. This means that the camber at the LE is shallower than intended.

Looking at the centre panel, it can be seen that the blade-to-blade differences are well within roughly ±0.4. The behaviour along the span is still quite noisy, however there seems to be a small reduction of scatter along the first third of the span. Like in the case of maximum thickness, the scatter is mostly high due to a few blades that have LE angles close to the design intent.

This can be seen more clearly in the right panel, which shows the behaviour of the lines. The lines can be seen to fan out towards the hub. Note also that the scatter decreased somewhat by removing the means of the lines.

The noise observed in figure 5.36 could be due to the blade LE shape being different to the design intent, or the blade actually being differently curved in this region. The former is due more to the actual LE shape, whereas

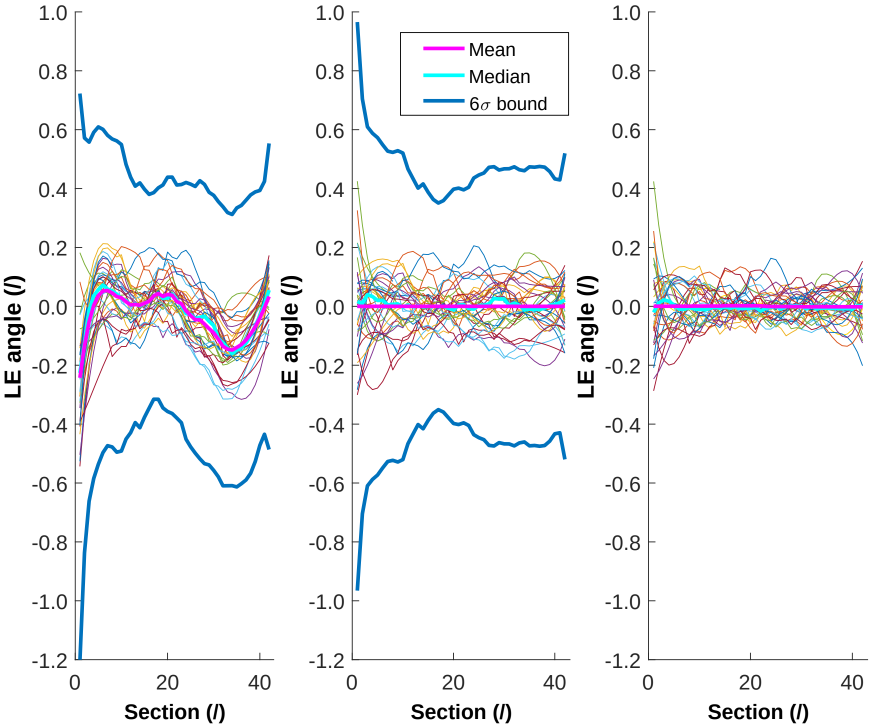

Figure 5.37: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design LE angle smoothed by a Savitzky-Golay filter.

the latter is more in line with the meaning of the LE angle. This is supported by the fact that the LE angles of the nominal blade are much less noisy than the LE angles of the scanned blades. To assess this the leading edges of the scanned blades and the nominal CAD design are plotted in figure 5.38. This image aims to show the reasons for the LE scatter at a single spanwise section.

Figure 5.38 shows the mid-span (section 21) LE geometries of all the scanned manufactured blades as well as the nominal blade. For ease of comparison these geometries have been rotated to a vertical position and aligned by the right contour of the geometry as well as their highest point. This figure most notably shows that at mid-span the manufactured blades are thicker than intended. Because of this, the radius of the LE changes as well. More importantly, it can be seen that the manufactured blades have a different curvature of the LE, with some blades achieving the highest point to the right of the nominal CAD design highest point. The scanned blades therefore appear to be slightly more pointy. Small differences in the pointiness could potentially cause larger deviations of the LE angle than warranted.

Furthermore, it can be seen that the left and right contours of the scanned blade LE geometries are not parallel to the nominal CAD design ones. At the left contour the difference between the highlighted blades at about where the axis label is located is slightly larger than at where the legend is located. The same can be observed at the left contour where the yellow line slowly crosses back over the brown line.

Figure 5.38 explains the variation in LE angles for a single section and figure 5.36 shows the magnitude of this variation section by section. It is unlikely that the blades would be manufactured with such high differences of the blade LE and TE angles between individual spanwise sections. Furthermore, such differences would also appear as features on the surface deviation plots if they were real, but they do not.

Figure 5.10 showed that even a small change in alignment caused high noise of the TE angle. In that case a single blade was parametrised twice, with only the alignment being different between the cases. As the alignment only changes the orientation of the blade it can be expected that any systematic differences would cancel, or at least dampen out one another. The amount of noise remaining in figure 5.10 supports the idea that the calculation process of the LE and TE angles is sensitive to small changes of the geometry.

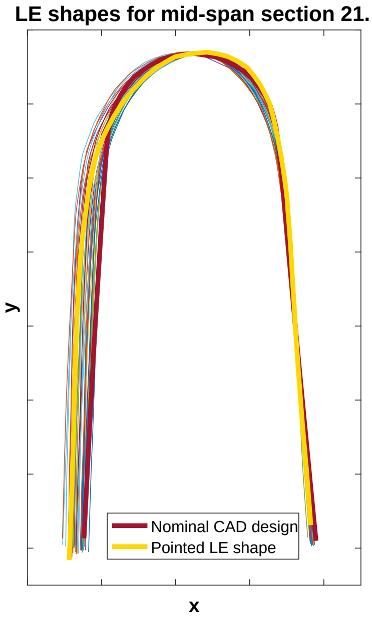

Figure 5.38: LE shapes rotated and grouped for ease of comparison.

Lastly, it should be noted that the LE angles are calculated on the basis of the camber distribution along the first 1.5% of the chord. For blades with chord length of about $4cm$ the deviations are therefore calculated based on data corresponding to about $600\mu m$ of the chord. The maximum LE and TE angle error that could be induced by the measurement uncertainty of the these regions of the blade is therefore about $E = tan\left(\frac{30\mu m}{600\mu m}\right) \approx 2.86^o$, which is about twice the scatter seen in figure 5.36.

Any significant blade specific spanwise behaviour and the behaviour of the 6σ bounds along the span can therefore be attributed to the manufacturing deviations, but the rest of the behaviour is most likely due to noise. In this specific case the behaviour due to manufacturing variability can be seen in the left panel of figure 5.37. The blade-to-blade offsets that can be determined by comparing the centre and right panels of figure 5.37 are most likely due to noise. If these offsets were compared with the mean of 42 points sampled from a normal distribution $D \sim \mathcal{N}\left(\mu = 0, \sigma = 1\right)$, it would be seen that the orders of magnitude are the same. Therefore, for this specific set of blades, the blade specific spanwise offsets can be attributed to the noise.

TE angle

Similarly to the LE angle, the calculated TE angles are noisy, however in the case of the TE angles the noise is larger. The differences between TE angles of manufactured scans and the nominal CAD design can be seen in figure 5.39. Like for the case of the LE angles, some behaviour can be seen despite the noise. Specifically, in the first half of the span, near the hub, there is a hump in the mean and median lines with a peak of roughly 1.

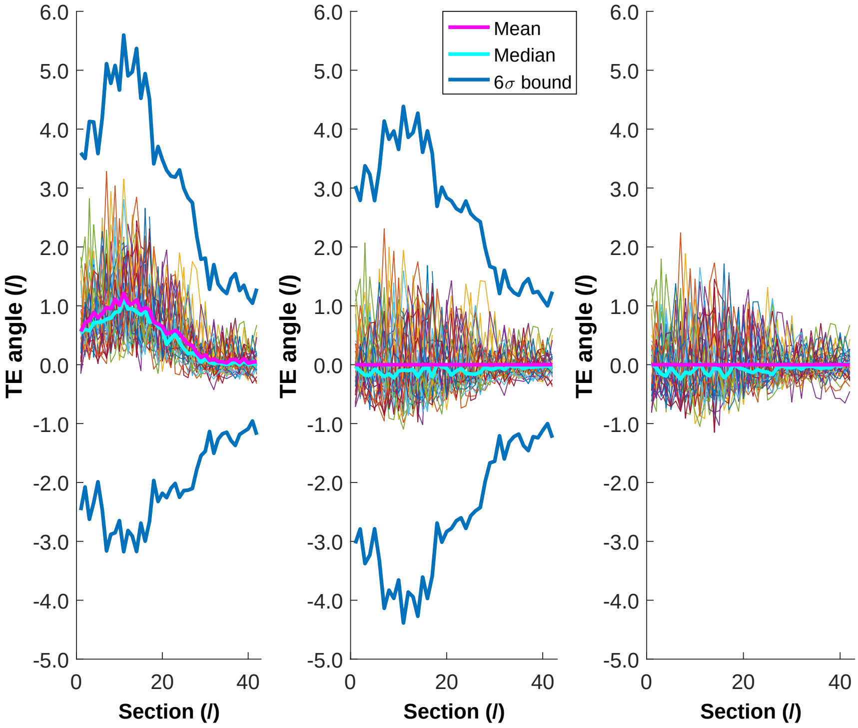

Figure 5.39: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design TE angle.

Furthermore, the left panel of this figure shows that the noise is now roughly within -0.4 and $+3$. The scatter can also be seen to increase up until approximately one third of the span, after which it decreases drastically towards the tip. This can be seen well in the centre panel. It is interesting to note that the scatter is roughly the same for points with the same TE angle mean offset in the left panel of figure 5.39. This suggests that there is a correlation between the size of the differences and their scatter. Lastly, the mean and median lines are still close together, with the median line lying slightly below the mean line.

The right panel of figure 5.39 shows practically the same scatter as the left panel. This suggests that the noise is the main cause of the blade-to-blade differences. In order to better understand the spanwise behaviour of the differences, and to better visualise the hump in the bottom half of the blade, the Savitzky-Golay filter has been applied to the data, as has been done for the case of LE angles in figure 5.37. The smoothed TE angles can be seen in figure 5.40.

Left panel of figure 5.40 shows that most blades feature the TE angle hump, however some do not, for example the green line originating at about 1.6. Contrary to the centre panel of figure 5.39, the centre panel of figure 5.40 shows the scatter decreasing along the entire span. This is a consequence of the filtering procedure.

Furthermore, the right panel is virtually the same as the left panel. Therefore it is reasonable to conclude that any blade-to-blade behaviour that might be present is lost in the noise. Furthermore the same arguments that were presented to explain the noise for the case of LE angles apply directly here as well, as the calculation procedure is the same.

The TE angle hump appears to be a significant blade specific spanwise behaviour. It can therefore be seen as a

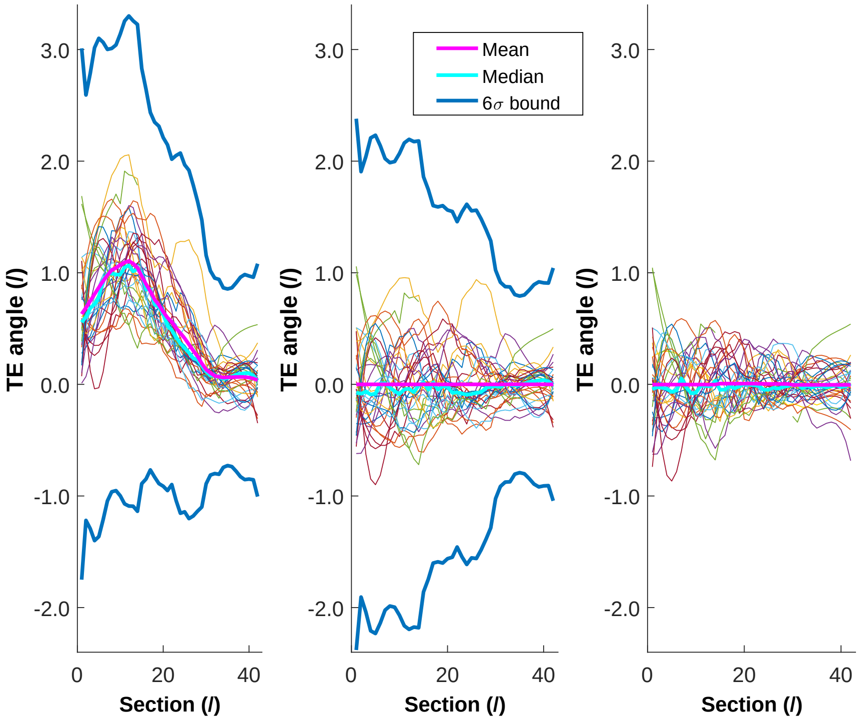

Figure 5.40: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design TE angle smoothed by a Savitzky-Golay filter.

systematic effect of the manufacturing deviations. The rest of the behaviour is most likely due to noise. The blade-to-blade offsets that can be determined by comparing the centre and right panels of figure 5.40 are most likely due to noise. For this specific set of blades, the blade specific spanwise TE angle offsets can be attributed to the noise.

LE thickness

Another important parameter of the LE is its thickness. Sharper LE edges help increase the performance in the design points, whereas blunter LE edges are able to operate in a wider range of inflow angles at the price of reduced efficiency.

Figure 5.41 shows the impact of manufacturing on the LE thickness. In the left panel it can be seen that the LE thickness differences originate at roughly 3 at the hub and drop to about 0 at the tip. The differences have a wavy behaviour that includes an inflection point at roughly midspan.

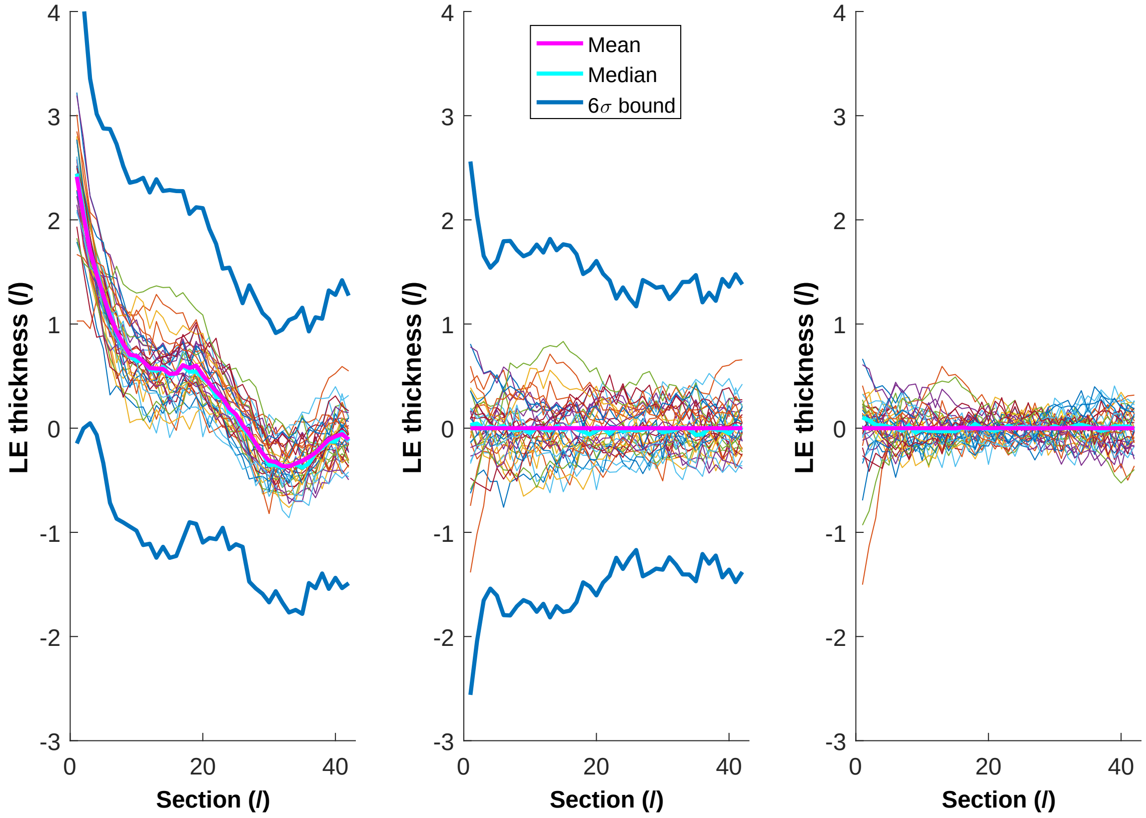

Figure 5.41: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design LE thickness.

The scatter appears to follow the behaviour of the differences and reduces somewhat from hub to tip. This can be seen in the centre panel. It is again interesting to note that the large scatter at the hub is due to one line that originates at about 1 on the left panel.

The right panel of figure 5.41 shows that the largest impact on the blade-to-blade differences are due to the behaviour of individual lines. The same blades as highlighted in previous figures are highlighted here as well, however the same behaviours as seen in figures 5.25, 5.26, and 5.27 are not present here.

TE thickness

TE thickness controls the flow as it is leaving the blade. Thicker TE mean that the wakes behind them will be larger, which in turn reduces the efficiency of the compressor.

Figure 5.42 shows the impact of manufacturing on the TE thickness. In the left panel it can be seen that the TE thickness differences originate at roughly 1 at the hub and drop to about 0 in a dip within the first 25% of the span. The differences then achieve a peak at about 30% of the span, after which they drop linearly toward the tip.

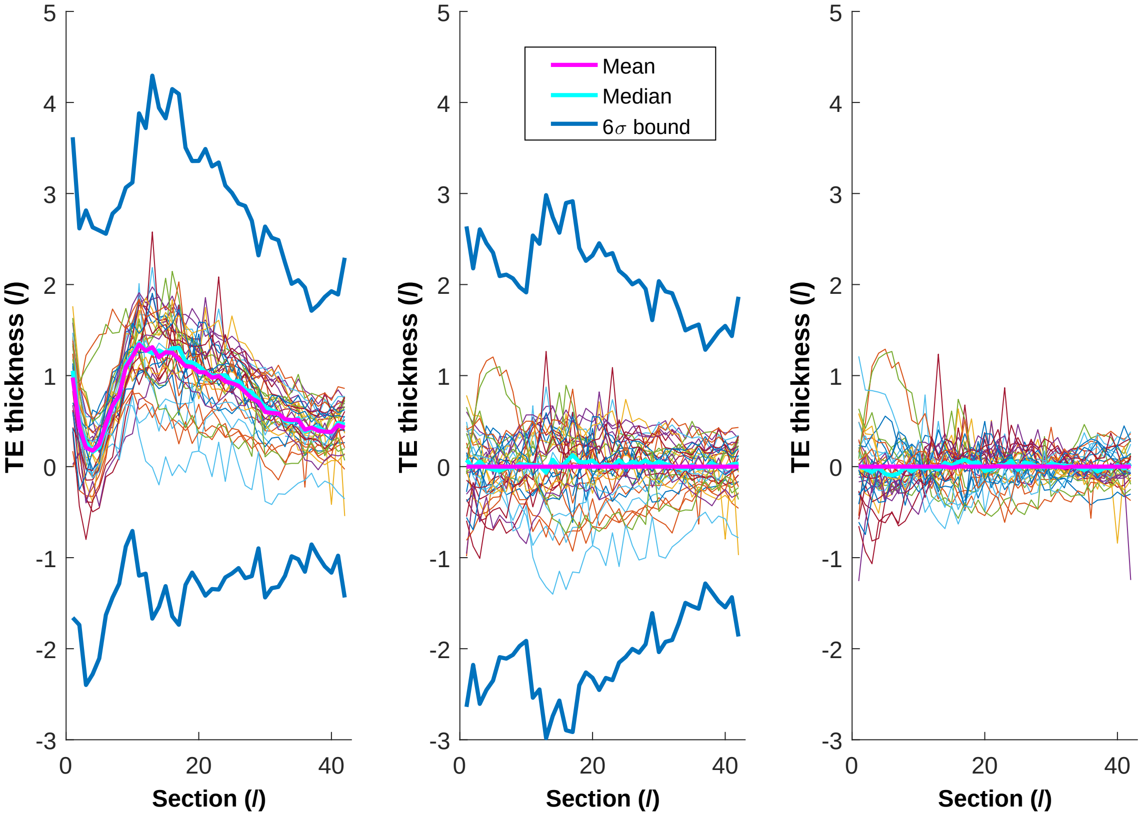

Figure 5.42: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design TE thickness.

The centre panel shows that the differences between the mean and the median lines can be neglected. Furthermore, the centre panel shows that the scatter reduces from the hub towards the lowest point in the dip, and again rises towards the peak. This suggests that whatever effect that caused the increase of the differences also caused the increase of the scatter.

The right panel shows that some of the blade-to-blade scatter is due to offsets. However, a large portion of it is due to the behaviour of the individual lines. Lastly, the lines in the right panel also demonstrate the noise of the TE thickness differences. As the LE and TE thicknesses are calculated as thicknesses at a specific hardcoded curvature of the LE and TE edge, any differences in the shape of these geometries might produce non-representative thicknesses. The differences in the shape refer to pointiness of the LE contours as shown in figure 5.38.

LE position

Position of the LE is a parameter that is used to determine at what percentage of the chord length the LE begins. These are the positions where the LE thickness has been calculated.

The impact of manufacturing variability on the LE position can be seen in figure 5.43. It can be seen that the lines in the left panel closely follow the shape seen in the left panel of figure 5.41. It is therefore clear that in this particular case the LE position and thickness are correlated. This is supported by figure 5.28 which shows that the chord lengths are more or less constant along the span. As the LE thickness decreases, so does its length. Therefore the LE at the hub are larger than the LE at the tip.

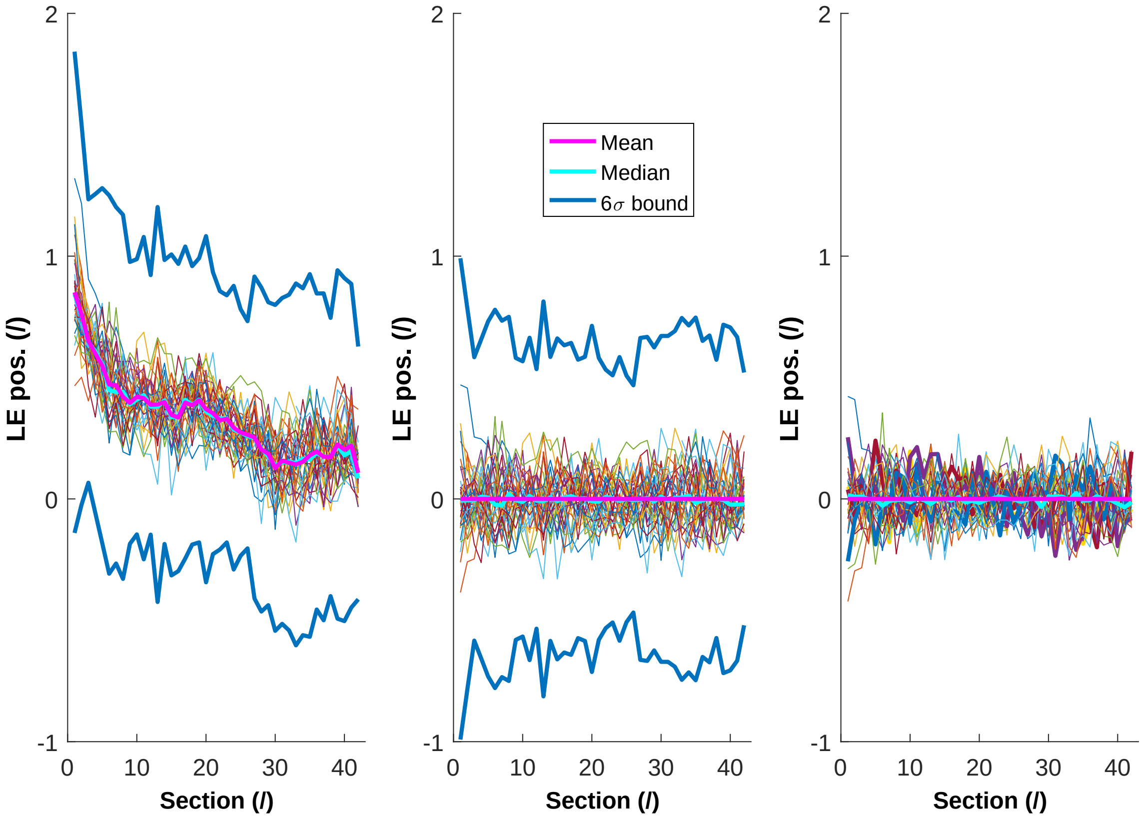

Figure 5.43: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design LE position.

The centre panel of figure 5.43 shows that the blade-to-blade differences are quite noisy. Furthermore, it shows that the blade-to-blade scatter is slightly decreasing towards the tip. Comparing the centre plot to the right plot no real differences can be seen. Any blade-to-blade behaviour is simply lost in the noise.

TE position

Position of the TE is a parameter that is used to determine at what percentage of the chord length the TE begins. These are the positions where the TE thickness has been calculated.

The impact of manufacturing variability on the TE position can be seen in figure 5.44. A comparison of this figure to figure 5.42 shows that the behaviours are exactly the opposite. This means that the TE position and thickness are correlated, which is also supported by the constant chord length shown in figure 5.28. Note that the TE position is measured as a percentage from the TE. An increasing TE position therefore means a reduction of the TE size, which is what can be seen in figure 5.44. Lastly, any blade-to-blade differences are lost in the noise, as a comparison of the centre and right plots shows no noticeable differences.

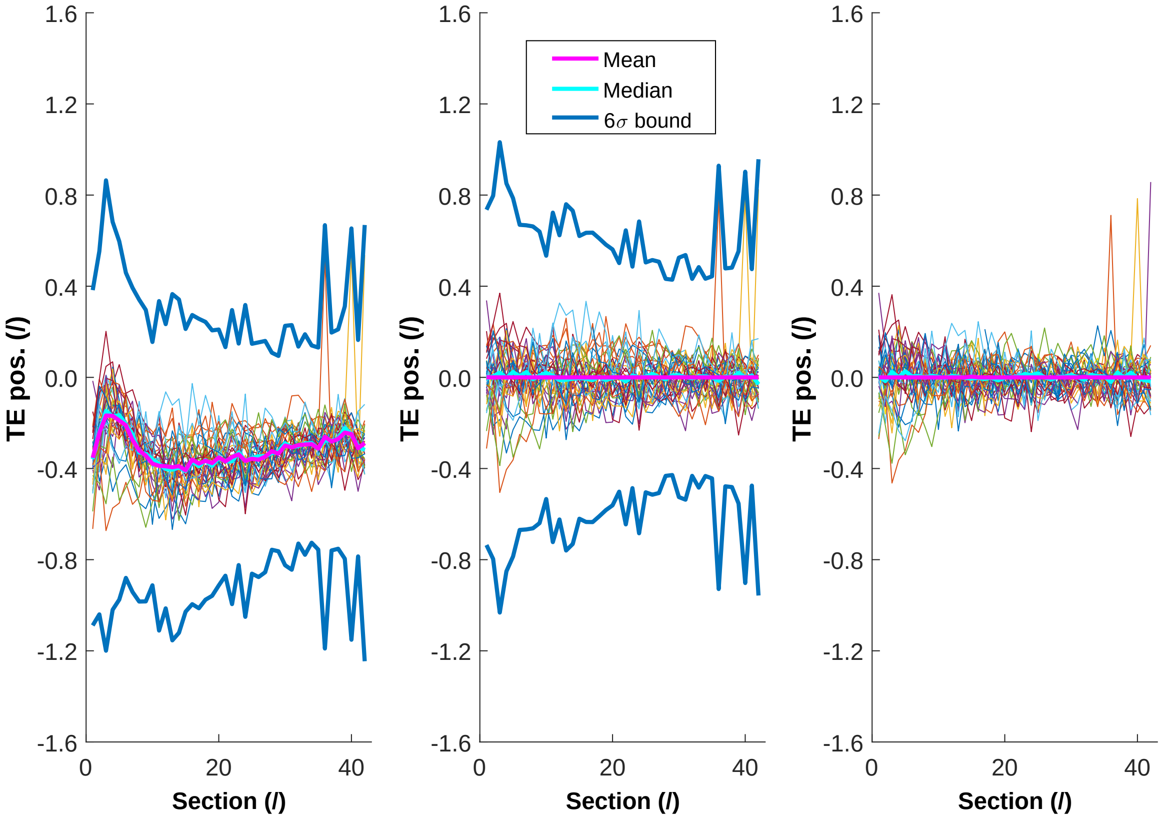

Figure 5.44: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design TE position.

Position of max. thickness

The position of maximum thickness helps compare the thickness distributions of multiple blades. The impact of manufacturing variability on position of maximum thickness can be seen in figure 5.45.

In the left panel of this figure it can be seen that the manufacturing deviations of the position of maximum thickness follow a wavy behaviour. Note that the difference of the maximum thickness positions are relatively large. The maximum deviations from the nominal design are about 10, at the peak at roughly 30% of the span. The lowest values are about -6, at the valley at roughly 60% of the span. This is ignoring the values that occur near the tip where the scatter is large, and focusing on the values that most blades will typically achieve.

Note that for a blade with a chord length of approximately $40mm$ this chord length variation translates into an absolute variation which is approximately an order of magnitude larger than the variation of the axial position. Therefore, it can be seen that it is the maximum thickness position that dominates the position of maximum thickness within the gas path channel. For a single blade evaluation this directly impacts the position of the narrowest throat area.

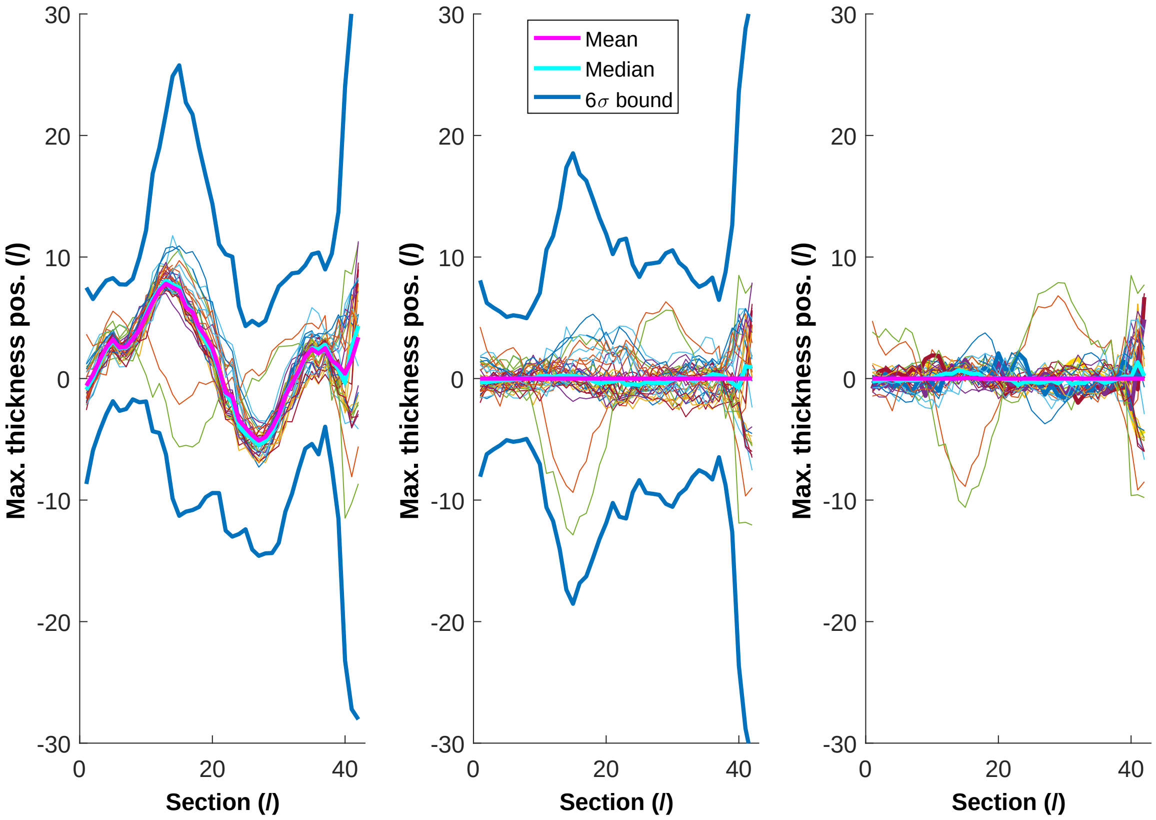

Figure 5.45: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design maximum thickness position.

The centre panel shows that the deviations change along the span. The peak at about 30% of the span is mostly due to two blades which dip down to about -10. These blades are actually closer to the design intent than the rest of the blades. Furthermore, the scatter increases significantly towards towards the tip. A comparison of the right and centre panels show that most of the scatter is due to the different blade spanwise behaviours, as opposed to simple offsets.

The scatter seen in the centre panel of figure 5.45 is either due to actual geometrical deviations, calculation procedure or its interpretation relative to the chord length. For example, it could be that the LE or the TE decrease at different rates. This could cause an instability of the position of maximum camber, even if it is located at the same coordinates, as the position is calculated relative to the chord length. The left panel shows that the differences appear to have a high signal to noise ratio, and it can be seen that calculation procedure does not produce the scatter present at the tip at any other section. There is nothing special about the highest section, as it is still located within the body of the blade. Therefore it can be concluded that the scatter at the tip is not due to the calculation method itself.

Secondly, figure 5.28 shows that the scatter of chord length is approximately constant all along the span. Therefore it cannot be responsible for the scatter seen at section 42 of the centre panel of 5.45. Therefore it is expected that this scatter is due to actual blade geometry deviations. Furthermore, the spanwise behaviour of the chord length is not responsible for the spanwise behaviour of the position of the maximum thickness.

Position of max. camber

Similarly to the position of maximum thickness, the position of maximum camber helps compare the camber distributions of multiple blades. The impact of manufacturing variability on position of maximum camber can be seen in figure 5.46. The left panel of this plot shows that the differences follow a wavy behaviour, with the scatter increasing dramatically in the last third of the span. Comparing the right and centre panels it can be seen that most of the scatter is due to different blade spanwise behaviours, as the panels show very similar plots. Similarly to the discussion of 5.45, this is attributed to actual geometry changes.

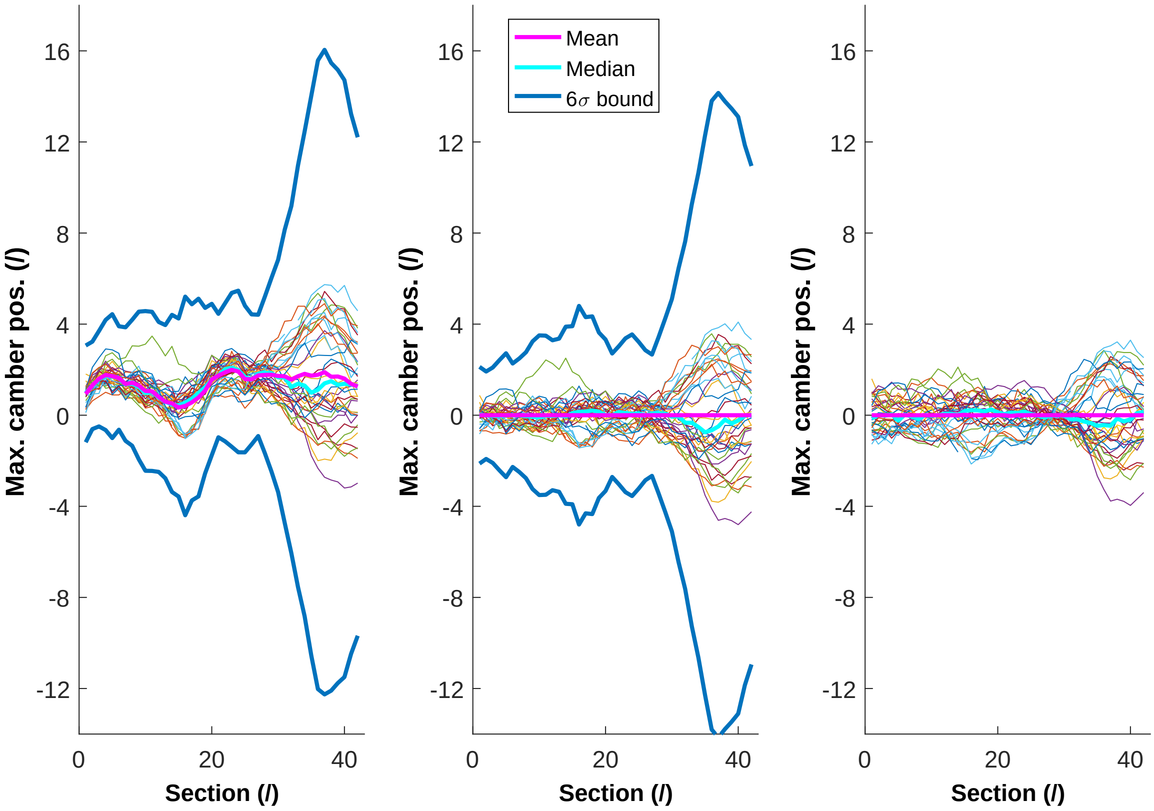

Figure 5.46: Absolute differences (left) and isolated difference profiles (right) between the manufactured and the nominal CAD design maximum camber position.

5.6.2. Geometrical deviations due to deterioration

The aim of this subsection is to quantitatively describe the deterioration scatter of the scanned blades. The analysis of the impact of deterioration is based on the discussion of the impact of manufacturing from subsection 5.6.1. This is done to maintain a perspective between the two sources of geometrical deviations investigated.

The presentation of the results also follows the format established in subsection 5.6.1. However, in this case the differences of geometrical parameters are calculated between the associated deteriorated and manufactured blade parameters. The deteriorated blades were aligned to the nominal CAD design in order to show the full impact of deterioration on the gas path channel.

Axial position

The impact of deterioration on the axial position of the blade LE can be seen in figure 5.47. The right panel of this figure has been removed as it did not contribute new insights. The left panel shows that the changes near the hub of the blade are very small, as all the lines originate at about 0. Along the span the differences rise, which means that the axial positions of the deteriorated blades have moved towards the rear of the engine, and are now behind the manufactured ones. Note that there is line with decreasing values along the span. The decrease of axial position and therefore an offset in front of the manufactured blade is a consequence of the alignment of the blade, as the blade is not expected to increase in size. It also suggests that the differences of other blades could be caused due to the deterioration of the dovetail, as all of the differences follow a linear behaviour.

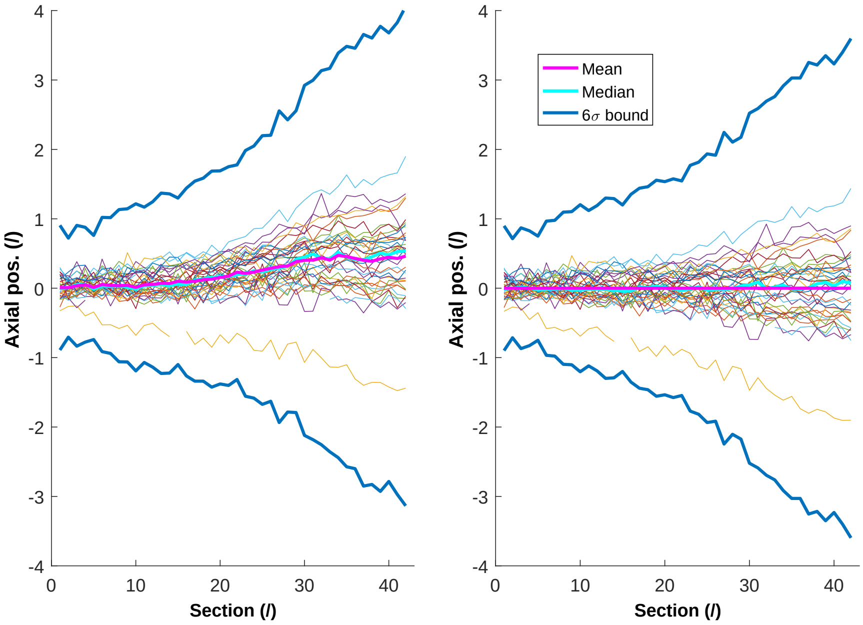

Figure 5.47: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured axial positions.

This is further supported by the right panel, which shows that the blade-to-blade scatter increases roughly linearly with the span. The blade spanwise behaviours in the right panel are representative of the behaviour in the left panel, as the subtracted mean line is also practically linear. The behaviours of the lines all appear to be linear in this panel, which further supports the idea that the differences in axial position occur due to dovetail deterioration.

The blade-to-blade scatter in this panel, at sections towards the tip, is approximately ±1. Comparing this scatter to the one seen in the centre panel of figure 5.25 it can be seen that for the investigated set of blades the scatter due to manufacturing is about twice the scatter due to deterioration.

Tangential position

The impact of deterioration on the tangential position of the blade LE can be seen in figure 5.48. This figure contains only one panel that corresponds to the left panel of the parameter difference figured. The figure shows two separate groups. As the groups of different behaviours can already be seen in the left panel, the centre and right panels do not provide additional insight and have therefore been omitted. Furthermore, the mean, median and the 6σ bounds are misleading in this case, therefore they were omitted as well.

Both groups originate just below 0. The first group increases somewhat along the span, with its scatter also increasing. The behaviour of the first group is linear and could therefore have been caused by the deterioration of the dovetail.

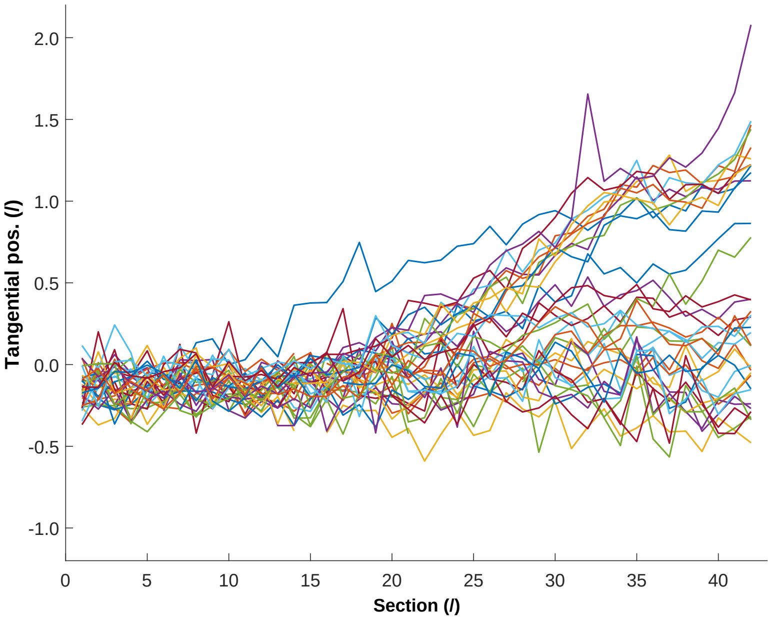

Figure 5.48: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured tangential positions.

The second group increases dramatically in the last half of the span. The increase in tangential positions means that these blades have their LE offset in the direction of the suction side relative to the the nominal CAD design. Furthermore, it can be seen that the behaviour of the second group is not linear, therefore it certainly cannot be just an effect of the dovetail deterioration. This suggests that the blades were bent somewhat towards the suction side.

Lastly, comparing the scatter due to deterioration, seen in figure 5.48, and the scatter due to manufacturing, seen in figure 5.26, it can be seen that manufacturing has a larger impact on the scatter of tangential positions of the LE for the investigated set of blades. The difference due to the behaviour of the second group is roughly comparable to the spanwise behaviour due to manufacturing.

Stagger angle

The impact of deterioration on the stagger angle of the blades can be seen in figure 5.49. The first thing that can be seen is that the differences of the stagger angle due to deterioration are noisy. Furthermore, from the left panel of this figure it can be seen that the stagger angle decreases from the hub, where it is slightly positive, towards the tip, where it is slightly negative. As the stagger angle increases from the hub to the tip by design, this means that the difference between the stagger angles at the hub and the tip has decreased. Therefore the blades have been untwisted somewhat as a consequence of deterioration.

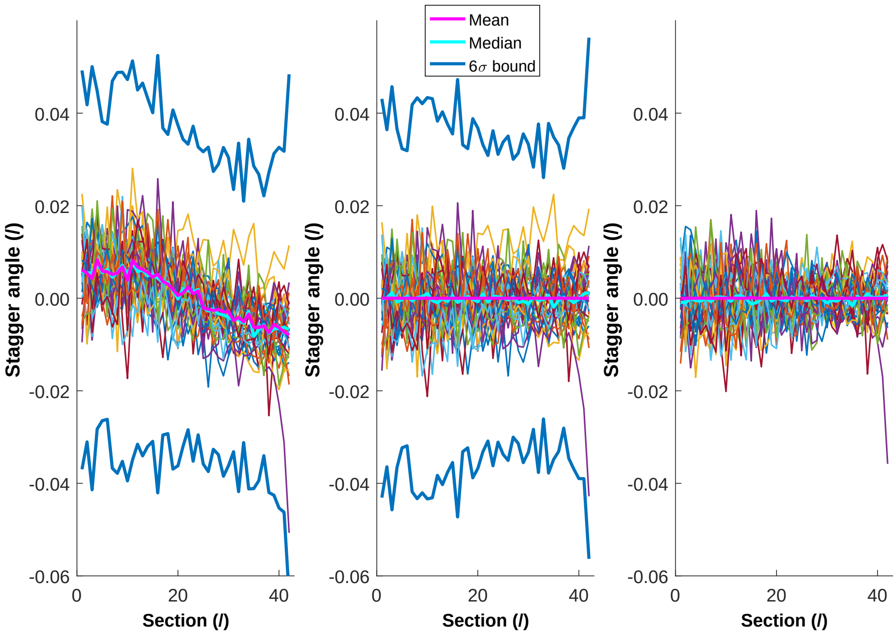

Figure 5.49: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured stagger angle.

The centre panel shows the blade-to-blade differences in the stagger angle behaviour. It can be seen that the noise dominates the plot and no behaviour can be discerned directly. The right panel shows that virtually all of the scatter is due to the noise of the stagger angle as the plots in the centre and right panel are virtually identical. Therefore any blade-to-blade differences are likely centred at 0 and concealed with the noise. However, it can be seen that the blade-to-blade scatter of stagger angle reduces slightly from the hub towards the tip.

The difference between the mean stagger angle differences at the hub and the tip is smaller than 0.01. The right panel shows that the scatter due to the noise is roughly ±0.014, which is about 1.5 times larger than the change due to the mean behaviour. Therefore, the accuracy of the observed impact of deterioration on stagger angle is questionable.

Lastly, by comparing figures 5.27 and 5.49 it can be seen that manufacturing has a larger impact on the uncertainty of the stagger angle. The behaviour seen in the left panel of figure 5.27 shows maximum mean differences of about -0.1, whereas for the case of deterioration the largest differences are less than 0.01. This means that the impact of manufacturing on stagger angle is an order of magnitude larger than the impact of deterioration.

Thickness distributions

The impact of deterioration on the blade thickness distributions along the span can be seen on the left side of figure 5.50. The far left panel shows the impact of manufacturing on the thickness distributions, as seen on the left panel of figure 5.30. The left panel shows the differences between the measured deteriorated and manufactured thickness distributions.

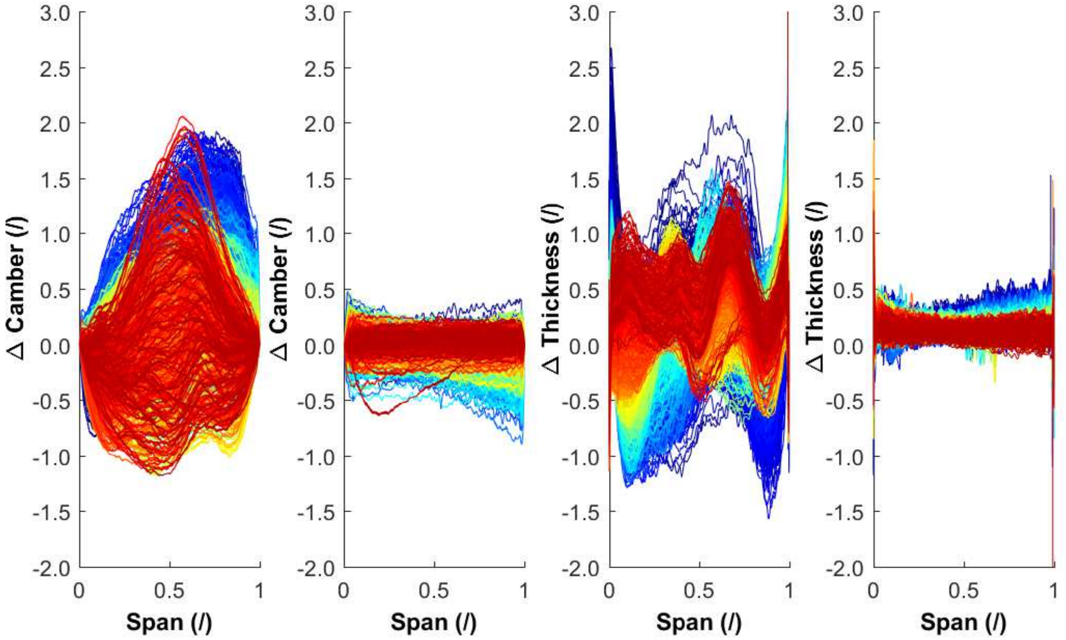

Figure 5.50: The comparison of impact of manufacturing (far left) and deterioration (left) on all thickness distributions, and the comparison of impact of manufacturing (right) and deterioration (far right) on all camber distributions.

It can be seen that the impact of deterioration is relatively small compared to the impact of manufacturing. However, it can be seen that the differences due to manufacturing general lie above 0, which indicated that the deteriorated blades are somewhat thicker. Furthermore, it can be seen that the blades are becoming thicker specifically in the sections near the TE of the hub. This is indicated by the blue lines climbing along the span. The increase in thickness is less than 0.5.

From this it can be seen that the blades are generally thicker near the hub TE due to deterioration, however these differences are small relative to the differences incurred by the manufacturing process. Deterioration therefore does not have a large effect on the thickness of the considered blade set.

Camber distributions

The impact of deterioration on the blade camber distributions along the span can be seen on the right side of figure 5.50. The right panel shows the impact of manufacturing on the camber distributions, as seen on the left panel of figure 5.32. The far right panel shows the differences between the measured deteriorated and manufactured camber distributions.

Similarly to the case of impact of deterioration on the thickness distributions, the far right panel of figure 5.50 shows that the impact of deterioration on the camber distributions is relatively small. Furthermore, it can be seen that the differences due to deterioration at the tip sections are constant along the chord, whereas at the hub sections they are more scattered. Therefore it can be concluded that deterioration impacts the shape of the sections near the hub more than the shape of the sections near the tip.

Chord length

The differences in the chord length as a result of deterioration can be seen in figure 5.51. Similarly to the case of the tangential position, 3 behavioural groups can be identified. As discussed for that case, the centre and right panels, as well as the mean, median and 6σ bounds have been omitted here as well.

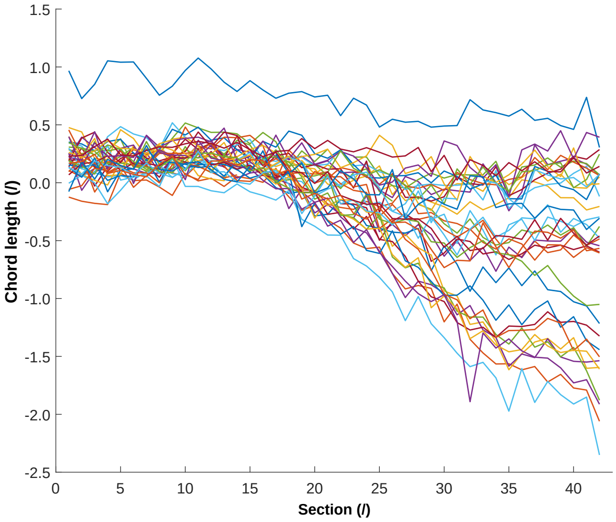

Figure 5.51: Absolute differences between the deteriorated and the manufactured chord length.

The three groups all originate between chord length difference of roughly 0 and 0.5, with one outlier originating at 1. For the first half of the span all lines stay together, and fan out into three separate groups at about mid-span. A single outlier originating at approximately 1 can be identified as well. As the outlier lies so far away from the other lines, it is likely that it is caused by the significant deterioration of the alignment surfaces.

It can be seen that the top most group remains roughly within the scatter bounds of the origin for the entire span. The mid group retains the same scatter, however falls down to about -0.5. The bottom group falls down to roughly -1.50 and has its scatter doubled.

The increase of the chord length that is observed in the first half of the span is most likely due to the deterioration of load bearing surfaces of the dovetail. As these are eroded, the blade is offset towards the tip, which in turn increases the chord length at any given radial position above the axis of rotation. The increase of the chord length of about 0.25 would mean that the blade has been offset in the radial direction by about a third of the spanwise distance between the radial sections, which translates into about 4 for a blade with the height of 600.

The top group therefore shows blades which have not had their chord lengths reduced, but have been offset in the radial direction. The mid and bottom groups show blades which were offset in the radial direction by the same amount, but have had some other effects reduce their chords by various amounts. The split into three groups could be due to the blades being operated in three different engines. As the operational data of the blades is unavailable this could not be checked.

Lastly, comparing the impacts of manufacturing and deterioration on the chord length, it can be seen that the maximum difference due to deterioration are roughly -2, and the maximum differences due to manufacturing are about -5. The scatter due to manufacturing is about 4, whereas the entire scatter due to deterioration, including the three distinct groups, is about 2.5. From this it can be seen that the impact of manufacturing on the chord length is about twice the impact of deterioration.

Maximum thickness

The differences in the maximum thickness as a result of deterioration can be seen in figure 5.52. The left panel of this figure shows that due to deterioration the blades are thicker at the end of their lives. The increase of thickness is linear with the span and is highest at the hub (0.2) and lowest at the tip (0.1). Note that these values are smaller than the resolution of the scanner.

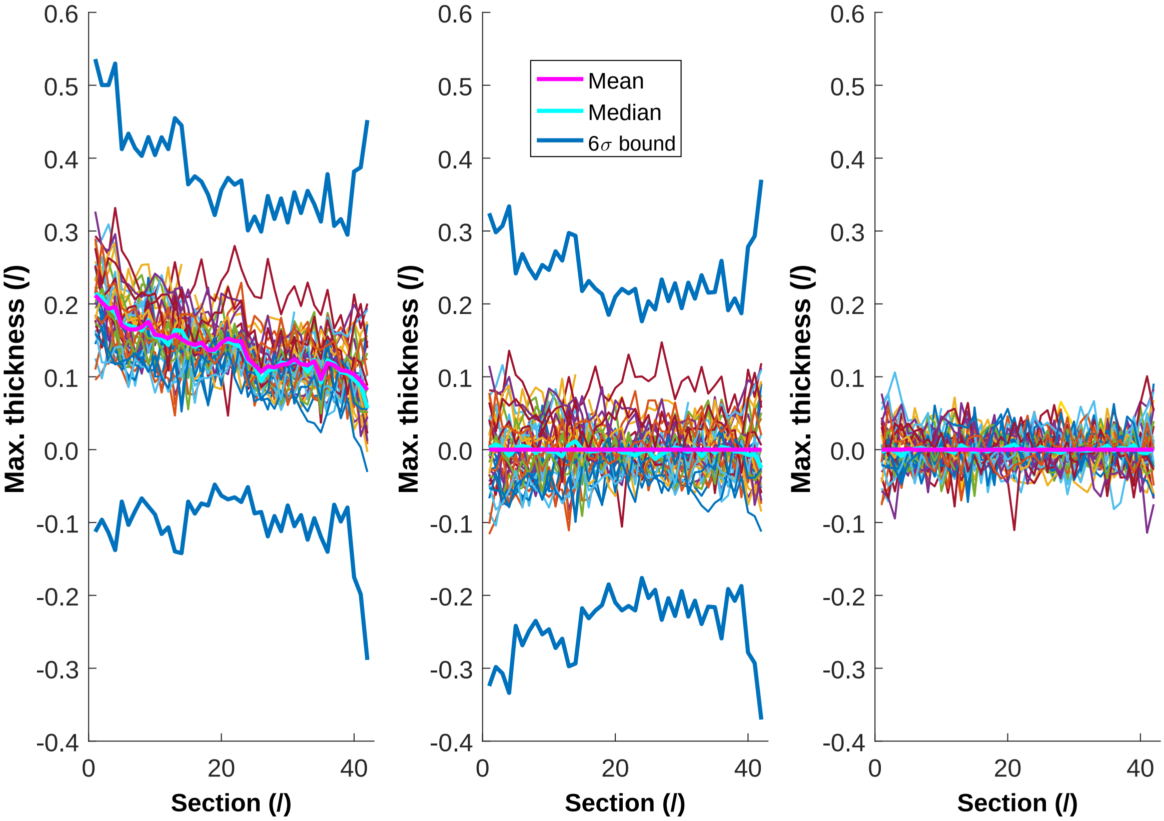

Figure 5.52: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured maximum thickness.

The left panel shows that the decrease of maximum thickness from the hub to the tip is about 0.1. The centre panel shows that the blade-to-blade differences are about twice that. Furthermore, by comparing the centre and right plots it can be seen that noise accounts for most of those differences. Therefore noise spans a larger range than the spanwise variation of the maximum thickness seen in the left panel. Therefore, the accuracy of the identified impact of deterioration on maximum thickness is questionable.

Comparing the impact of deterioration and manufacturing on maximum thickness, it can be seen that the maximum mean differences due to deterioration are about 0.2, while the mean maximum difference due to manufacturing is roughly 0.8. The scatter due to manufacturing is about 2, while the scattering due to deterioration is about 0.2. Therefore, manufacturing has an impact on maximum thickness that is roughly 4 times larger than the impact of deterioration.

Maximum camber

The differences in the maximum camber as a result of deterioration can be seen in figure 5.53. The left panel of this figure shows that the mean differences have the shape of a quadratic function, with the minimum at around section 10. This shape shows that near the hub the maximum camber has been reduced due to deterioration, whereas near the tip it has been slightly increased.

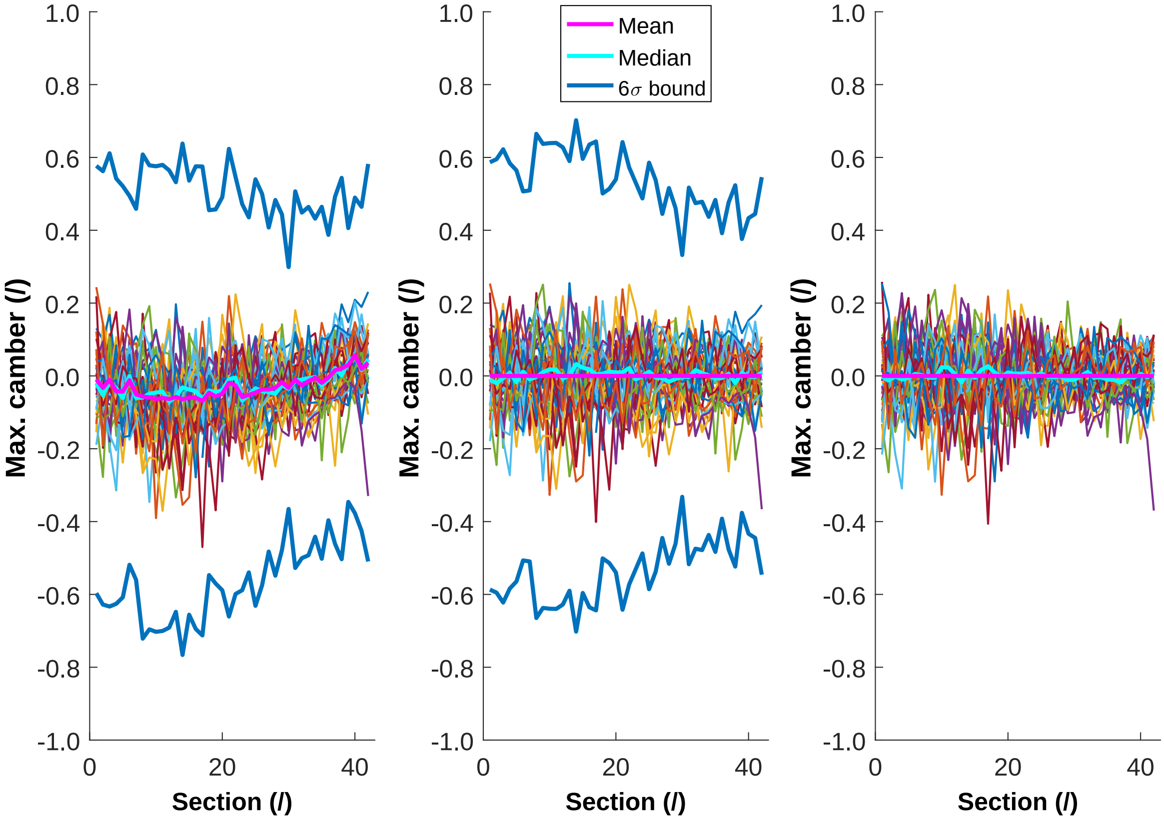

Figure 5.53: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured maximum camber.

Figure 5.52 shows that the maximum thickness of the blades has increased. The increased thickness and reduced camber together suggest that the accumulation of the thickness has therefore occured such that the camber has decreased. Specifically, the thickness of the blades has increased because of the additional geometry on the pressure side of the blades. Furthermore, the maximum camber of the blades decreases from the hub to the tip by design. This means that any offsets in the radial direction due to the abrasion of the load bearing surfaces of the dovetail would increase the maximum camber, not decrease it. This suggests that the actual reduction of camber at comparable spanwise positions is larger than the one seen in figure 5.53.

However, the calculated maximum camber differences are small and very noisy. This can be seen best by comparing the plots of the centre and right panels. This comparison shows that the scatter of the values is due to noise as opposed to some behaviour. Furthermore, the noise accounts for about ±0.20. Therefore, the accuracy of the observed impact of deterioration on maximum camber is questionable.

Lastly, comparing the impact of deterioration and manufacturing, it can be seen that the maximum mean difference due to deterioration is roughly -0.05, whereas the maximum mean difference due to manufacturing is roughly 0.8. The blade-to-blade scatter due to deterioration, including the noise, is roughly 0.4, whereas the scatter due to manufacturing is at least 1 and increases along the span. Therefore, manufacturing has a larger impact on the maximum camber of the blades than deterioration.

LE angle

Similarly to the discussions of the impact of manufacturing on LE and TE angles, the differences due to deterioration are very noisy. This can be seen in figure 5.54. The mean differences due to deterioration are centred about 0 along the entire span.

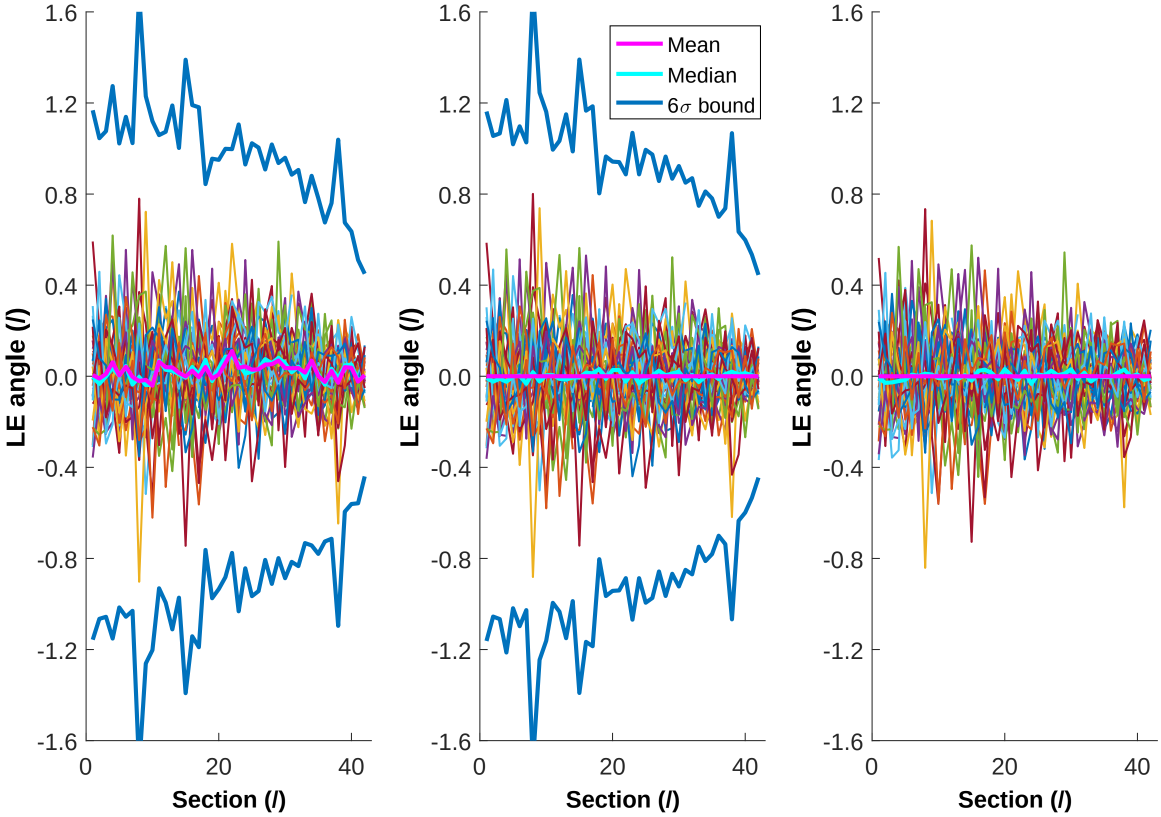

Figure 5.54: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured LE angle.

It can be seen that the scatter reduces from the hub to the tip. The comparison between the centre and right panels shows that virtually all of the scatter is due to the noise and not due to any inherent blade-to-blade behaviour differences. The same reasons for the noise as mentioned in the discussion of the impact of manufacturing on the LE and TE angles apply here as well. Lastly, as the mean behaviour appears to be linear along the entire span the plot of smoothed LE angle differences has been omitted here.

Lastly, the impact of deterioration compared to the impact of manufacturing is quite small for the considered blade set. Due to the noise it is not possible to determine any real difference of the LE angle, whereas the maximum mean difference due to manufacturing is roughly -0.2. It can therefore be concluded that manufacturing has a larger impact on the LE angle than deterioration. However, the LE angle differences are very noisy, therefore it is not possible to accurately determine by how much the impact of manufacturing is larger than the impact of deterioration.

TE angle

The impact of deterioration on the TE angle can be seen in figure 5.55. Unlike for the LE angle, the deterioration does seem to have a distinct impact on the TE angle. In the left panel of this figure it can be seen that there is a dip in the differences of the TE angle in the first half of the span. The minimum of the dip is about −0.8.

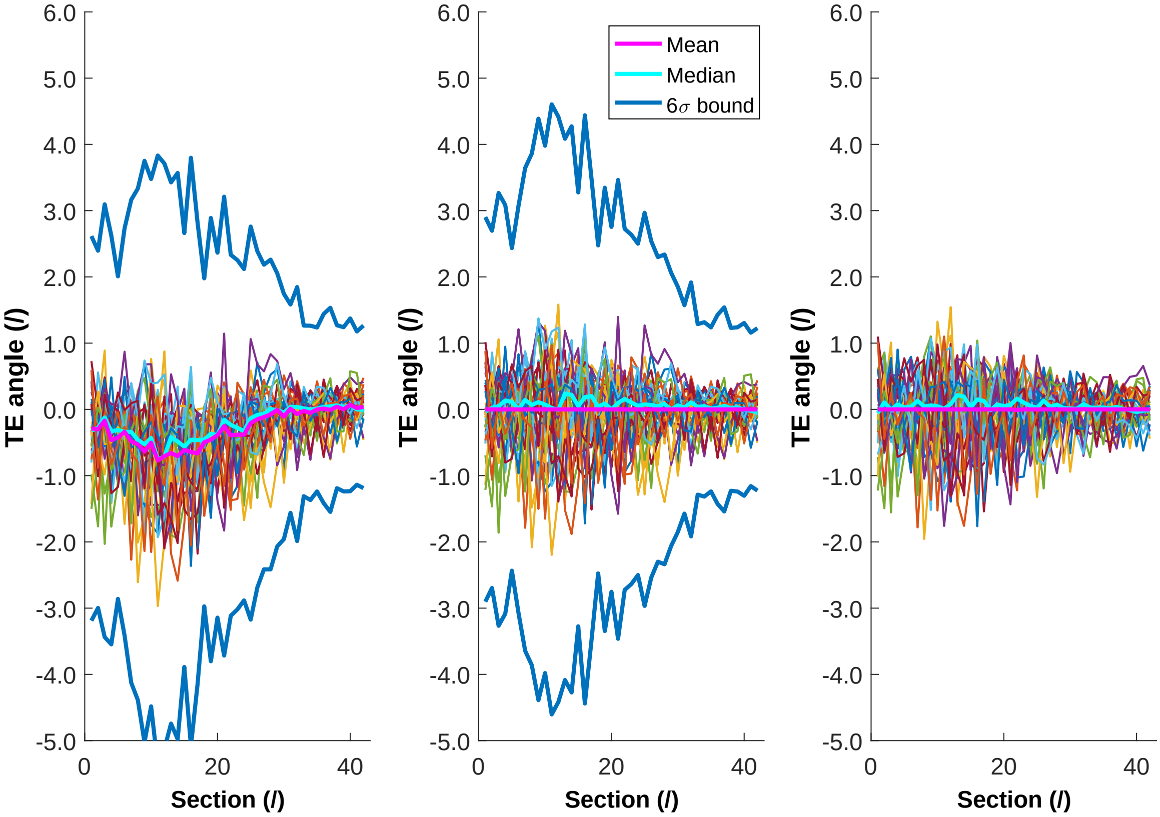

Figure 5.55: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured TE angle.

The differences are very noisy however, which can be observed by comparing the centre and right panels. This comparison also shows that the scatter of values is about twice the magnitude of the dip observed in the left panel. In order to understand the behaviour better the values smoothed with the Savitzky-Golay filter could be plotted. However, the smoothed plots showed conclusions similar as in the cases where the impact of manufacturing on LE and TE angles has been discussed. Therefore this plot and its discussion have been omitted here.

Lastly, comparing the impact of deterioration and manufacturing, it can be seen that the maximum mean difference due to deterioration is roughly −0.8, whereas the maximum mean difference due to manufacturing is roughly 1.2. The blade-to-blade scatters due to deterioration, including the noise, are also comparable. Furthermore, it can be seen that the impact of deterioration is exactly the opposite of the effect of manufacturing.

LE thickness

The differences of the LE thickness due to deterioration can be seen in figure 5.56. The left panel of this plot shows that for the first half of the span the differences of LE thickness are centred about $0 \mu m$. In the second half of the span the LE thickness differences show that the thickness of the LE has increased due to deterioration.

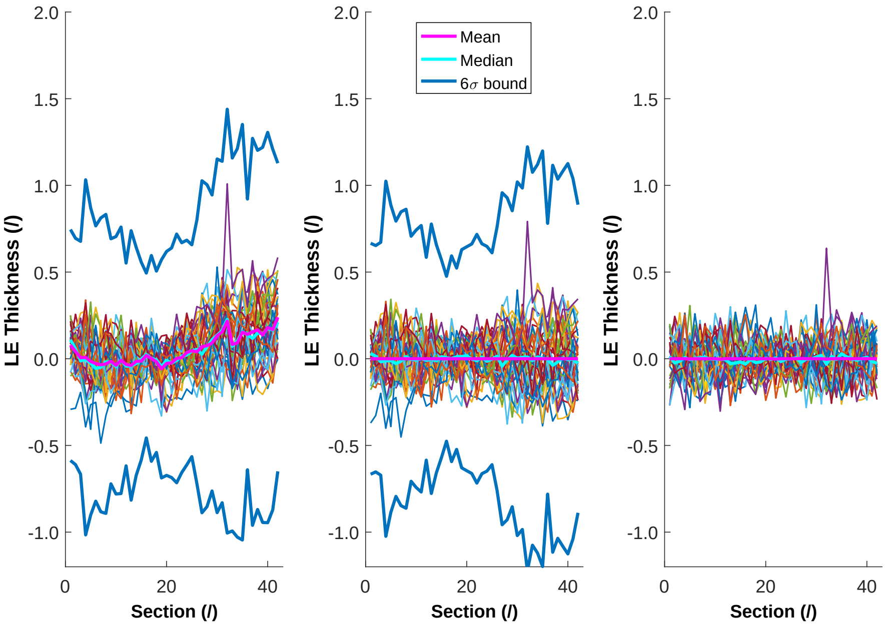

Figure 5.56: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured LE thickness.

The reason for this could be the offset in the radial direction due to the abrasion of the load bearing dovetail surfaces, as the lower sections of the manufactured blades had thicker LE than the upper sections. However, in the expected offset of 1/3 of the spanwise sections the increase of the LE thickness would be less than 0.1. Furthermore, this increase would depend on the spanwise behaviour of the LE thickness, specifically on the derivative of the behaviour with the span. Figure 5.41 shows that the decrease of the LE thickness with the span is highest for the sections near the hub. This means that if the radial offset would be the cause of the LE thickness increase due to deterioration, the increase would be largest at the hub sections. The left panel of figure 5.56 shows that this is not so, therefore other effects must have caused the increased LE thickness. These could be due to fouling and dulling of the LE.

The scatter of the LE thickness values is quite large. This can be seen best in the centre and right panels of figure 5.56. The centre panel shows that the blade-to-blade scatter first reduces from the hub towards the mid-span, before it increases towards the tip, where it is larger than at the hub. The comparison with the right panel shows that most of this scatter is due to noise, however a smaller part could be due to differences in the blade-to-blade behaviours. Notice also that the range between the LE thickness differences shown in the left panel at the hub and the tip is less than approximately 0.25, whereas the scatter easily has a range of 0.5. This casts doubts on the accuracy of the mean behaviour shown in the left panel.

Comparing the impact of manufacturing and deterioration it can be seen, that the impact of deterioration on the LE thickness is smaller. The maximum mean differences due to deterioration are less than 0.25, whereas the maximum mean differences due to manufacturing are about 0.25. Similarly, the scatter due to manufacturing is more than 2, whereas the scatter due to deterioration, including the noise, is about 0.5.

TE thickness

The differences of the TE thickness due to deterioration can be seen in figure 5.57. Similarly to the differences of the LE thickness due to deterioration, the differences of the TE thickness are very noisy. In the left panel of this figure it can be seen that the differences are positive along the entire span. Furthermore it can be seen that the differences are somewhat higher in the first quarter of the span near the hub than in the other spanwise positions.

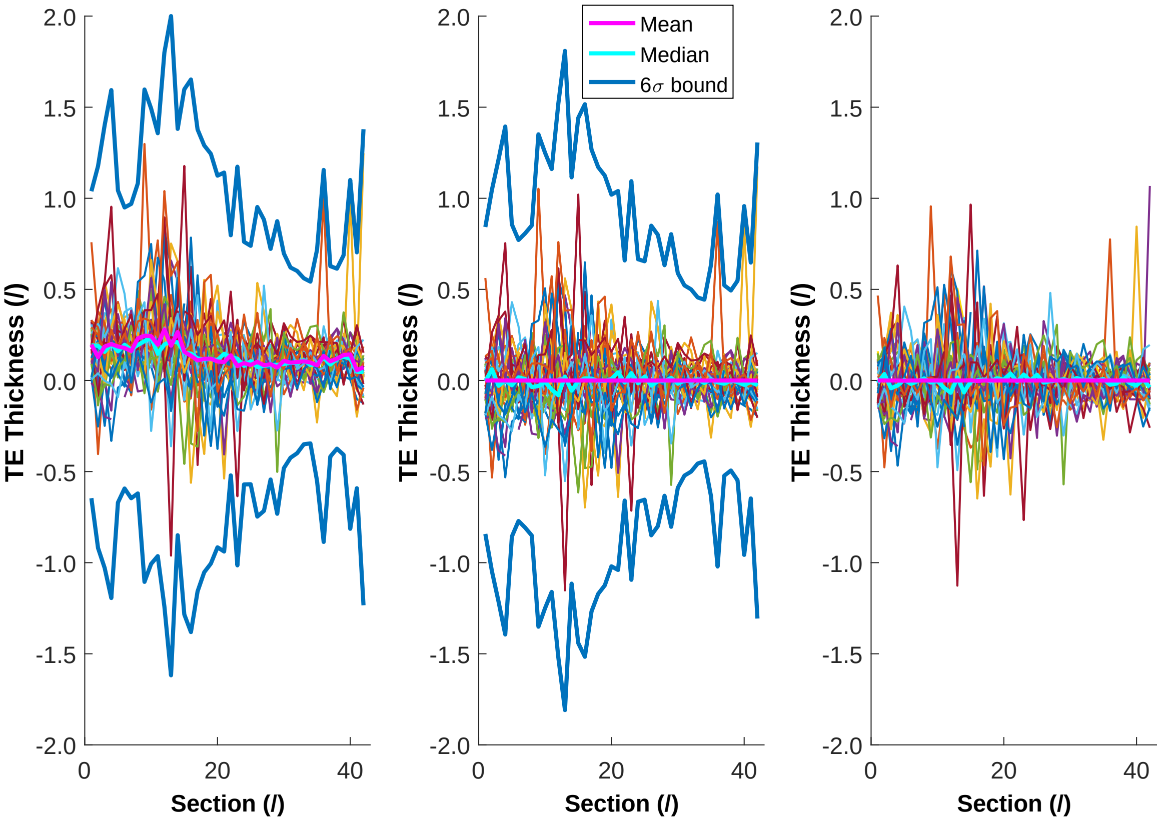

Figure 5.57: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured TE thickness.

Like in the case of the impact of deterioration on the LE thickness, the impact of deterioration on the TE thickness is not due to a radial offset of the blade. This can be seen by applying the discussion given for the case of the impact of deterioration on the LE thickness to the TE thickness. Similarly to that case, the possible explanations for this behaviour are fouling or dulling of the TE as a consequence of the reduction of the chord length. Lastly, the noise of the calculated values is much higher than the magnitude of the mean behaviour seen in the left panel.

Lastly, comparing the impacts of manufacturing and deterioration on the TE thickness it can be seen, that the maximum mean differences due to deterioration are roughly 0.25, whereas the maximum mean differences due to manufacturing are roughly 1.2. It can be seen that the blade-to-blade scatter is also larger in the case of manufacturing.

LE position

The differences of the LE position due to deterioration can be seen in figure 5.58. The left panel of this figure shows that the LE position moves backwards all along the span due to deterioration. The shape of the mean difference behaviour appears to be quadratic, with a minimum around section 30. The lowest mean difference is not more than -0.2 normalised percent of the chord length. For a blade with a chord length of approximately $40mm$ this translates to roughly equal to the scanner resolution.

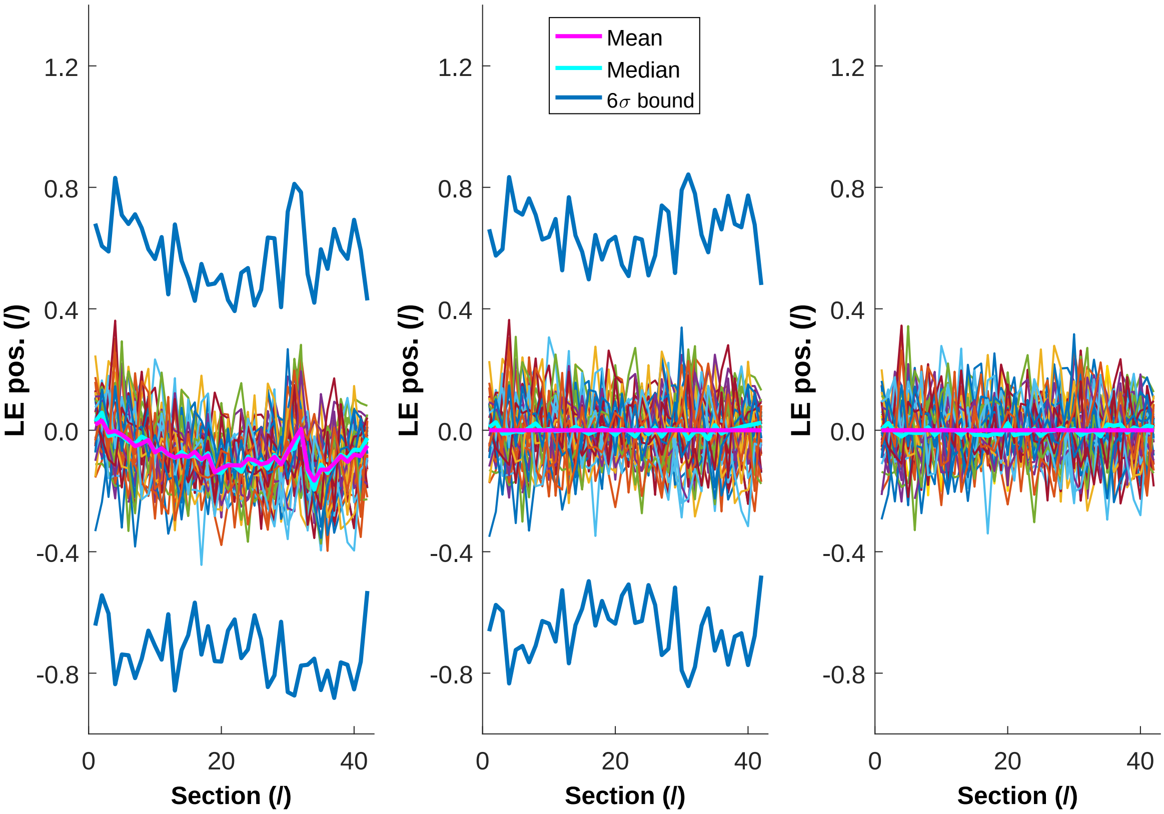

Figure 5.58: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured LE position.

The scatter appears to be relatively constant along the span. This can be seen best in the centre and right plots. The comparison of these panels shows that virtually all blade-to-blade differences are due to noise. Due to the small magnitude of the mean deviations and the high amount of noise the accuracy of the mean behaviour shown in the left panel is questionable.

The maximum mean difference of the LE position is smaller than 0.2 normalised percent of the chord length, whereas the maximum mean difference due to manufacturing is just under 1 normalised percent of the chord length. The blade-to-blade scatters are comparable, and mostly due to noise. It can therefore be seen that the impact of manufacturing on the position of the LE is larger than the impact of deterioration. This is amplified by the fact that the deteriorated chord lengths are smaller than the manufactured chord lengths.

TE position

Similarly to the impact of deterioration on the LE position, the differences of the TE position due to deterioration are noisy. Furthermore, the mean differences behaviour is centred around 0 normalised percent of the chord length, with magnitudes smaller than 0.1 normalised percent of the chord length. This translates into a movement of the position of the LE which is roughly equal to the blade measurement resolution. Therefore, the accuracy of the identified behaviour is questionable. This is even more so in this case, than in the case of the differences of the LE positions due to deterioration.

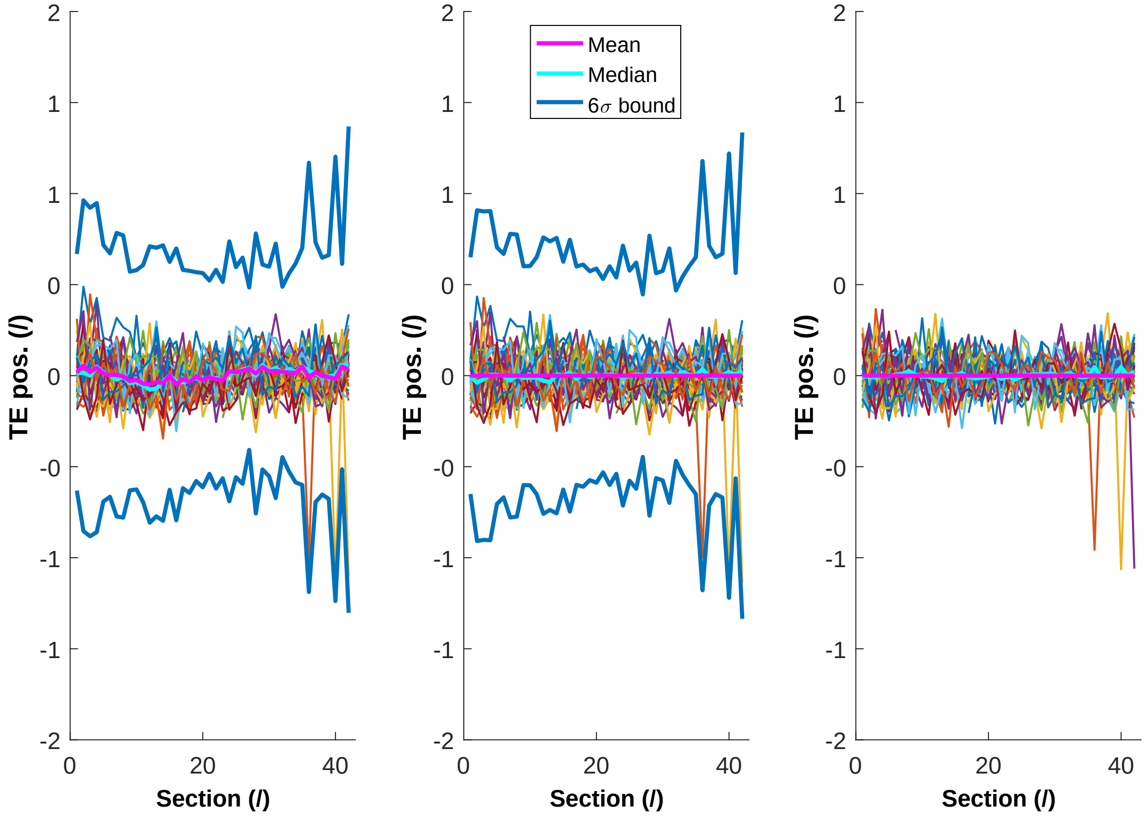

Figure 5.59: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured TE position.

The maximum mean difference of the TE position due to manufacturing, as seen in figure 5.44, is roughly −0.4 normalise percent of the chord length. The impact of deterioration on the TE position much smaller than that along the entire span. Therefore, it can be concluded that manufacturing has a larger impact on the TE position.

Position of max. thickness

The impact of deterioration on the position of the maximum thickness can be seen in figure 5.60. The left panel shows that for the majority of the span the differences of the position of maximum thickness due to deterioration are roughly ±1 normalised percent of the chord length. The differences are much larger for the last few sections near the tip. However, the mean behaviour can be contained by a much smaller interval. Furthermore, the mean behaviour appears to be centred about 0 normalised percent chord length with any deviations from it appearing random.

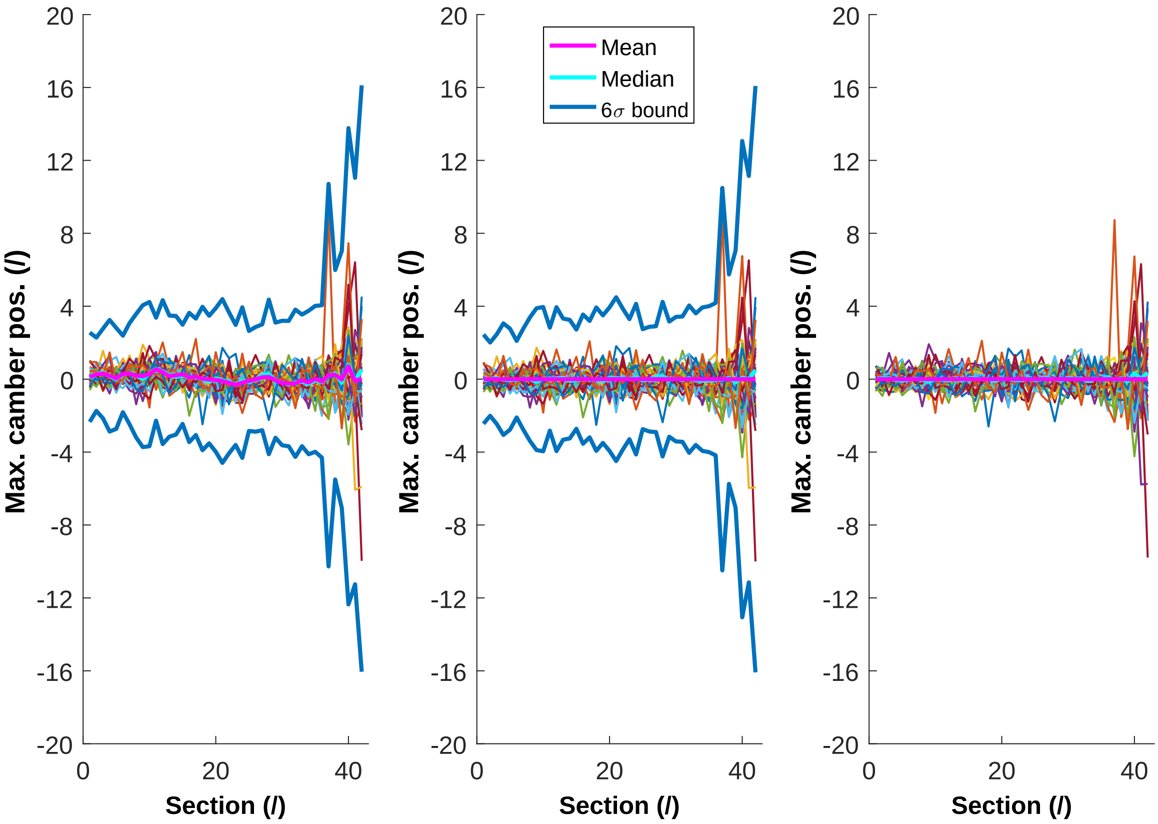

Figure 5.60: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured position of maximum thickness.

The scatter can be seen to be increasing along the span, even if the final few sections where the differences spike is not considered. This can be seen in the centre panel. Furthermore, by comparing the centre and right panels it can be seen that the noise dominates the blade-to-blade differences.

Lastly, comparing the impacts of manufacturing and deterioration on the position of maximum thickness, it can be seen that the maximum mean differences due to manufacturing are roughly 8 normalised percent of the chord length, whereas the maximum mean differences due to deterioration are roughly 0 normalised percent of the chord length for the entire span. Therefore, it can be concluded that manufacturing has a larger impact on the position of the maximum thickness.

Position of max. camber

The impact of deterioration on the position of the maximum camber can be seen in figure 5.61. The left panel of this figure shows that the difference of the max camber position due to deterioration are negative for most of the span. This means that the position of maximum camber has moved backwards, towards the TE.

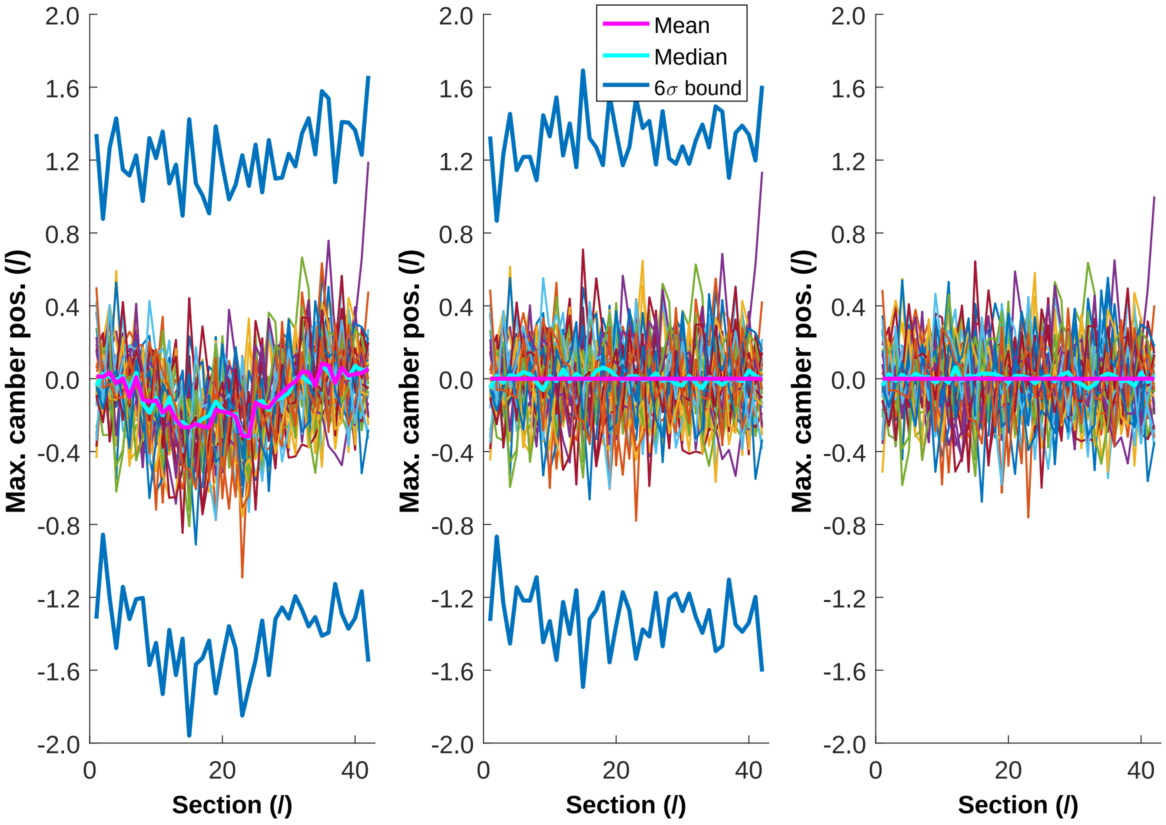

Figure 5.61: Absolute differences (left) and isolated difference profiles (right) between the deteriorated and the manufactured position of maximum camber.

The maximum shift due to deterioration appears at about mid-span, where the difference of the maximum mean camber position is less than −0.4 normalised percent of the chord length, which translates into less than −0.8.

The centre panel shows that the blade-to-blade scatter remains roughly constant along the span. The range of the blade-to-blade scatter is roughly ±0.8, which means that the noise is about twice the magnitude of the largest mean difference. The right panel further shows that the blade-to-blade scatter is mostly just noise.

Comparing the impacts of manufacturing and deterioration on the position of maximum camber, it can be seen that the impact of manufacturing is larger. The maximum mean differences due to manufacturing are about 2 normalised percent of the chord length, whereas the maximum mean differences due to deterioration are about −0.2 normalised percent of the chord length.

5.6.3. Deteriorated blade height analysis

The tip gaps could not be calculated for both the manufactured state and the deteriorated state due to the lack of information as discussed in subsection 5.2.3.

However, the expectation that the blade heights will not change significantly with regards to deterioration can still be evaluated by comparing the deteriorated blade heights to the manufacturing tolerances. The offset of the mean blade height from the design value would be due to deterioration, assuming that the blade heights were normally distributed about the design value when entering operational use.

The blade heights have been calculated using the B2P functionality presented in the section 5.2.2. This calculation outputs the tip contour points in 3D Cartesian coordinates. For plotting purposes these are converted into 2D coordinates using equation 5.3.

$$f(x, y, z) = (x, \sqrt{y^2 + z^2}) = (x, r) \tag{5.3}$$

As the scanned blades were aligned to the nominal blade, the combined coordinate "r" is the height of individual points above the axis of revolution of the compressor.

The heights of the deteriorated and the nominal blade are presented in figure 5.62 along with the blade height tolerance bounds. Note that the nominal blade design features extra blade height to accommodate final machining during assembly. Therefore the nominal geometry has been manually reduced according to the assembly instructions for comparison with the deteriorated blades.

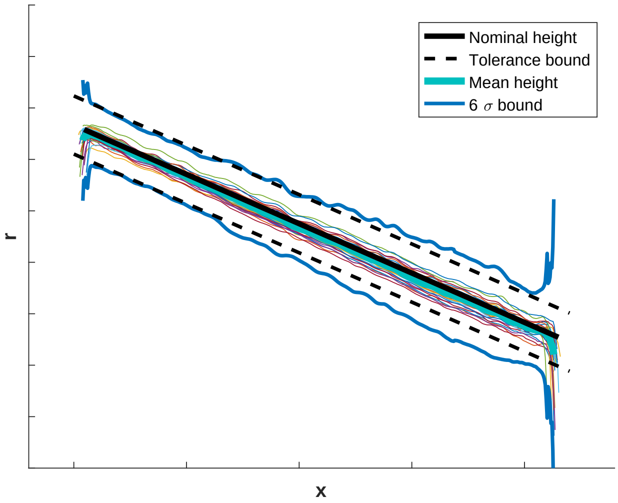

Figure 5.62: Tip heighs of deteriorated BR715 R3 blades in radial coordinates.

To assess the offset of the group from the nominal blade a mean contour has also been calculated. This has been calculated as the mean of the contours interpolated at the same positions. Based on these interpolated values the standard deviation at all positions has also been calculated to plot the 6σ bounds. Note that if the median blade contour is added it lies directly on top of the mean blade. This means that the distribution of the tip contour values at none of the positions is skewed.

In figure 5.62 the deteriorated blade heights are represented by coloured thin lines, and the mean, 6σ bound, nominal height and tolerance bounds as thick cyan, blue, black and dashed black lines respectively.

Near the LE and the TE the contours drop suddenly. In the LE both the scanned and the nominal blade contours drop together – this is where the leading edge truncation begins. At the trailing edge only the scanned blades drop. This suggests that at least near the blade tip the blade chord of actual blades is shorter than the nominal chord.

Between these LE and TE regions all the blade tip contours fall well within the tolerances. Furthermore also the nominal height and the mean height lie very close together and partially overlap. It can be seen that near the LE they are closest and they are more apart nearer the TE. It can also be seen that near the LE the scatter between individual blades is smaller than near the TE. The 6σ bounds mostly exceed the manufacturing tolerances, more prominently towards the TE. However as the tip machining is not enforced long the entire contour it cannot be claimed that the breach of the 6σ bounds are due to deterioration.

In order to better visualise this behaviour the nominal contour can be subtracted from the deteriorated contours to show only the variation about it. This is done in figure 5.63. This figure uses the same format as 5.62, with the addition of a new y-axis. This axis contains the p-values of the hypothesis that the differences of individual blades, at a certain position, come from a distribution with its mean at 0. The p-values are plotted by a thin green line.

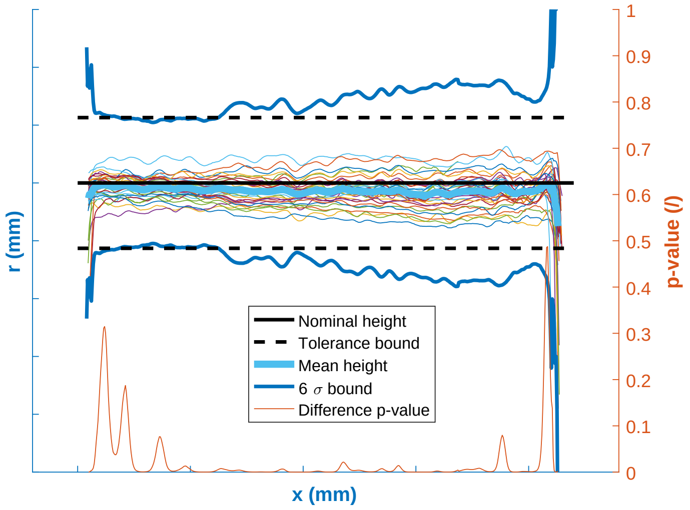

Figure 5.63: Tip height difference between deteriorated and nominal BR715 R3 blades in radial coordinates.

In figure 5.63 it can be seen that the scatter almost doubles from the LE to the TE if we ignore the highest contour near the LE as an outlier. Furthermore it can be seen that the mean height difference is consistently negative.

The p-values along the x-axis show the probability of extracting the blade differences of that position from a distribution centred about 0. It can be seen that the p-values are high near the LE and near the TE where the mean and the nominal line are close together. In the mid-chord region they are well below the 2% level. These values suggest that the mean offset is statistically significant at least at the 5% level.

However, the mean offset is only about 2 times larger than the resolution of the scanner that measured the blades. Furthermore it is important to note that compared to the estimated tip gap in cruise conditions this mean offset is insignificant.

It can also be seen that the blade contours roughly follow a slope from the LE to the TE which can be either positive or negative. Correlations of the slopes with the heights at the LE, mid chord and TE were calculated for both the manufactured and deteriorated tip heights. As no significant correlations have been found their discussion is omitted here. Furthermore as the average offsets are insignificant it is unlikely that significant changes in the slope would be caused by deterioration.

A separation of the observed behaviours into the contributions from the axial root and aerofoil were not performed. For the case of manufacturing the separation could not be done as the blades cannot be aligned such that these effects could be isolated. For the case of deteriorated blades the contributions were not separated as the overall effects were too small.

To summarise, possible reduction of the blade heights due to deterioration has been identified. This offset appeared to be statistically significant, but insignificant compared to the expected values of the entire tip gap.

The manufactured blades are expected to have less scattered tip heights than deteriorated blades. Therefore both the manufactured tip gaps can be modelled as normal PDFs about the design intent value. The $6\sigma bounds$ based on the differences between the deteriorated and nominal tip heights fall approximately on top of the tolerance bounds near the LE, but are relatively far away from these bounds near the TE. As it is was determined to be unlikely that deterioration would change the slopes of the tip contours this behaviour can be used when simulating the manufacturing tip gaps.

5.7. Lessons learned in comparison of deterioration and manufacturing and recommendations

Most of the geometrical uncertainty for the considered measured blades is due to manufacturing as opposed to deterioration. Furthermore, many of the parameter differences due to manufacturing appeared to have linear behaviours along the span. In the case of linear and constant behaviours, the imperfection is likely due to the positioning of the blade on the platform. These are the deviations that could potentially be reduced by adjusting the machining settings of the dovetail. In the case of piecewise linear, or other non-linear behaviours, only a part of the deviations could be recovered by merely adjusting the dovetail.

Many of the parameters have shown a large amount of noise in their values, which reduces the accuracy of the analysis of the geometrical uncertainty. In order to improve the quality of similar analyses more robust methods for extracting the parameters could be developed.

Deterioration has been found to have some spanwise effect, however for most of parameters the impact of this has been very small. The most notable changes have been observed for the positioning parameters and the chord length, which exhibited three different deterioration groups. Furthermore, the stagger angle showed a small untwist of the blades.

The analysis of the deteriorated blade heights showed that the heights are still within the manufacturing tolerances after being used in service. Therefore, it can be concluded that these blades did not have their heights reduced. The tip gap clearance behaviour over time could not be determined as information on them was not available. Future work could include scanning the liner, disk and casing geometries in order to support the determination of the tip gaps of the engine.

6

Performance analysis of manufacturing variability and deterioration

This chapter aims to demonstrate what the impact of manufacturing and deterioration uncertainty on compressor performance is. Section 6.2 discusses the settings of the Hydra 3D solver that were used to obtain performance results.

The preprocessing of geometry for use with Hydra, which is done in subsection 6.1. This subsection covers the creation of the mesh used within Hydra from the measured point cloud measurements. As some of the preprocessing steps cause a loss of some of the measured geometry, the convergence of the mesh to the original measured blades is discussed.

Furthermore, the preprocessing of geometry, selection of the boundary conditions, as well as the limitations of and shortcomings of the CFD evaluation are discussed.

Section 6.3 presents the results of the CFD calculations as a comparison between the manufactured, deteriorated and nominal CAD design blade performances. Furthermore, a comparison of the calculation convergences is given.

Lastly, the influence of individual geometrical parameters on the performance scatter is given in section 6.4. This section helps identify the geometrical parameters that cause most of the performance scatter and offset using Coefficients of Influence. The coefficients also give insight on why the performance scatter and offset occur.

6.1. Preprocessing of measured blade geometry

Subsection 6.1.1 gives a general overview of the problem geometry selection for the CFD calculations. It describes the number of blades used in a single evaluation and how that geometry is prescribed within Hydra.

Subsection 6.1.2 discusses the preprocessing of the blades before the meshing takes place.

Subsection 6.1.3 discusses the convergence of the preprocessed blades to the original measured blades. This is done by comparing the rebuilt and measured geometries, and by comparing the actual mesh to the measured geometry.

6.1.1. Formulation of the CFD problem geometry

Within the CFD the evaluation of single blades is performed. The main advantage of the measured dataset is that it enables an evaluation of the impact of manufacturing and deterioration on individual blades. Using different blade geometries to construct the gas path passages introduces effects cross blade effects. These might not be significant, however to isolate the quantification of deterioration only individual blades are evaluated in Hydra. Note that this means that simple blade-to-blade offsets of axial and tangential positions have no impact on performance.

Additionally, the tip gap must be provided as an input to Hydra separately from the blade geometry. The blade geometry input to Hydra must be extended to span the entire annulus, and the tip gap is given as a percentage of this span. As the tip gap analysis performed in subsection 5.6.3 showed that the tip heights do not seem to deteriorate, the blade heights are kept constant in the performance evaluation.

6.1.2. Blade geometry preprocessing

Geometry preprocessing is required as the measured geometry is stored in a different format than expected by Hydra. As mentioned in section 4.1, the measured blades are stored as data point clouds, whereas the Hydra preprocessing expects a NX part file geometry input. Therefore, the measured blade needs to be converted to the NX part format.

This can be done by first rebuilding the blade based on its parameters and distributions using B2P, and then using Parablading to create a NX part file based on the output files of the B2P rebuild process. As the distributions and parameters of the measured blade are sued as a basis to rebuild the blades the delta parameters are all zero. Furthermore, this means that there is no spanwise parameter flushing or distribution morphing. The parametrisation of blades rebuilt this way therefore results in the same parameters and distributions as those used in the rebuild.

However, as mentioned already in section 5.2.1, this also means that any information regarding the blades between the spanwise sections, for which the parameters are provided, is lost. The need to use rebuilt blade geometries in the CFD carries the information lost due to spanwise sampling over to the performance evaluation as well.

6.1.3. Convergence of the preprocessed blades to the actual measured blade scans

This subsection discusses the results of the blade geometry rebuilt as a part of the Hydra preprocessing. The convergence of the rebuilt geometry to the measured geometry is important as it influences the subsequent performance analysis. Furthermore, if the convergence is poor the subsequent performance results are not representative of the measured blades.

The convergence of the geometrical parameters and distributions has been ensured by using their parameters and distributions as the Base Line Model (BLM), and using all zero delta parameters. Therefore the distributions and parameters used in the rebuild should be exactly the same as those obtained when parametrising the blades. The assessment of the convergence of the rebuilt is therefore performed by inspecting the surface deviations of the rebuilt blades relative to the measured blades. The surface deviation plots were made by subtracting the surface deviations of the measured blades from the surface deviations of the rebuilt blades.

Convergence of the rebuilt measured manufactured blades

Figure 6.1 shows the mean surface deviation differences that occur due to the rebuilding process of the manufactured blades. Note that the color range of this figure is the same as the color range of figure 5.14, which shows the mean surface deviations due to manufacturing.

Figure 6.1 shows that the majority of the surface deviation differences is well within $\pm 50\mu m$. Note that this approximately equals the measurement resolution of $30\mu m$. The hub and the tip are regions of higher deviation differences, which occurs as no geometry has been sampled in those regions. The trailing edge also shows some mean deviations that are roughly $-50\mu m$ in magnitude.

The deviations of the TE occur as the TE and LE geometries are rebuilt using a predetermined shape as opposed to rebuilding the real geometry. This causes the rebuilt blades to be slightly shorter than the measured blades, depending on the local LE and TE geometries.

Furthermore, comparing the deviations seen in figure 6.1 with the deviations seen in figure 5.14, it can be seen that the geometrical deviations caused by the rebuild of the manufactured blades are much smaller than those caused by the manufacturing process.

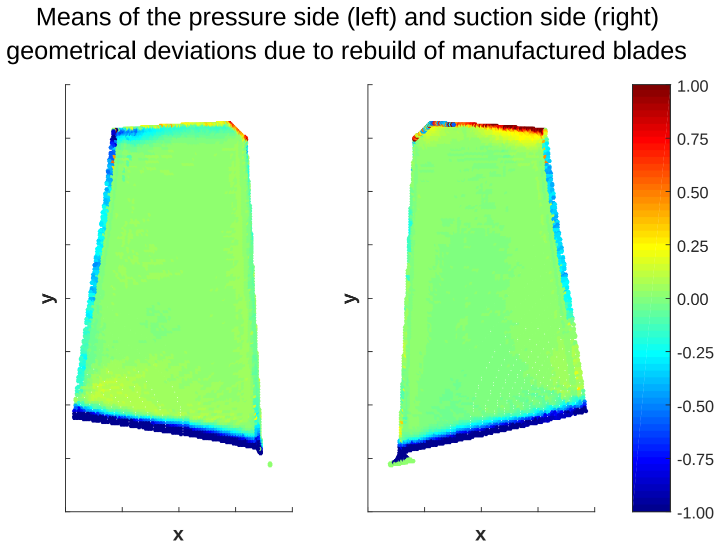

Figure 6.1: The mean geometrical deviation differences due to the preprocessing rebuild of the manufactured blades for the pressure side (left) and the suction side (right).

The standard deviations of the geometrical deviation differences between the rebuilt and the measured manufactured blades can be seen in figure 6.2. This figure shows that the geometrical deviation differences are constant for virtually the entire blade, with the exception of the hub, tip and LE and TE contours. For the blade the standard deviations are roughly $20\mu m$. At the TE the standard deviations are higher, and reach values of roughly $150\mu m$.

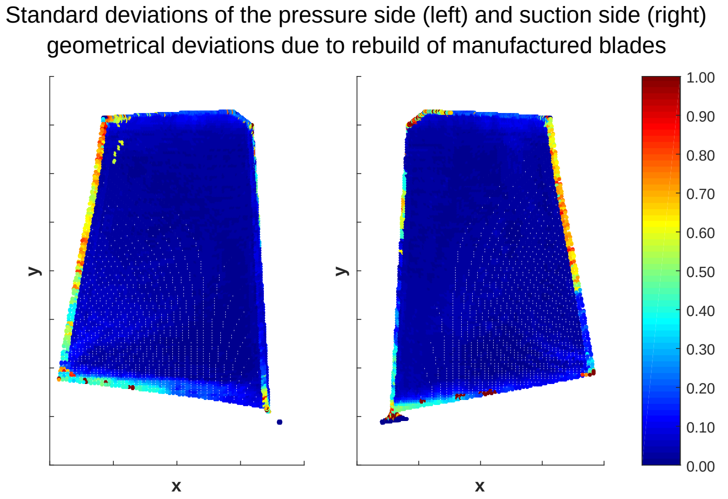

Figure 6.2: The standard deviations of the geometrical deviation differences due to the preprocessing rebuild of the manufactured blades for the pressure side (left) and the suction side (right).

Comparing the standard deviations of the deviations due to the rebuild process, as seen in figure 6.2 to those caused by the manufacturing process, as seen in figure 5.15, it can be seen that the standard deviations due to the rebuild process are relatively small.

To summarise, figures 6.2 and 6.1 show that geometrical deviations due to the rebuilding of the manufactured blades are very small compared to the geometrical deviations due to manufacturing. Therefore, the rebuilt manufactured blades can be considered converged. Because of this, the subsequent performance analysis is representative of the manufactured blades.

Convergence of the rebuilt measured deteriorated blades

Figure 6.3 shows the mean surface deviation differences that occur due to the rebuilding process of the deteriorated blades. Note that the color range of this figure is the same as the color range of figure 5.22, which shows the mean surface deviations due to deterioration.

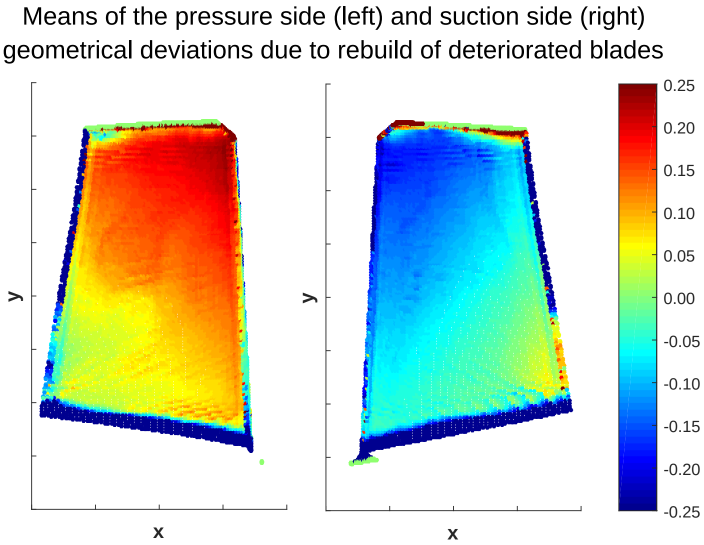

Figure 6.3: The mean geometrical deviation differences due to the preprocessing rebuild of the deteriorated blades for the pressure side (left) and the suction side (right).

Similarly to figures 6.2 and 6.1, figure 6.3 shows that the hub, tip and TE regions have relatively large deviations. This is due to the same reasons given in the discussion of figure 6.1. It can be seen that the deviation differences due to rebuild of deterioration blades change along the span. The largest differences of the blade are at the LE tip, where they are roughly $-40\mu m$ on the suction side, and roughly $+50\mu m$ on the pressure side. This suggests that the rebuilt deteriorated blades are roughly $10\mu m$ thicker than their measured counterparts.

Note that the mean geometrical deviation differences due to the rebuild process of deteriroated blades are roughly the same size as the geometrical differences due to deterioration, as seen in figure 6.3. Furthermore, the differences caused by the rebuild and deterioration are opposite in sign. This suggests that the rebuild process did not account for the changes caused by deterioration.

The standard deviations of the geometrical deviation differences between the rebuilt and the measured deteriorated blades can be seen in figure 6.4. This figure shows that the standard deviations and means of the geometrical deviation differences follow a similar pattern, which also matches the patter of the geometrical deviations caused by deterioration, as seen in figure 5.22. The standard deviations near the LE tip are roughly $40\mu m$.

In the case of the impact of deterioration on blade geometry, the pattern of geometrical deviations was not very uncertain, as seen in figure 5.23, where the blade deviation uncertainty originated from how much the deteriorated blades are angled in the tangential direction. This shows that if the deteriorated blades would have been rebuilt simply missing the pattern of deviations, it would not appear on figure 6.4.

This suggests that the differences due to the rebuild process of deteriorated blades are caused by exaggerated, or underrated, rebuilds of the deterioration effects. Note that the actual deviations caused by deterioration, and the mean and standard deviations of geometrical deviation differences due to the rebuild of deteriorated blades are all relatively small, with values of $\approx 50\mu m$. Therefore the variability near the LE tip could be due to the process accuracy.

To summarise, figures 6.3 and 6.4 show that geometrical deviations due to the rebuilding of the deteriorated blades are approximately equal to the geometrical deviations due to deterioration. Therefore, the rebuilt deteriorated blades cannot be considered well converged. Because of this, the subsequent performance analysis likely does not show the true impact of deterioration.

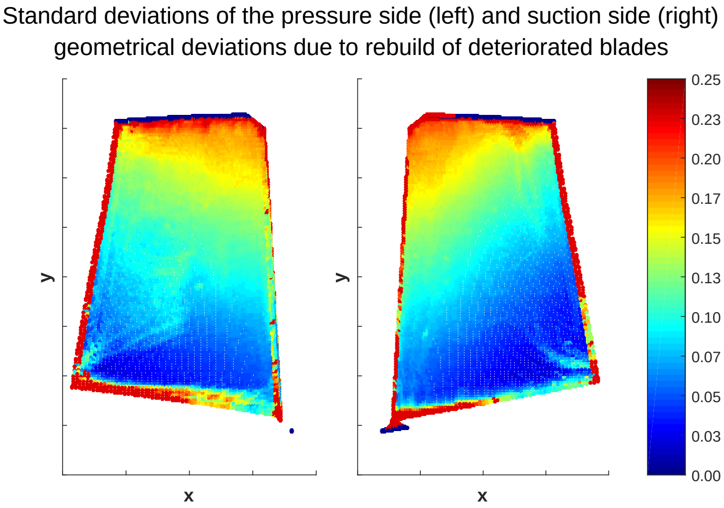

Figure 6.4: The standard deviations of the geometrical deviation differences due to the preprocessing rebuild of the deteriorated blades for the pressure side (left) and the suction side (right).

Convergence of the actual mesh to the measured blades

As the mesh is created based on the preprocessed rebuilt geometry it will carry with it the differences induced by the preprocessing rebuild process. As the rebuilt of the deteriorated blades was deemed insufficient, the mesh convergence analysis is performed only for the manufactured blades.

The evaluation of geometrical convergence of the mesh to the measured geometry is done for an arbitrarily selected example manufactured blade. However, the observed differences between the mesh and measured geometries for the example blade are representative for other blades as well.

Figure 6.5 shows a comparison of the original measured blade geometry and its corresponding CFD mesh visualised in Paraview software. The measured blade geometry is drawn in cyan, and the mesh blade geometry is drawn in magenta. This comparison only allows to judge which geometry is in front of the other for the given point of view.

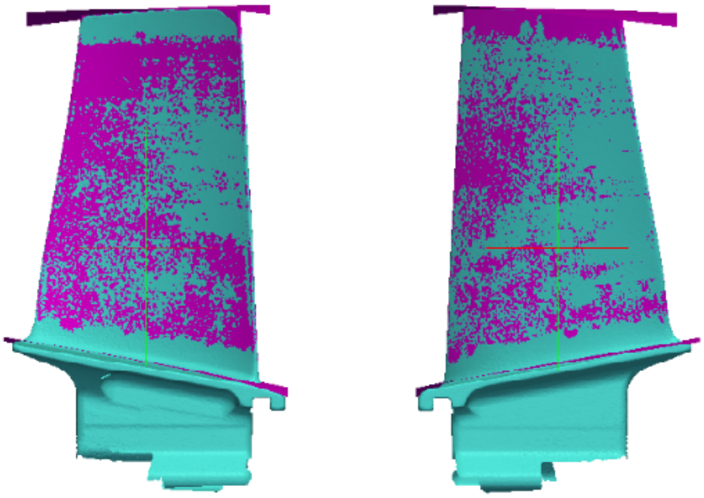

Figure 6.5: The convergence of the mesh to the measured manufactured blade geometry PS (left) and the SS (right). The image aspect ratio has been distorted.

It can be seen that the PS, shown on the left half of figure 6.5, is colored mostly in magenta, with cyan dots scattered over its surface. The SS displays a pattern opposite to that of the PS. In the regions where dots can be seen scattered over the surface the deviations between the mesh and the measured geometry are minimal, and only local. In regions where no dots can be seen the offsets are more systematic.

The most obvious regions where there are systematic offsets are at the tip and hub. Here the geometry has not been sampled and has been only extended to fill the annulus. The regions of more systematic offsets on the blades are near the TE, where the blades appears slightly too thick. However, even in these regions individual dots can be seen.

As the differences of the mesh appear to be relatively small, it can be considered geometrically representative of the measured manufactured blades. Therefore the subsequent analysis of the performance QoI is representative as well.

6.2. Selection of boundary conditions for 3D Hydra CFD

This section describes the selection of boundary conditions used in the evaluation of manufactured and deteriorated blade scans. As a direct comparison of the blade performances is desired, identical boundary conditions have been used for both. First, the selection of the operational conditions under which the blades are evaluated is discussed is subsection 6.2.1.

Subsection 6.2.2 shows the computational domain, discusses the types of boundary conditions used, and visualises the prescribed inlet and outlet flow parameters.

Lastly, subsection 6.2.3 discusses the limitations of the performed calculations. This subsection lists the features of the problem that were not modelled and therefore do not have their impacts included in the performance results.

6.2.1. Operational conditions

As the engines operate under cruise conditions most of the time, these conditions have been selected for the evaluation of the blades. A single operational point on the 100% speed line, at the maximum efficiency has been selected for CFD evaluation.

6.2.2. Boundary condition visualisation

The problem geometry of a single rotor blade evaluated by Hydra can be seen in figure 6.6. This figure shows the blade outline in black, the tangential faces in blue and the axial faces in brown. For the tangential faces periodic boundary conditions are prescribed. For the axial faces the inflow parameter profiles are selected by the user. The top and bottom surfaces are not coloured, and are modelled as walls.

As such only the flow parameters on the axial faces need to be prescribed. It is assumed that the inflow and outflow flow parameters are constant in the direction of rotation. Therefore, only a single spanwise profile needs to be prescribed for each of the flow parameters in order to describe the entire inflow and outflow fields.

The inflow and outflow fields can be seen in figure 6.7. This figure shows that for the inlet several flow parameters have been chosen, which include total pressure and temperature, the whirl and pitch angles and a turbulence parameter. For the outlet only the flow function f is prescribed, and is kept constant along the span.

The flow function is essentially a measure of the axial Mach number, and its derivation can be seen in equation 6.1. The Mach number is a non-dimensional number that allows the comparison of flow fields. As such, setting the inlet temperature and mass flow, the blade rotational velocity and the outlet flow function, a direct comparison of the flow fields can be made. This allows for a comparison of performance of individual blades.

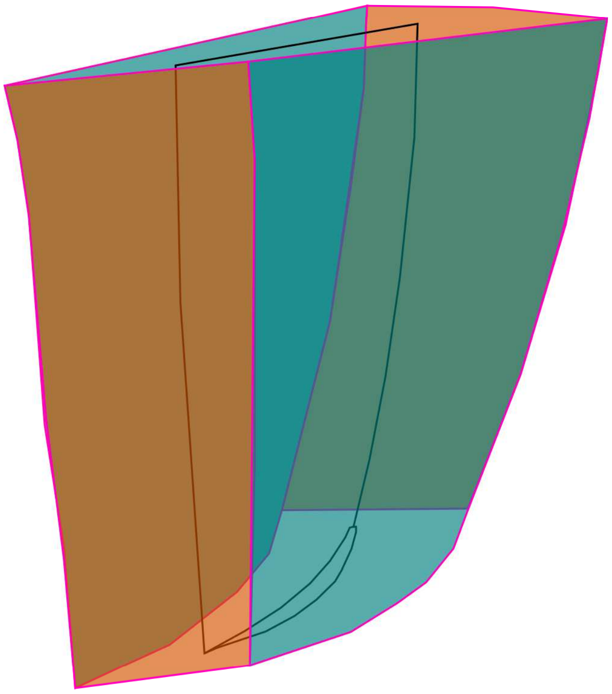

Figure 6.6: The problem geometry of 3D CFD of a single rotor blade.

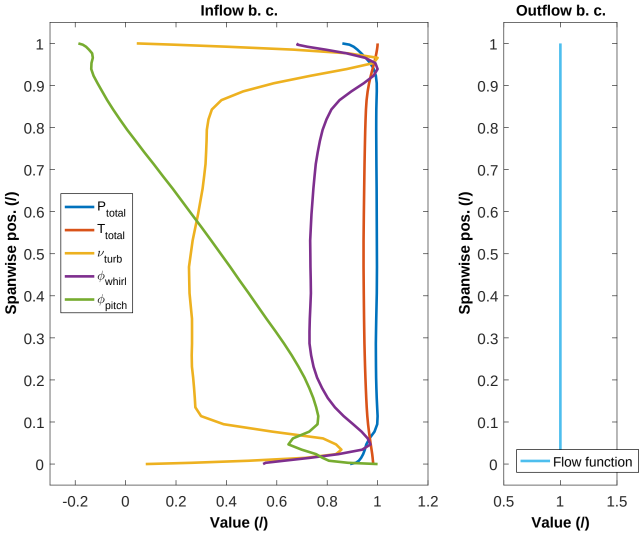

Figure 6.7: The spanwise inflow and outflow profiles of the boundary condition flow parameters.

$$Ma_{axial} = \frac{v_{axial}}{\sqrt{\kappa R T_{total}}} = \frac{\dot{m}}{\rho A} \frac{1}{\sqrt{\kappa R T_{total}}} = \frac{\dot{m} R T_{total}}{p_{total} A} \frac{1}{\sqrt{\kappa R T_{total}}} = \dot{m} \frac{\sqrt{T_{total}}}{p_{total}} \frac{\sqrt{R}}{\sqrt{\kappa} A} \tag{6.1}$$

$$where \quad v_{axial} = \frac{\dot{m}}{\rho A} \quad and \quad \rho = \frac{p_{total}}{R T_{total}} \tag{6.2}$$

$$f = \dot{m} \frac{\sqrt{T_{total}}}{p_{total}} \quad and \quad \frac{\sqrt{R}}{\sqrt{\kappa} A} = const. \tag{6.3}$$

However, the comparison of the rotor performance does not directly show what happens to the engine performance. As the performance of a single component drops or rises, the performance of other components changes as well. This can change the inflow and outflow conditions of the initial component as well, changing its performance as well. This can lead to different compressor blades operating in different inflow and outflow conditions in the same flight regime. This would require the simulation of other engine modules as well, which is complex and computationally expensive. A direct aerodynamic analysis of the compressor is therefore preferred. However, the changes in performance seen in the results of the compressor calculation cannot be claimed to be whole engine performance changes.

6.2.3. Limitation of the performance evaluation

Within Hydra, the geometry is prescribed separately as the evaluated geometry, boundary conditions, the tip gap, and the surface roughness of the blade.

Subsection 5.6.3 showed that the blades considered do not appear to have their blade heights reduced. As no other information is available to calculate the tip gaps, as per subsection 5.2.2, the impact of tip gap uncertainty cannot be included in the calculation. Therefore the tip gaps were specified as a constant percentage of span. Specifically, the tip gaps were specified as the tip gap of the nominal blades under cruise operating conditions.

Subsection 4.1.3 explains that the measurement of the surface roughness could not be captured with the scanner used in the measurement campaign. As no other measurements of the surface roughness were available it was excluded from the analysis.

As such, the results of the performance evaluation include only the effects of the blade geometrical uncertainty, without including the tip gaps and surface roughness. This means that the actual performance differences due to deterioration will be larger than the ones obtained from the performance analysis.

6.3. CFD results analysis

This section shows the comparison of the CFD results of the nominal CAD design and the measured manufactured and deteriorated blades. First, subsection 6.3.1 establishes the QoI for comparison. These are the quantities used in the subsequent discussion of the impacts of manufacturing and deterioration scatter.

Subsection 6.3.2 gives an overview of the convergences of the individual calculations. The understanding of convergence is important as comparing results at different convergence levels might lead to inaccurate conclusions.

Subsection 6.3.3 then presents the impacts of manufacturing and deterioration on the QoI. These are obtained by comparing the individual blade performances to the nominal CAD performance.

6.3.1. Quantities of interest

The flow parameters used in the analysis of the performance of the measured blades are presented here. Furthermore the formulae that is used to calculate them is given as well.

Total pressure ratio

The main task of the compressor is to increase the pressure of the working fluid. A measure for the pressure increase is the pressure ratio, which is calculated as the ratio of the total inlet and outlet pressures as seen in equation 6.4.

$$\pi_{total} = \frac{p_{total\ out}}{p_{total\ in}} \tag{6.4}$$

The total pressures are calculated for the entire inlet and outlet area. These values were then aggregated into a single value by mass averaging over the inlet or outlet area. The ratio of the mass averaged pressure ratio is used in the subsequent performance analysis.

Relative inlet and outlet angles and flow turning

The pressure rise in the compressor is achieved, roughly speaking, by accelerating the flow in the rotor to increase its dynamic pressure component, and decelerating in the stator, to convert the dynamic pressure to static. The flow is mostly accelerated in the direction the compressor is turning, and the amount by which it is accelerated can be determined by the amount of turning that occurred.

The turning can be calculated as the difference of relative inlet and outlet angles of the blade, as seen in equation 6.5. Note that this assumes that the turning of the flow is the same as the change of the relative angles, which is not necessarily true. The relative inlet and outlet angles can in turn be calculated from the axial and tangential velocities of the flow at the inlet and outlet.

$$\Delta\beta = \beta_{out} - \beta_{in} \tag{6.5}$$

$$\beta = atan\left(\frac{v_{rotation} - v_{tangential}}{v_{axial}}\right) \tag{6.6}$$

The tangential and axial velocity outputs of Hydra, $v_{tangential}$ and v_{axial}, are also mass averaged to give a single outlet value, expressed in m/s. The blade rotational velocity, $v_{rotation}$, is prescribed as an input to Hydra, and is given in radian per second. In order to convert to m/s the average annulus radius was calculated by assuming that the mass flow is uniformly distributed over the inlet and outlet areas. The average radius was calculated such that the are above it and below it are the same.

Loss coefficient

The loss coefficient, Y, is defined as the change of stagnation pressure over a blade, which is non-dimensionalised by dividing with the inlet dynamic pressure. It is calculated as per equation 6.7.

$$Y = \frac{p_{total\ in} - p_{total\ out}}{p_{total\ in} - p_{static\ in}} \tag{6.7}$$

Note that for a compressor the outlet pressure is higher than the inlet pressure, therefore the loss coefficient will be negative. In this case, the loss coefficient is the measure of the total pressure increase as a percentage of the inlet dynamic pressure. For the calculation of the loss coefficient the mass averaged values output from Hydra were used.

Static pressure coefficient

The static pressure coefficient is similar to the loss coefficient Y, and is defined as the change of static pressure over the blade, which is non-dimensionalised by the inlet dynamic pressure. In this case, the static pressure coefficient is the measure of the static pressure increase as a percentage of the inlet dynamic pressure. For the calculation of the static pressure coefficient the mass averaged values output from Hydra were used.

$$Cp_{static} = \frac{p_{static\ out} - p_{static\ in}}{p_{total\ in} - p_{static\ in}} \tag{6.8}$$

Total temperature ratio

As a result of work being done on the working gas, its pressure and temperature change. In the case of compressors, the increase of pressure is desired, but the increase of temperature is not. The ratio of the inlet and outlet temperatures is therefore another QoI, and can be calculated using equation 6.9.

$$\tau_{total} = \frac{T_{total\ out}}{T_{total\ in}} \tag{6.9}$$

Furthermore, the efficiency of the stage can be calculated using the total inlet and outlet temperatures, if the adiabatic temperature corresponding to the outlet pressure and inlet conditions is known. Therefore the total temperature ratio can be used as an estimate for efficiency. In the case where the inlet total temperatures are prescribed, low values of τ_{total} mean better efficiency than high values of τ_{total}. The calculation of τ_{total} is performed using Hydra mass averaged output total temperature values.

Isentropic efficiency

The compressor isentropic efficiency is the measure of how well the actual compressor approximates an isentropic compressor. It can be calculated using equation 6.10. Obviously, higher efficiency values are preferred over lower ones.

$$\eta_{isentropic} = \frac{\frac{P_{total\,out}}{P_{total\,in}}^{\frac{\kappa-1}{\kappa}} - 1}{\frac{T_{total\,out}}{T_{total\,in}} - 1} \tag{6.10}$$

Mass flow

For a compressor the achieved pressure ratio is important, but the amount of the working fluid delivered at the exit of the compressor is also important. This is measured as the inlet mass flow. Note that the boundary conditions prescribe only the inlet total pressure, and through the outlet capacity the outlet Mach number. Due to this, the mass flow is not prescribed as a part of the boundary conditions, and is therefore calculated by Hydra. In the following analysis the Hydra averaged output mass flow is used.

6.3.2. Convergence of the 3D CFD calculations

The results of the CFD converge over a number of iterations. As only converged solutions should be used in the analysis of impact of geometrical variability on performance, this section shows the convergence of the calculations made. This is shown for the residuals of the calculations as well as the QoI for both the evaluation of the manufactured and deteriorated blades.

Convergence of residuals

The flow residuals are changes between the solutions at subsequent iterations. The residuals are given as logarithms of the actual values, therefore their values are indicative of the order of magnitude of the actual residuals. Figure 6.8 shows the residuals of the solutions for the manufactured and deteriorated blades in the left and right panels respectively.

It can be seen in figure 6.8 that the residuals of individual calculations all behave similarly. Specifically, all residuals are virtually the same for roughly the first 100 iterations. In this region the residuals experience two steep descends separated by a very short region of slower descend. After roughly 100 iterations the residuals decrease linearly at a shallower rate than previously. The exact rate differs from calculation to calculation. This behaviour is the same for the residuals corresponding to evaluations of manufactured and deteriorated blades.

In both the left and the right panels the residuals reach values of approximately -14, with the residuals on the left panel reaching somewhat lower values than those on the right panel. The residuals on the left panel are also more grouped, whereas those on the right panel are more scattered. In both cases a single clear outlier is visible, however it does not correspond to the evaluation of the same blade in manufactured and deteriorated states.

Convergence of quantities of interest

The convergence of the QoI for the evaluation of manufactured blades can be seen in figure 6.9. This figure shows the convergence of QoI normalised by their converged value. The converged value has been calculated as the average value of the last 50 iterations. The left panel shows the convergence for all 500 iterations, while the right panel shows the convergence above 200 iterations.

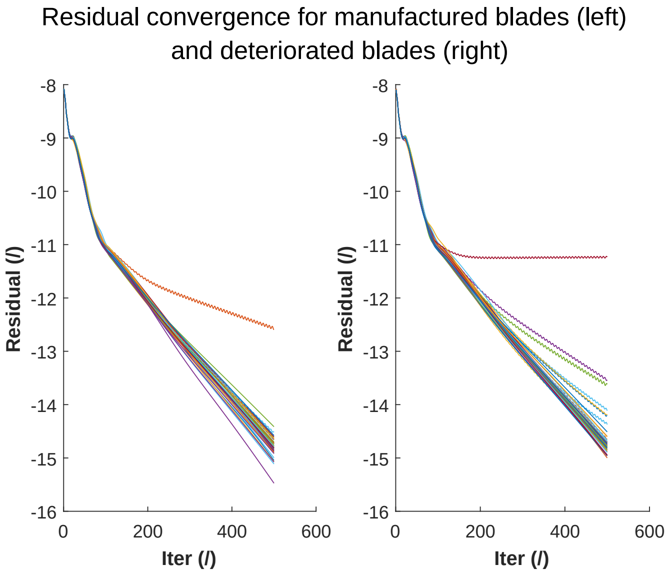

Figure 6.8: Convergence of flow field residuals for the evaluations of manufactured blades (left) and deteriorated blades (right)

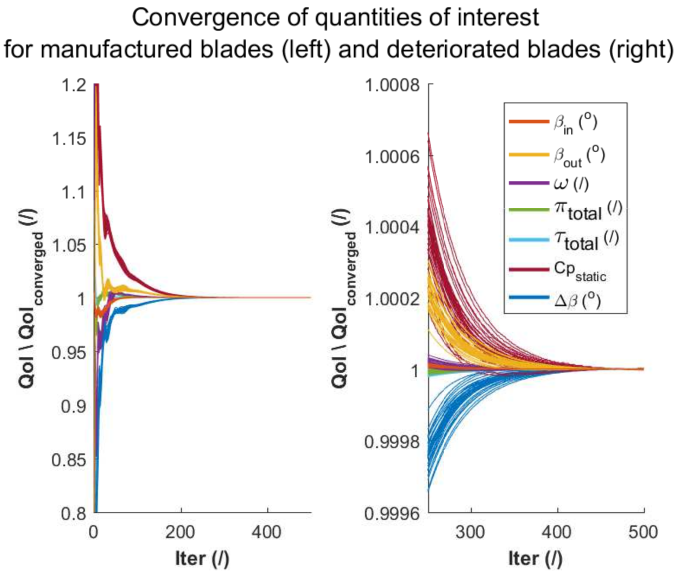

Figure 6.9: Convergence of QoI for the evaluations of manufactured blades.

In the left panel it can be seen that initially the values can be more than ±20% of the final converged value. Within 200 cycles the QoI appear to converge to their final value. The right panel shows that at 200 cycles the differences to the final values are −0.02% to +0.07% of the final values. Furthermore, at 400 cycles the differences are well within ±0.01% of the final values and at 500 cycles no differences in the current order of magnitude can be observed.

The convergences of the QoI for the evaluation of manufactured and deteriorated blades are virtually the same. Because of this, the plot showing the convergence of the latter has been omitted here as it does not add to the discussion. Based on the observed convergence of the QoI, the CFD evaluations of manufactured and deteriorated blades can be assumed converged for the purposes of the subsequent performance study.

6.3.3. Performance impact of manufacturing and deterioration uncertainty

In this section the performance impacts of manufacturing and deterioration are discussed by comparing the QoI identified in subsection 6.3.1. The impact of manufacturing can be determined by subtracting the converged nominal QoI value from the values from the CFD evaluation of manufactured blades. The impact of deterioration can be obtained by subtracting the converged manufactured QoI values from the deteriorated values. The converged QoI values are calculated by averaging the values of the last 50 iterations. All values plotted have been normalised by an appropriate length scale.

The comparison of manufacturing and deterioration is performed by comparing the histograms of the converged value differences. The figures showing the comparison of the manufacturing and deterioration impacts show the differences to the nominal performance on the left panel, and the individual impacts of manufacturing and deterioration on the right panel.

The individual impact of manufacturing is calculated as the difference between the QoI of the manufactured and the nominal blades and are plotted in blue. The individual impact of deterioration is calculated as the difference between the QoI of the deteriorated and the manufactured blades, and is plotted in brown. The differences are calculated separately for each individual blade.

Note that as the rebuild of scanned deteriorated blades did not converge to the measured blades, as shown in section 6.1.3, the results shown in this plot are not representative of the real impact of geometry deterioration. The actual impact is likely exaggerated, but is nevertheless included for completeness and as a reference point for the analysis of the impact of manufacturing.

The histograms are normalised such that their areas equal 1. They are complemented with a fitted Gaussian distribution, and a line that denotes the mean of the calculated QoI values. Lastly, a line denoting the nominal performance is drawn in black.

Total pressure ratio

The left panel of figure 6.10 shows that the manufactured and deteriorated blades appear to achieve, on average, higher total pressure ratios. This is more so for manufactured blades than for deteriorated blades. Furthermore, the deteriorated blades appear to have less scatter as the brown PDF corresponding to the deteriorated blades is higher than the blue PDF corresponding to the manufactured blades. The maximum total pressure differences are of the order of roughly ±0.05% of the mass averaged inlet total pressure. For a hypothetical inlet total pressure of $2bar$ that amounts to an increase of the mass averaged total outlet pressure of $0.01bar$.

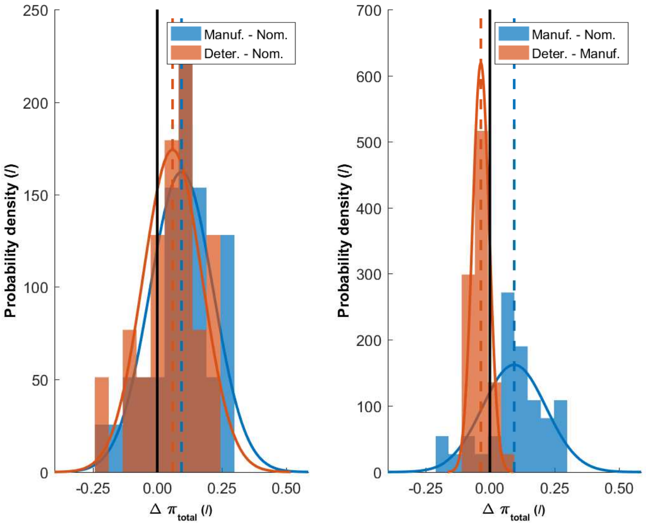

Figure 6.10: Impact of manufacturing and deterioration on total pressure ratio relative to the nominal pressure ratio (left) and as individual impacts (right).

The right panel of figure 6.10 shows a direct comparison of impacts of deterioration and manufacturing. In this panel the impact on the total pressure ratio due to deterioration can be seen more clearly. It can be seen that this impact is significantly smaller than the impact due to manufacturing even though the geometrical scatter due to deterioration has likely been exaggerated in the preprocessing. Lastly, deterioration introduces less performance scatter than manufacturing.

Relative inlet angle

The left panel of figure 6.11 shows that the manufactured and deteriorated blades appear to operate, on average, with smaller than nominal, and nominal relative inlet angles respectively. The deteriorated blades appear to have less scatter as the brown PDF corresponding to the deteriorated blades is higher than the blue PDF corresponding to the manufactured blades. The maximum inlet relative angle differences to the nominal are of the order of roughly $\pm 0.1^o$ of the mass averaged relative inlet angles.

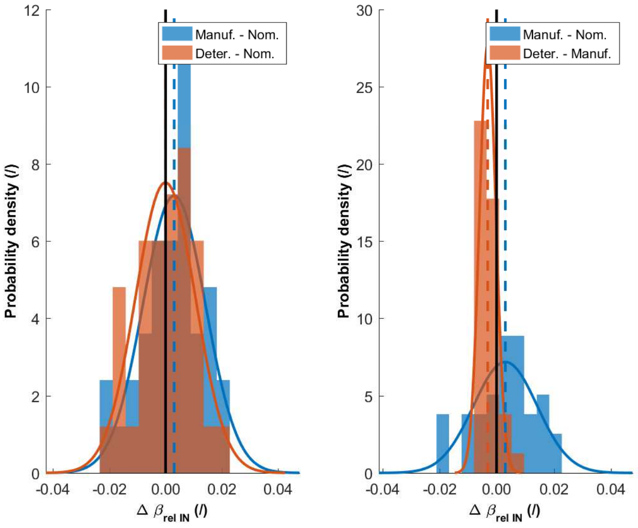

Figure 6.11: Impact of manufacturing and deterioration on the relative inlet angle with respect to the nominal relative inlet angle (left) and as individual impacts (right).

The observed on average increase of the inlet dynamic pressure means that the inlet velocity has also increased on average. Because of this, and because the prescribed rotational velocity is constant, the relative inlet angle decreases. In turn, the angle of attack experienced by the aerofoil also decreases, which in turn decreases the lift coefficient. In this case, as the changes in velocity are relatively small, the changes in the relative inlet angle are also small.

The right panel of figure 6.11 shows that the impact of deterioration on the relative inlet angle is opposite to the impact of manufacturing. Deterioration appears to have a more constant impact on the inlet relative angles than manufacturing.

Relative outlet angle

The left panel of figure 6.12 shows that the manufactured and deteriorated blades appear to operate, on average, with higher than nominal relative outlet angles. The deteriorated blades appear to have less scatter as the brown PDF corresponding to the deteriorated blades is higher than the blue PDF corresponding to the manufactured blades. The maximum outlet relative angle differences to the nominal are of the order of roughly $\pm 0.5^o$ of the mass averaged relative outlet angles. It can be seen that these values are roughly 5 times larger than the relative inlet angles shown in figure 6.11.

The right panel shows that the offset of the outlet relative angle is roughly 0.2^o. Furthermore, it can be seen that deterioration does not introduce a offset in the relative outlet angles. Lastly, the scatter introduced by manufacturing is larger than that introduced by deterioration, even though that is likely exaggerated.

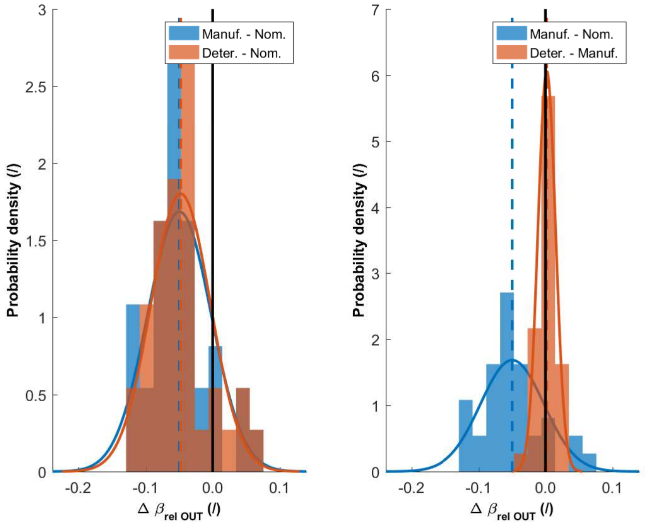

Figure 6.12: Impact of manufacturing and deterioration on the relative outlet angle with respect to the nominal outlet relative angle (left) and as individual impacts (right).

Turning

The left panel of figure 6.13 shows that the manufactured and deteriorated blades appear to operate, on average, with higher than nominal turning. The deteriorated blades appear to have less scatter as the brown PDF corresponding to the deteriorated blades is higher than the blue PDF corresponding to the manufactured blades. The maximum turning differences to the nominal are roughly -0.2 ± 0.4^o of the mass averaged relative outlet angles. As turning is just a subtraction of the relative outlet angle from the relative inlet angle, most of the turning differences come from the relative outlet angle differences.

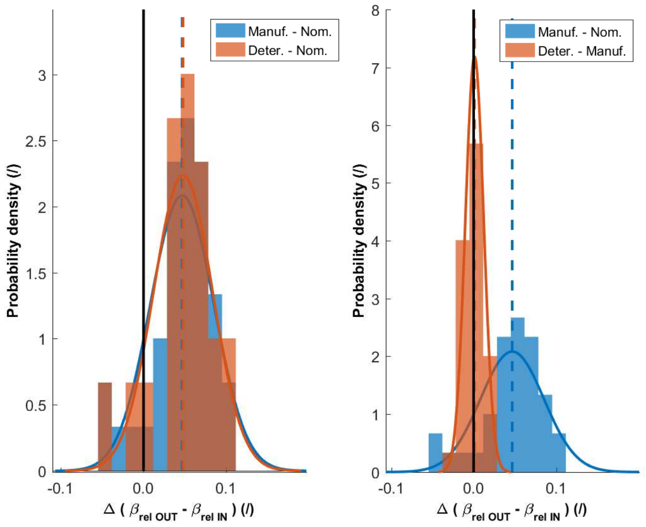

Figure 6.13: Impact of manufacturing and deterioration on the turning with respect to the nominal turning (left) and as individual impacts (right).

The right panel shows that manufacturing causes a turning offset of roughly -0.2^o, while deterioration does not cause a turning offset. The scatter due to manufacturing is larger than the scatter due to deterioration.

Loss coefficient

The left panel of figure 6.14 shows that the manufactured and deteriorated blades appear to operate, on average, with lower than nominal loss coefficient. The deteriorated blades appear to have less scatter as the brown PDF corresponding to the deteriorated blades is higher than the blue PDF corresponding to the manufactured blades. The maximum loss coefficient differences to the nominal are roughly -0.01 of the mass averaged loss coefficient. This is a reduction of 1% of the total pressure change relative to the inlet dynamic pressure.

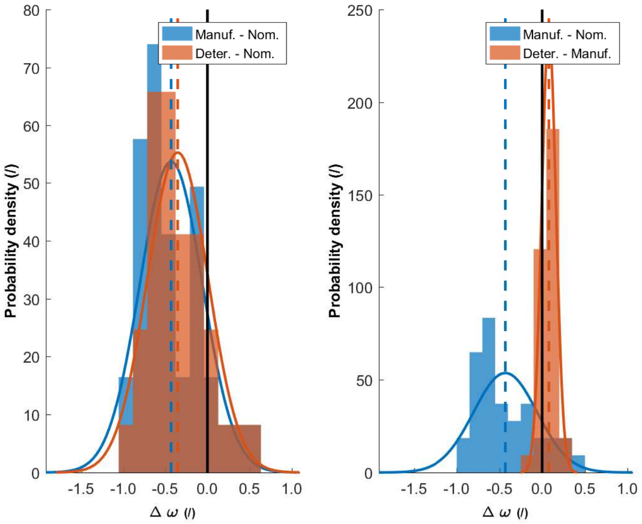

Figure 6.14: Impact of manufacturing and deterioration on the loss coefficient with respect to the nominal loss coefficient (left) and as individual impacts (right).

The loss coefficient decreases despite an increase of the outlet total pressure, relative to the nominal values. This is caused by an increase of the inlet dynamic pressure. Proportionally to the nominal values, the increase of the inlet dynamic pressure is higher than the increase of the outlet total pressure.

The right panel shows that manufacturing causes a loss coefficient offset of roughly -0.01, while deterioration causes a loss coefficient offset that is positive, but much smaller than 0.005. The scatter due to manufacturing is larger than the scatter due to deterioration.

Static pressure coefficient

The left panel of figure 6.15 shows that the manufactured and deteriorated blades appear to operate, on average, with lower than nominal static pressure coefficient. The deteriorated blades appear to have less scatter as the brown PDF corresponding to the deteriorated blades is higher than the blue PDF corresponding to the manufactured blades. The maximum static pressure coefficient differences to the nominal are roughly -0.008 of the mass averaged static pressure coefficient. This is a 0.8% reduction of the static pressure difference relative to the inlet dynamic pressure.

The static pressure coefficient drops is the static pressure difference across the blade drops, or if the inlet dynamic pressure increases. In this case, both of this conditions are fulfilled, and both effects reduce the static pressure coefficient together. This shows that, on average, for a given inlet dynamic pressure the manufactured blades deliver a smaller static pressure increase.

The right panel shows that manufacturing causes a static pressure coefficient offset of roughly -0.008, while deterioration causes a static pressure coefficient offset that is negative, but much smaller than 0.005. The scatter due to manufacturing is larger than the scatter due to deterioration.

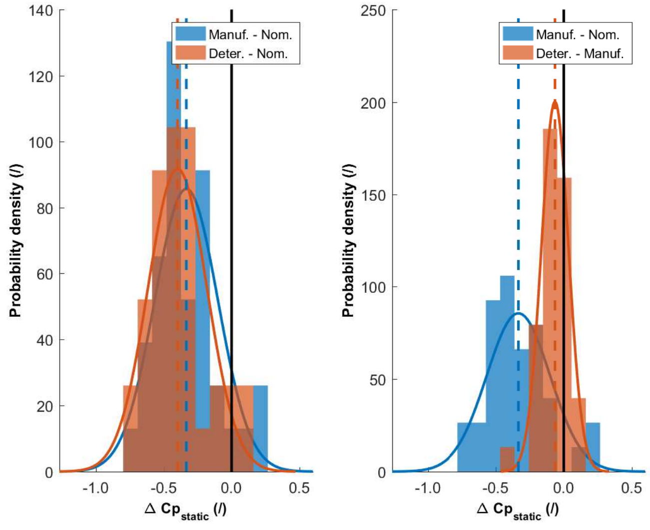

Figure 6.15: Impact of manufacturing and deterioration on the static pressure coefficient with respect to the nominal static pressure coefficient (left) and as individual impacts (right).

Total temperature ratio

The left panel of figure 6.15 shows that the manufactured and deteriorated blades appear to operate, on average, with a higher than nominal total temperature ratio. The deteriorated blades appear to have less scatter as the brown PDF corresponding to the deteriorated blades is higher than the blue PDF corresponding to the manufactured blades. The maximum total temperature ratio differences to the nominal are roughly $0.5 \cdot 10^{-3} K$.

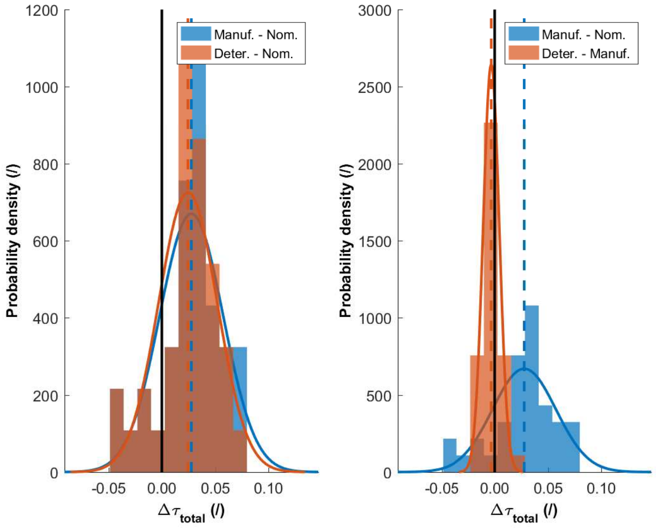

Figure 6.16: Impact of manufacturing and deterioration on the total temperature ratio with respect to the nominal total temperature ratio (left) and as individual impacts (right).

The right panel shows that manufacturing causes a total temperature ratio offset of roughly $0.5 \cdot 10^{-3} K$, while deterioration causes a total temperature ratio offset that is negative, but much smaller than $0.5 \cdot 10^{-3} K$. The scatter due to manufacturing is larger than the scatter due to deterioration.

Isentropic efficiency

The left panel of figure 6.17 shows that the manufactured and deteriorated blades appear to operate, on average, with a lower than nominal isentropic efficiency. The deteriorated blades appear to have more scatter as the brown PDF corresponding to the deteriorated blades is lower than the blue PDF corresponding to the manufactured blades. The maximum isentropic efficiency differences to the nominal are roughly $-3 \cdot 10^{-3}$, which is equal to an isentropic efficiency reduction of roughly 0.3%.

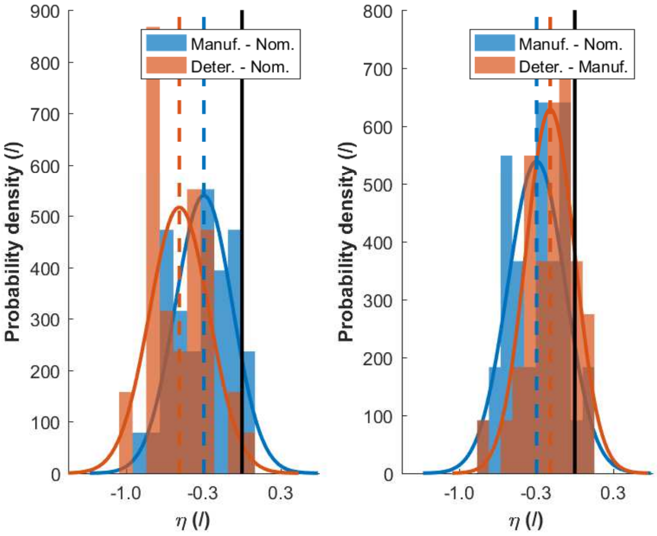

Figure 6.17: Impact of manufacturing and deterioration on the isentropic efficiency with respect to the nominal isentropic efficiency (left) and as individual impacts (right).

The right panel shows that manufacturing causes an isentropic efficiency offset of roughly $-1 \cdot 10^{-3} K$, while deterioration causes a isentropic efficiency offset that is also negative and somewhat smaller than $-1 \cdot 10^{-3} K$. The scatter due to manufacturing is smaller than the scatter due to deterioration, although this is likely exaggerated. Overall, manufacturing and deterioration appear to have a similar impact on isentropic efficiency.

Mass flow

The left panel of figure 6.18 shows that the manufactured and deteriorated blades appear to operate, on average, with a higher than nominal mass flow. The deteriorated blades appear to have more scatter as the brown PDF corresponding to the deteriorated blades is lower than the blue PDF corresponding to the manufactured blades. The maximum mass flow differences to the nominal are less than $0.1 kg/s$.

The right panel shows that manufacturing causes a mass flow offset of roughly $0.02 kg/s$, while deterioration causes a mass flow offset that is negative and somewhat smaller than the manufacturing offset. The scatter due to manufacturing is larger than the scatter due to deterioration. As the inlet temperatures and pressures are prescribed, and the inlet cross sectional area remains constant, the mass flow is directly representative of the inlet Mach number as well.

The outlet flow function is prescribed in the boundary conditions, and establishes a relation between the outlet total temperature and pressure, and mass flow. If the outlet total temperature and pressures change the mass flow must also change to accommodate for this. From figures 6.16 and 6.10 it can be seen that the total pressure ratio increases proportionally more than the total temperature ratio. Therefore the temperature and pressure ratio from equation 6.1 reduces, and the mass flow increases to compensate.

6.4. Influence of geometrical parameters on performance

The aim of this section is to present the impact of the individual NACA like geometrical parameters on the performance QoI. The impact of individual geometrical parameters is assessed using CoI (Coefficients of

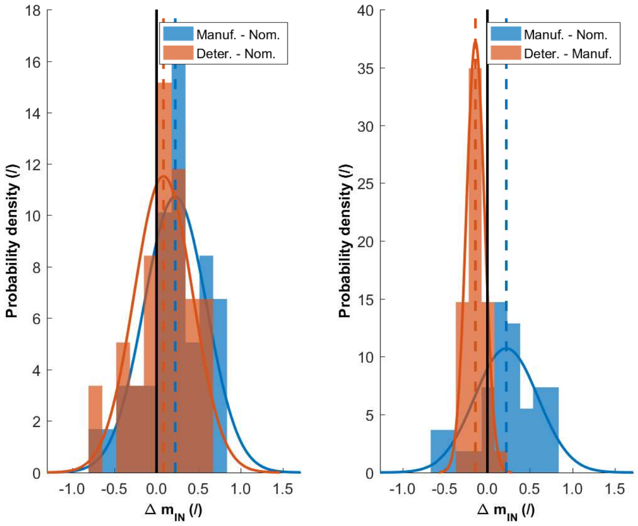

Figure 6.18: Impact of manufacturing and deterioration on the mass flow with respect to the nominal mass flow (left) and as individual impacts (right).

Influence). The calculation of the CoI is discussed in subsection 6.4.1.

Subsection 6.4.2 then presents the CoI that show which aspects of manufacturing and deterioration variability cause the most of the performance scatter.

6.4.1. Calculation of Coefficients of Influence

In order to understand the importance of the geometry on the performance the so called CoI (coefficients of importance) can be calculated based on a polynomial regression of the parameters. The polynomial regression creates a response surface with some residual differences to the data that it is based on. These residuals can be used as a measure of how well the response surface fits the real data by calculating the CoD (coefficient of deterioration) using equation 6.11. In this equation R^2 is the CoD, S_R is the sum of the squared residuals and S_y is the sum of the squared original values. This can be seen as a normalisation of the residuals by the magnitude of the real value. Small R^2 mean that the residuals are a large part of the response surface approximation of the real value, and vice versa.

$$R^2 = 1 - \frac{S_R}{S_y} \tag{6.11}$$

The coefficient of influence is the difference of two such coefficients of determination. Specifically, the CoI of a particular parameter is calculated by subtracting the CoD calculated when all the parameters except the particular parameter is used, from the CoD when all the parameters are used, using equation 6.12. In this equation i is the subscript of the performance QoI, and j is the subscript of the geometrical parameter.

$$CoI_{i,j} = R_i^2 - R_{i,j}^2 \tag{6.12}$$

Within the CoI, the contribution from R_i^2 carries the information on the total residual values of the polynomial regression. The contribution from $R_{i,j}^2$ introduces the information of how much of the residuals are due to parameters other than the particular parameter. The contribution of $R_{i,j}^2$ is calculated only on a subset of the data used in the calculation of R_i^2. The removal of one dimension of a problem will produce a better fit, and therefore $R_{i,j}^2 < R_i^2$. This also means that the CoI ranges between 0 and 1.

If this $R^2_{i,j}$ is small, removing the particular parameter from the regression improved the fit. Therefore this parameter has been predominantly a source of noise in the regression as opposed to a source of the behaviour. Therefore it will have a relatively large CoI denoting a smaller influence on the QoI. Conversely, if a parameter has a large CoI, it also has a larger influence on the QoI.

6.4.2. Impact of individual geometrical parameters on performance quantities of interest

The analysis of the impact of individual geometrical parameters on performance QoI is done separately for manufacturing and deterioration. For each of these the analysis first considers the CoI calculated on the basis of the parameters extracted directly from the measured blades and their corresponding CFD evaluations. Secondly, the CoI calculated based on delta parameters to the reference geometry and performance are considered. The first approach gives a global perspective on the importance of parameters, and the second directly analyses the impact of the differences caused by an uncertain process.

To calculate the CoI the geometrical parameters must first be selected. The selection must be made such that the blade-to-blade behaviour, as seen in section 5.6, is captured. As most of the blade-to-blade differences appeared to be linear, and to reduce the number of parameters in the CoI analysis, the spanwise parameter means and the parameter spanwise slopes were selected. These were calculated individually for all the evaluated blades. The slopes were calculated by fitting a linear function to the parameter spanwise position and its value using least squares regression. The slopes are denoted with a k, and their subscripts show the parameter they represent. The calculation of CoI was performed using Rolls-Royce in-house software ProSi.

CoI of manufactured blades

Figure 6.19 shows the CoI, as calculated by ProSi, of all the selected manufacturing geometrical parameters on their corresponding performance QoI. The geometrical parameters are on the x-axis whereas the performance QoI are on the y-axis. The cells that are coloured blue denote higher CoI, with darker shades of blue denoting CoI above 0.8. On the far left hand of the matrix a column with green coloured cells can be seen. This is the column of R^2, calculated using equation 6.11. Note that all the values in this column are very close to 1, or are 1. This means that the polynomial approximation captured the behaviour well leaving very low residuals.

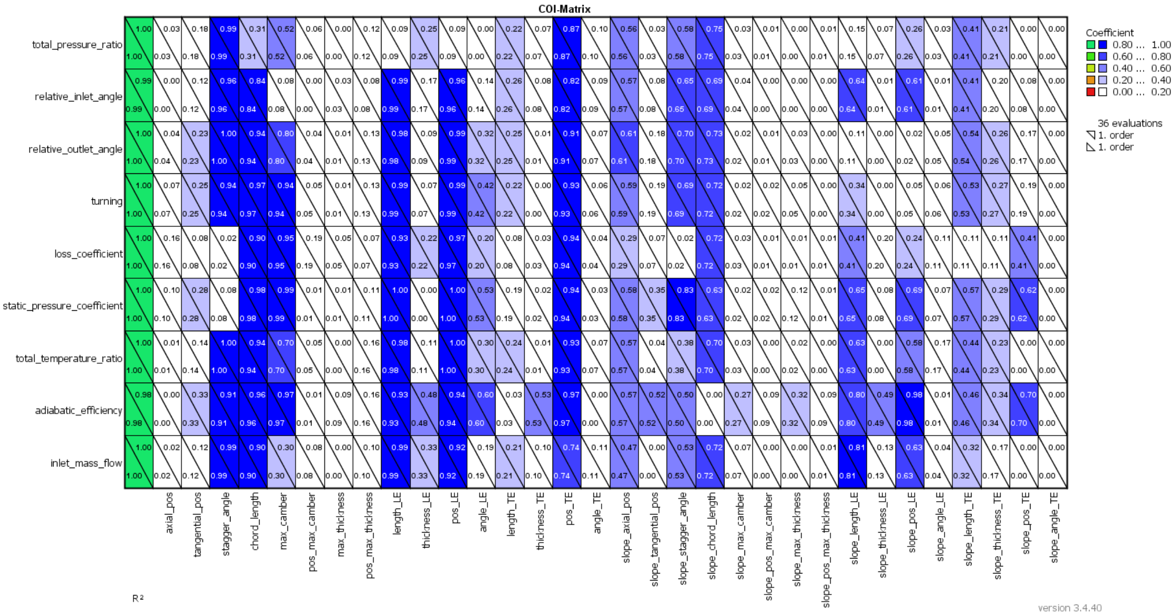

COI-Matrix

	R²	axial_pos	tangential_pos	stagger_angle	chord_length	max_camber	pos_max_camber	max_thickness	pos_max_thickness	length_LE	thickness_LE	pos_LE	angle_LE	length_TE	thickness_TE	pos_TE	angle_TE	slope_axial_pos	slope_tangential_pos	slope_stagger_angle	slope_chord_length	slope_max_camber	slope_pos_max_camber	slope_max_thickness	slope_pos_max_thickness	slope_length_LE	slope_thickness_LE	slope_pos_LE	slope_angle_LE	slope_length_TE	slope_thickness_TE	slope_pos_TE	slope_angle_TE
total_pressure_ratio	1.00	0.03	0.18	0.99	0.31	0.52	0.06	0.00	0.12	0.09	0.25	0.09	0.00	0.22	0.07	0.87	0.10	0.56	0.03	0.58	0.75	0.03	0.01	0.00	0.01	0.15	0.07	0.26	0.03	0.41	0.21	0.00	0.00
relative_inlet_angle	0.99	0.00	0.12	0.96	0.84	0.08	0.00	0.03	0.08	0.99	0.17	0.96	0.14	0.26	0.08	0.82	0.09	0.57	0.08	0.65	0.69	0.04	0.00	0.00	0.00	0.64	0.01	0.61	0.01	0.41	0.20	0.08	0.00
relative_outlet_angle	1.00	0.04	0.23	1.00	0.94	0.80	0.04	0.01	0.13	0.98	0.09	0.99	0.32	0.25	0.01	0.91	0.07	0.61	0.18	0.70	0.73	0.02	0.01	0.03	0.00	0.11	0.00	0.02	0.05	0.54	0.26	0.17	0.00
turning	1.00	0.07	0.25	0.94	0.97	0.94	0.05	0.01	0.13	0.99	0.07	0.99	0.42	0.22	0.00	0.93	0.06	0.59	0.19	0.69	0.72	0.02	0.02	0.05	0.00	0.34	0.00	0.05	0.06	0.53	0.27	0.19	0.00
loss_coefficient	1.00	0.16	0.08	0.02	0.90	0.95	0.19	0.05	0.07	0.93	0.22	0.97	0.20	0.08	0.03	0.94	0.04	0.29	0.07	0.02	0.72	0.03	0.01	0.01	0.01	0.41	0.20	0.24	0.11	0.11	0.11	0.41	0.00
static_pressure_coefficient	1.00	0.10	0.28	0.08	0.98	0.99	0.01	0.01	0.11	1.00	0.00	1.00	0.53	0.19	0.02	0.94	0.03	0.58	0.35	0.83	0.63	0.02	0.02	0.12	0.01	0.65	0.08	0.69	0.07	0.57	0.29	0.62	0.00
total_temperature_ratio	1.00	0.01	0.14	1.00	0.94	0.70	0.05	0.00	0.16	0.98	0.11	1.00	0.30	0.24	0.01	0.93	0.07	0.57	0.04	0.38	0.70	0.03	0.00	0.02	0.00	0.63	0.00	0.58	0.17	0.44	0.23	0.00	0.00
adiabatic_efficiency	0.98	0.00	0.33	0.91	0.96	0.97	0.01	0.09	0.16	0.93	0.48	0.94	0.60	0.03	0.53	0.97	0.00	0.57	0.52	0.50	0.00	0.27	0.09	0.32	0.09	0.80	0.49	0.98	0.01	0.46	0.34	0.70	0.00
inlet_mass_flow	1.00	0.02	0.12	0.99	0.90	0.30	0.08	0.00	0.10	0.99	0.33	0.92	0.19	0.21	0.10	0.74	0.11	0.47	0.00	0.53	0.72	0.07	0.00	0.00	0.01	0.81	0.13	0.63	0.04	0.32	0.17	0.00	0.00

Figure 6.19: CoI of manufactured geometrical NACA like parameter on the performance QoI.

It can be seen that there are several parameters that have a strong influence on all the QoI. These parameters are the stagger angle, chord length, maximum camber, LE length, LE and TE positions, and several parameters characterising spanwise blade-to-blade differences. These are the slopes of axial position, stagger angle and chord length. Some other parameters appear to have strong influence on specific performance parameters. It is interesting to note for example that the blade-to-blade differences in the spanwise distributions of axial position have an influence on the performance. This suggests that 3D flow behaviour changes significantly due to the spanwise blade-to-blade differences in the axial position.

Figure 6.19 shows the impact of the manufactured geometrical parameters on the overall performance of the manufactured blades. In order to isolate the impact of the changes of these parameters due to manufacturing on the performance QoI, a plot of the CoI calculated based on the delta values was made. In this case the delta values were calculated between the manufactured and nominal geometrical parameters, and performance QoI. The CoI based on these delta parameters is shown in figure 6.20.

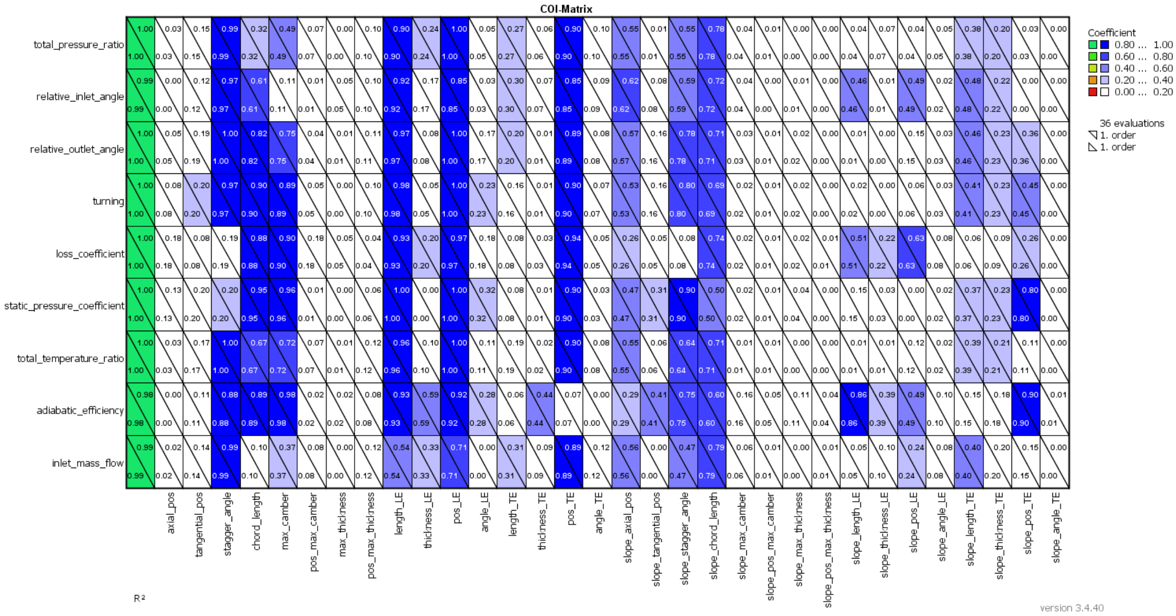

	R²	axial_pos	tangential_pos	stagger_angle	chord_length	max_camber	pos_max_camber	max_thickness	pos_max_thickness	length_LE	thickness_LE	pos_LE	angle_LE	length_TE	thickness_TE	pos_TE	angle_TE
total_pressure_ratio	1.00	0.03	0.15	0.99	0.32	0.49	0.07	0.00	0.10	0.90	0.24	1.00	0.05	0.27	0.06	0.90	0.10
relative_inlet_angle	0.99	0.00	0.12	0.97	0.61	0.11	0.01	0.05	0.10	0.92	0.17	0.85	0.03	0.30	0.07	0.85	0.09
relative_outlet_angle	1.00	0.05	0.19	1.00	0.82	0.75	0.04	0.01	0.11	0.97	0.08	1.00	0.17	0.20	0.01	0.89	0.08
turning	1.00	0.08	0.20	0.97	0.90	0.89	0.05	0.00	0.10	0.98	0.05	1.00	0.23	0.16	0.01	0.90	0.07
loss_coefficient	1.00	0.18	0.08	0.19	0.88	0.90	0.18	0.05	0.04	0.93	0.20	0.97	0.18	0.08	0.03	0.94	0.05
static_pressure_coefficient	1.00	0.13	0.20	0.20	0.95	0.96	0.01	0.00	0.06	1.00	0.00	1.00	0.32	0.08	0.01	0.90	0.03
total_temperature_ratio	1.00	0.03	0.17	1.00	0.67	0.72	0.07	0.01	0.12	0.96	0.10	1.00	0.11	0.19	0.02	0.90	0.08
adiabatic_efficiency	0.98	0.00	0.11	0.88	0.89	0.98	0.02	0.02	0.08	0.93	0.59	0.92	0.28	0.06	0.44	0.07	0.00
inlet_mass_flow	0.99	0.02	0.14	0.99	0.10	0.37	0.08	0.00	0.12	0.54	0.33	0.71	0.00	0.31	0.09	0.89	0.12

	slope_axial_pos	slope_tangential_pos	slope_stagger_angle	slope_chord_length	slope_max_camber	slope_pos_max_camber	slope_max_thickness	slope_pos_max_thickness	slope_length_LE	slope_thickness_LE	slope_pos_LE	slope_angle_LE	slope_length_TE	slope_thickness_TE	slope_pos_TE	slope_angle_TE
total_pressure_ratio	0.55	0.01	0.55	0.78	0.04	0.01	0.00	0.00	0.04	0.07	0.04	0.05	0.38	0.20	0.03	0.00
relative_inlet_angle	0.62	0.08	0.59	0.72	0.04	0.00	0.01	0.00	0.46	0.01	0.49	0.02	0.48	0.22	0.00	0.00
relative_outlet_angle	0.57	0.16	0.78	0.71	0.03	0.01	0.02	0.00	0.01	0.00	0.15	0.03	0.46	0.23	0.36	0.00
turning	0.53	0.16	0.80	0.69	0.02	0.01	0.02	0.00	0.02	0.00	0.06	0.03	0.41	0.23	0.45	0.00
loss_coefficient	0.26	0.05	0.08	0.74	0.02	0.01	0.02	0.01	0.51	0.22	0.63	0.08	0.06	0.09	0.26	0.00
static_pressure_coefficient	0.47	0.31	0.90	0.50	0.02	0.01	0.04	0.00	0.15	0.03	0.00	0.02	0.37	0.23	0.80	0.00
total_temperature_ratio	0.55	0.06	0.64	0.71	0.01	0.01	0.00	0.00	0.01	0.01	0.12	0.02	0.39	0.21	0.11	0.00
adiabatic_efficiency	0.29	0.41	0.75	0.60	0.16	0.05	0.11	0.04	0.86	0.39	0.49	0.10	0.15	0.18	0.90	0.01
inlet_mass_flow	0.56	0.00	0.47	0.79	0.06	0.01	0.00	0.01	0.05	0.10	0.24	0.08	0.40	0.20	0.15	0.00

Figure 6.20: CoI of manufactured geometrical NACA like delta parameters on their corresponding delta performance QoI.

Figure 6.20 is very similar to figure 6.19. The main difference is that most of the medium strength CoI are not present in figure 6.20. This can be seen best in the right part of the figure, where there are now fewer highlighted CoI, and those that are are generally shaded more lightly. This reduction can be observed in the left part of figure 6.20 as well, for example in the case of TE length, which lost influence on some QoI.

The reduction of the CoI in the right part of figure 6.20 means that their impact on the overall performance is relatively strong, but their impact due solely to the manufacturing scatter is not as strong. Furthermore, as the right half of figure 6.20 has few high CoI the blade-to-blade spanwise behaviour due to manufacturing does not appear to result in a performance scatter.

In order to more accurately compare the parameters with highest CoI on the QoI, the exact values of the CoI for the geometrical parameters are plotted as bars in figure 6.21. In this figure the CoI shown in figure 6.19 are used. The selected QoI are the isentropic efficiency, total pressure ration and turning. However, as the geometrical parameters with high CoI generally have an impact on all the QoI, the most important parameters are likely to be identified using just the listed QoI.

Figure 6.21 shows that in the case of manufacturing, the most important parameters are k_{xLE} for the isentropic efficiency, γ for the total pressure rise, and the LE position x_{LE} for turning. Note that the impact of manufacturing on the LE positions does not appear to have introduced any observable blade-to-blade offset differences, or blade-to-blade spanwise behaviour differences, as shown in figure 5.43. Therefore k_{xLE} and x_{LE} likely have high CoI parameters as their removal in the calculation of R^2_{ij} reduces the noise, and thus increases the CoI.

Additionally, note that in the cases of the isentropic efficiency and turning the first few parameters have

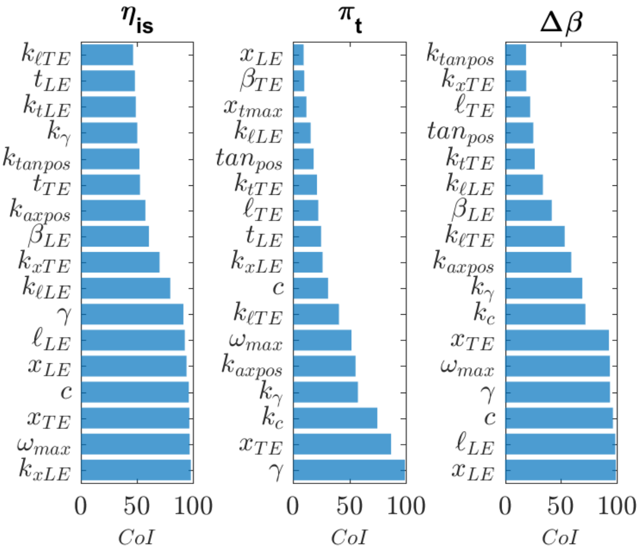

Figure 6.21: Ordered CoI of manufactured geometrical NACA like parameters on their corresponding performance QoI.

virtually the same CoI. As γ is consistently among them it is also the most important of them.

CoI of deteriorated blades

Figure 6.22 shows the CoI, as calculated by ProSi, of all the selected deteriorated geometrical parameters on their corresponding performance QoI. The values in the far left green shaded column are very close to 1, or are 1. The polynomial approximation has therefore captured the behaviour well.

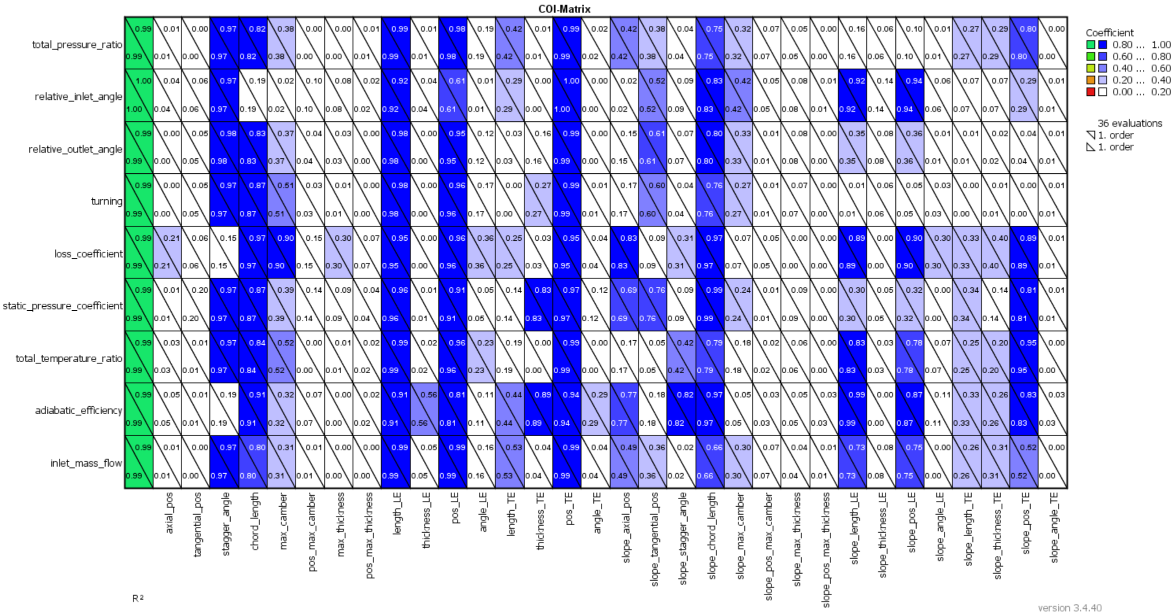
COI-Matrix

	R²	axial_pos	tangential_pos	stagger_angle	chord_length	max_camber	pos_max_camber	max_thickness	pos_max_thickness	length_LE	thickness_LE	pos_LE	angle_LE	length_TE	thickness_TE	pos_TE
total_pressure_ratio	0.99	0.01	0.00	0.97	0.82	0.38	0.00	0.00	0.01	0.99	0.01	0.98	0.19	0.42	0.01	0.99
relative_inlet_angle	1.00	0.04	0.06	0.97	0.19	0.02	0.10	0.08	0.02	0.92	0.04	0.61	0.01	0.29	0.00	1.00
relative_outlet_angle	0.99	0.00	0.05	0.98	0.83	0.37	0.04	0.03	0.00	0.98	0.00	0.95	0.12	0.03	0.16	0.99
turning	0.99	0.00	0.05	0.97	0.87	0.51	0.03	0.01	0.00	0.98	0.00	0.96	0.17	0.00	0.27	0.99
loss_coefficient	0.99	0.21	0.06	0.15	0.97	0.90	0.15	0.30	0.07	0.95	0.00	0.96	0.36	0.25	0.03	0.95
static_pressure_coefficient	0.99	0.01	0.20	0.97	0.87	0.39	0.14	0.09	0.04	0.96	0.01	0.91	0.05	0.14	0.83	0.97
total_temperature_ratio	0.99	0.03	0.01	0.97	0.84	0.52	0.00	0.01	0.02	0.99	0.02	0.96	0.23	0.19	0.00	0.99
adiabatic_efficiency	0.99	0.05	0.01	0.19	0.91	0.32	0.07	0.00	0.02	0.91	0.56	0.81	0.11	0.44	0.89	0.94
inlet_mass_flow	0.99	0.01	0.00	0.97	0.80	0.31	0.01	0.00	0.00	0.99	0.05	0.99	0.16	0.53	0.04	0.99

	angle_TE	slope_axial_pos	slope_tangential_pos	slope_stagger_angle	slope_chord_length	slope_max_camber	slope_pos_max_camber	slope_max_thickness	slope_pos_max_thickness	slope_length_LE	slope_thickness_LE	slope_pos_LE	slope_angle_LE	slope_length_TE	slope_thickness_TE	slope_pos_TE	slope_angle_TE
total_pressure_ratio	0.02	0.42	0.38	0.04	0.75	0.32	0.07	0.05	0.00	0.16	0.06	0.10	0.01	0.27	0.29	0.80	0.00
relative_inlet_angle	0.00	0.02	0.52	0.09	0.83	0.42	0.05	0.08	0.01	0.92	0.14	0.94	0.06	0.07	0.07	0.29	0.01
relative_outlet_angle	0.00	0.15	0.61	0.07	0.80	0.33	0.01	0.08	0.00	0.35	0.08	0.36	0.01	0.01	0.02	0.04	0.01
turning	0.01	0.17	0.60	0.04	0.76	0.27	0.01	0.07	0.00	0.01	0.06	0.05	0.03	0.00	0.01	0.00	0.01
loss_coefficient	0.04	0.83	0.09	0.31	0.97	0.07	0.05	0.00	0.00	0.89	0.00	0.90	0.30	0.33	0.40	0.89	0.01
static_pressure_coefficient	0.12	0.69	0.76	0.09	0.99	0.24	0.01	0.09	0.00	0.30	0.05	0.32	0.00	0.34	0.14	0.81	0.01
total_temperature_ratio	0.00	0.17	0.05	0.42	0.79	0.18	0.02	0.06	0.00	0.83	0.03	0.78	0.07	0.25	0.20	0.95	0.00
adiabatic_efficiency	0.29	0.77	0.18	0.82	0.97	0.05	0.03	0.05	0.03	0.99	0.00	0.87	0.11	0.33	0.26	0.83	0.03
inlet_mass_flow	0.04	0.49	0.36	0.02	0.66	0.30	0.07	0.04	0.01	0.73	0.08	0.75	0.00	0.26	0.31	0.52	0.00

Figure 6.22: CoI of deteriorated geometrical NACA like parameter on their corresponding performance QoI.

it can be seen in figure 6.22 that the CoI of the deteriorated geometrical parameters and QoI are very similar to those shown in figure 6.19. This is likely due to the fact that section 6.3.3 showed that deterioration has a relatively small impact on performance, compared to manufacturing.

However, it can be seen that the CoI shown in figure 6.22 are generally stronger than those show in figure 6.19. Note that the deteriorated blades did not converge geometrically to the measured geometries during the CFD preprocessing. Therefore the enhanced CoI could be a remnant of the preprocessing.

Figure 6.23 shows the CoI calculated based on the delta geometrical parameters and delta QoI. In this case the deltas are calculated by subtracting the manufacturing values from the deteriorated ones. It can be seen that in this case virtually all the parameters have a high influence on all the performance QoI. However, due to the small changes in the geometrical parameters and performance QoI due to deterioration, the high CoI are likely due to the noise.

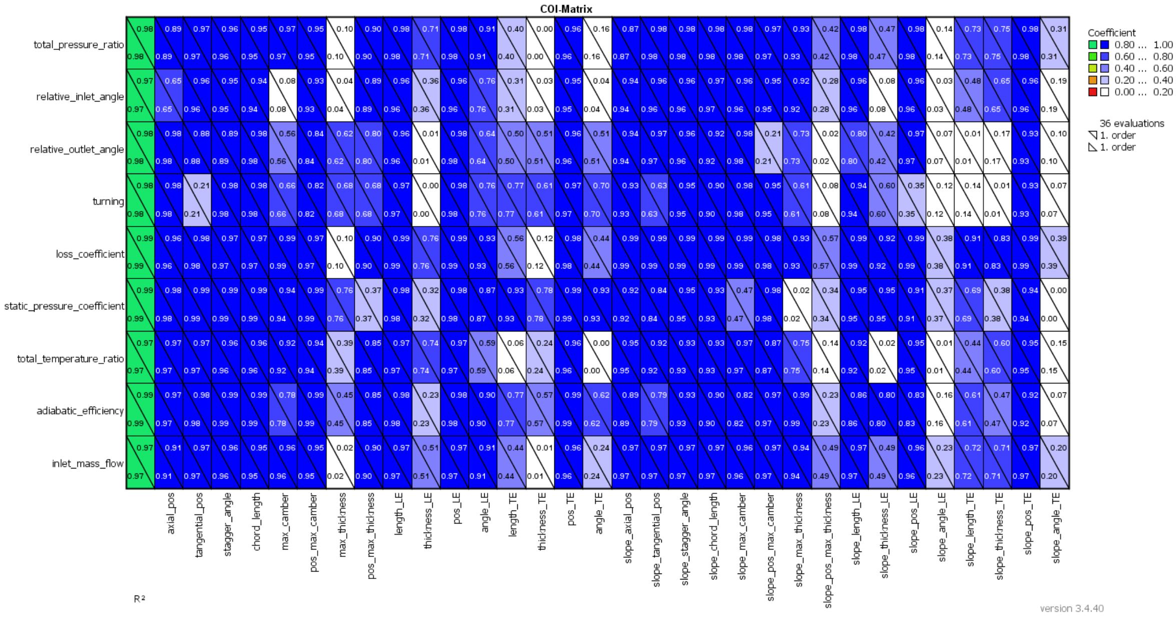

Figure 6.23: CoI of deteriorated geometrical NACA like delta parameters on their corresponding delta performance QoI.

As all the geometrical parameters appear to have high CoI on all the performance QoI the ordering of the parameters per their CoI, as seen in figure 6.21, is omitted here.

6.5. Lessons learned in performance quantification of manufactured and deteriorated blades

The aim of this chapter was to compare the performance scatter due to manufacturing variability and deterioration. Section 6.3 showed that manufacturing variability has a larger impact on performance than deterioration. However, also the impact of manufacturing variability on performance was relatively small.

Furthermore, section 6.1 showed that the geometrical variability due to deterioration is roughly the same order of magnitude than the error caused by the geometry preprocessing. Therefore, the current process cannot accurately quantify the impact of the observed magnitude of deterioration, as seen in chapter 5.

Section 6.4 showed that different QoI are influenced by different geometrical parameters. For the selected QoI stagger angle appeared to be the most important parameter.

A separation of the observed behaviours into the contributions from the axial root and aerofoil were not performed. For the case of manufacturing the separation could not be done as the blades cannot be aligned such that these effects could be isolated. For the case of deteriorated blades the contributions were not separated as the overall performance impact was too small and the impact of deterioration on performance could not be evaluated accurately.

As the measurements required to calculate the tip gaps were not available for this assessment, they have been left out. In future work these could be included is the measurements of other parts that make up the tip gap are provided. The tip gaps are likely to be a geometrical feature that reduces performance significantly.

It is not necessary that all compressor blade designs would experience manufacturing variability and deterioration the same way, or that the geometric variability observed in chapter 5 would cause the same performance scatter for them. The evaluation of the impact of geometrical variability on performance should therefore be performed for other compressor blade designs as well. This would help show the relationship between the design parameter values and their geometrical variability due to manufacturing and deterioration.

A blade set available for this is the R9 blade set, which comes from the same measurement campaign as the R3 blades considered in this thesis. However, care should be taken to account for different mechanisms that deteriorate the blade. For example, the R9 blade analysis cannot be directly compared with the R3 analysis, as the R9 blades rub their height away, while the R3 rub away the liner. The R9 blades are therefore expected to have more variable stagger angle along the span. Analysing similar blades across different engine types should help avoid these considerations.

7

Simulation of manufacturing and deterioration effects for proposed blade designs

The aim of this chapter is to show how to simulate the effects of manufacturing variability and deterioration in order to support the robust optimisation of compressor blades.

The impacts of manufacturing and deterioration can be simulated by applying their parameter difference effects on a set of proposed design parameters. For this approach the crucial step is to build the geometry to be evaluated by the CFD. Section 7.1 shows how the blades can be rebuilt based on a set of parameters.

The subsequent step is to determine how to model the effects of manufacturing and deterioration in therms of delta parameters. Furthermore, any correlation between the manufacturing and deterioration effects must be taken into account. This is discussed in section 7.2.

Manufacturing and deterioration have 3D effects need to be taken into account while modelling their impact on the geometry. The individual NACA like parameters will have their own PDFs that have to be captured properly. This results in a large number of uncertain parameters to be used in the optimisation, which is impractical. Section 7.3 discusses the reduction of parameters and the subsequent convergence of the blade rebuild.

7.1. Rebuild process of blades

The aim of this section is to give an overview of how the NACA like geometrical parameters can be used to create a full blade. The methodology described is based on the one implemented within B2P based on Lange et. al. [35] and the B2P user manual [62].

That methodology combines some user defined baseline geometry (BLM) and the desired delta parameters of that geometry to rebuild a blade. Subsection 7.1.1 presents how the delta parameters are applied to a BLM.

It is important to have a BLM such that the applied delta parameters are as small as possible. Therefore the creation of a BLM based on the measured blades is discussed in subsection 7.1.2. This is advantageous as the major spanwise behaviour observed in chapter 5 can be contained within the BLM.

Before the delta parameters that represent manufacturing and deterioration variability can be applied to a new proposed design, the BLM for this design must be provided. In this case, the measured blades on which to perform the calculation are not available. The method to arrive at the BLM in that case is discussed in subsection 7.1.3.

7.1.1. Application of delta parameters

The delta parameters update the BLM which consists of section parameters and their corresponding chord, camber and thickness distributions. Therefore the parameter update is performed in two steps. First the parameters are updated with the delta parameters. Then, the distributions are updated to reflect the new updated parameters.

Application of delta parameters to parameters

The delta parameters supplied by the user are applied within B2P. It allows the user to specify the parameters for different averaging sections in which the parameters should be applied. This allows a reduction in the number of delta parameters to be used in the blade rebuild.

The delta parameters are applied within their prescribed bounds, as per equation 7.1 [62]. In equation 7.1 P^i are the updated parameters, P^{BLM} are the parameters of the BLM and $\Delta\bar{P}$ are the applied average delta parameters.

$$P^i = P^{BLM} + \Delta\bar{P} \tag{7.1}$$

Note that this might lead to cases where a sudden increase in the parameter would be observed. For example, increasing the axial position by different amounts for two spanwise halves would result in a leading edge with a slanted step between the two sections that belong to different halves. An example of this is shown in figure 7.1. This figure shows the design BLM as a thick black line, with the sections at which it is discretised denoted with the circular markers. The delta parameters are plotted with a thick blue line, and are -1 for first half of the span and $+1$ for the second half of the span. It can be seen that the delta parameters are discontinuous at the mid-span position. This occurs as they are specified as averages within their individual averaging sections and are therefore generally not C^0 smooth.

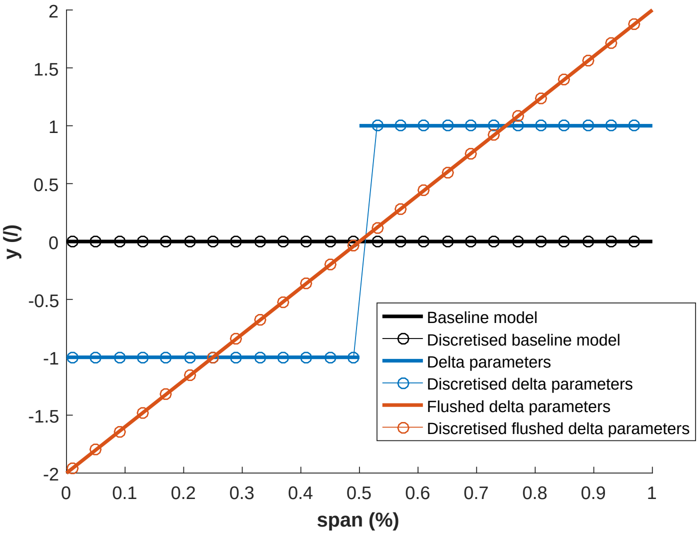

Figure 7.1: Spanwise flushing of delta parameters.

To account for this the delta parameters are flushed along the spanwise positions such that they are C^0 smooth. This is done by applying the prescribed delta parameter value the centres of their respective averaging sections, and interpolating the delta parameter values of other sections in between. Not that this might cause that the average of the parameters in an averaging section of the updated geometry might not match the delta parameter assigned to it. An example of this would be a delta parameter set that is -1 for the first and last thirds of the span and $+1$ for the middle third. The mid-point of the middle third would therefore have a delta parameter of 1 with the delta parameters decreasing towards both the tip and the hub. This means that the average of the middle third would be less than 1.

Application of delta parameters to distributions

After the delta parameters have been applied to the BLM parameters, the morphing vectors for the distributions are calculated. An example of such a morphing vector is shown in figure 7.2, as taken from the B2P user guide [62]. This figure shows two distributions. The lower black line denotes the distribution of a BLM section, whereas the upper black line denotes the target distribution. The change of the peak is shown by the red morphing vector, whose components are the delta position of maximum thickness and the delta maximum thickness.

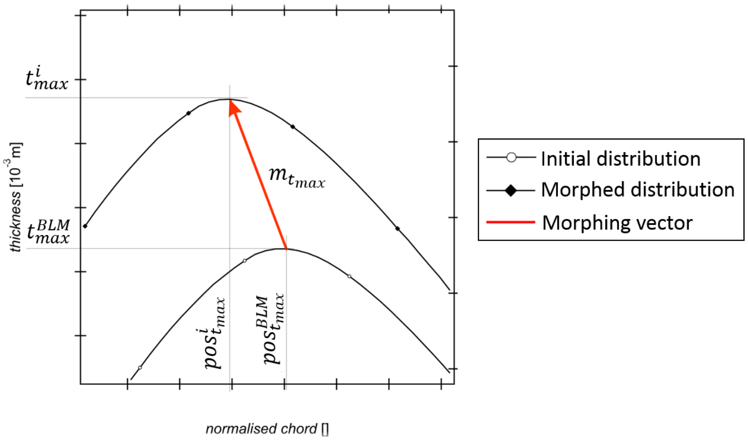

Figure 7.2: Morphing vector for the peak of the thickness distribution [62].

Figure 7.2 shows the change for only one distribution point. The morphing vectors are also calculated for the LE and the TE points. These points correspond to the points for which the LE and TE positions and thicknesses are calculated in the parametrisation of the blade . They are shown in figure 7.3 drawn in yellow and green respectively. The morphing vectors are sufficient to adjust these points, however no information is present in the delta parameters for the distribution points in between.

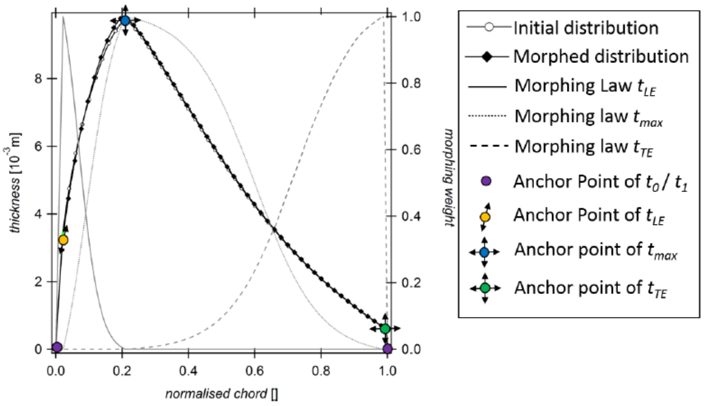

Figure 7.3: The morphing points of the thickness distributions and the corresponding weighting functions [62].

For these points a combination of the morphing vectors is therefore applied. The combinations are performed using weighing functions, which determine the impact of the specific morphing vectors on a particular distribution point. These are drawn as the gray full, dotted and dashed lines in figure 7.3. These lines show that the LE morphing vector has an impact on the points up until the maximum position, and that the TE morphing vector has impact beyond the maximum point. The morphing vector of the maximum point has an impact over the entire chord. The morphed distribution can be expressed as [62]:

$$t^i_j(l_{ch}) = t^{BLM}_j(l_{ch}) + \sum_{i=1}^{m} m_{t,i}\omega_{t,i}(l_{ch}) \tag{7.2}$$

In equation 7.2 t_j^i is the morphed distribution based on the new set of parameters, l_{ch} is a position along the chord, t_j^{BLM} is the distribution of the base geometry, $m_{t,i}$ represent the morphing vectors and $\omega_{t,i}$ are the weighting functions.

7.1.2. Creation of the BLM based on measured scans

For the rebuild process the BLM needs to be supplied by the user. It needs to be supplied in the form of four separate files containing the spanwise parameters, chord, camber and thickness distributions. This geometry should be such that the deviations introduced to it by the delta parameters are minimal. Therefore the BLM should be based on the measured blades.

Specifically the BLM could be created as a "mean" or "median" blade of the measured blades. That would ensure that the average delta parameters applied to this geometry will be minimised. In this subsection the approach will be discussed on the example of creating a median blade, but the approach is the same for the mean blade as well.

The creation of the median BLM blade is based on the results of the measured blade parametrisation. These results include the spanwise section parameters and their corresponding chord, camber and thickness distributions. As discussed in section 5.2.1 most of the spanwise section NACA like parameters are calculated directly from these distributions. The positioning parameters such as the axial and tangential positions, chord length and stagger angle are calculated separately.

The calculation of the median positioning parameters and the stagger angle is straight forward. The values of all the blades are collected from their separate files and their median calculated. This median is calculated for each of the spanwise sections separately, as the median BLM geometry needs to be created for all sections.

The calculation of the parameters based on the distributions is more involved. Simply calculating the median parameters might lead to a situation where the parameters do not agree with the distributions they are intended to describe. This leads to poor rebuild quality results. Instead, the median distributions can first be calculated, after which the parameters can be calculated from them exactly as they have been in the parametrisation process. These are then the median aerofoil section distributions and parameters.

To do so the median distributions for each of the spanwise sections are first calculated. This is done by pairing the values of chord distributions with the, for example, camber distribution values of a particular spanwise section of a particular blade. These can then be imagined as x and y coordinates respectively. These are then collected for all the scanned parameters, and for each x location the corresponding median y value is calculated.

Note that the chord distributions of different sections or of different blades are not necessarily the same. Therefore the values cannot be averaged directly. Instead a new set of chord positions can be selected along the chord forming a set of new x coordinates, x_n. The original x and y coordinates are then interpolated at locations x_n to give the values y_n for each of the blades. The values of y_n are then averaged at each of the positions x_n to give the average distribution.

Based on the median distribution the median aerofoil section parameters are calculated. The calculation procedure for those haas been discussed in section 5.2.1.

Note that in the case of the measured blades the BLM calculated with the methodology presented will contain the average spanwise behaviours as seen in chapter 5. Therefore only the blade-to-blade differences need to be simulated.

7.1.3. Creation of the BLM for a proposed design

A new proposed design is a new nominal CAD design. As such it does not contain the inherent differences of the camber and thickness distributions due to manufacturing and deterioration. Therefore, before this CAD design can be updated with the delta parameters that simulate the effects of manufacturing and deterioration, its distributions, and therefore also parameters, need to be updated for the inherent differences as well. The updated geometry is the BLM of the new proposed design.

For the case of simulating the effects of manufacturing and deterioration on the nominal CAD design this update has been done as the calculation of the median parameters. In the case of the design proposed in the optimisation, measured blades are not available. Therefore the differences of the camber and thickness distributions are assumed to be the same as those observed for the nominal CAD design, and applied to the proposed design in order to obtain the proposed BLM.

The calculation of the distribution differences is similar to the process of calculating the average distributions described in subsection 7.1.2. First the methodology described in subsection 7.1.2 is followed in order to obtain the median or mean camber and thickness distributions. Then the distributions of the nominal CAD design are subtracted from these average distributions, which gives the differences of the distributions. Lastly, these differences are applied to the new proposed design distributions. Note that also in this case the chord distributions are not necessarily at the same, and therefore interpolation of the distributions at new positions is required, as described in subsection 7.1.2.

The parameters of the base geometry need to reflect the distributions as well. Therefore the parameters of the proposed design must also be updated by the differences of the parameters. The procedure is analogous to the calculation of the distribution differences. First the average parameters of the measured blades are calculated and the parameters of the nominal CAD design are subtracted to give the parameter differences. The differences are then applied to the positioning parameters. The parameters that characterise the distributions are recalculated from the updated distributions. This procedure calculates the parameter differences for all the spanwise sections individually.

Note that this process assumes that the new proposed design blades will experience the same effects of manufacturing and deterioration as the measured blades. For small differences between the proposed and the nominal designs this is a reasonable assumption, however for proposed designs that are drastically different this assumption might not hold.

7.2. Modelling manufacturing and deterioration effects

The aim of this section is to present the modelling of the manufacturing and deterioration effects of the aerofoil section parameters. Subsection 7.2.1 discusses the basic approach of modelling the delta parameters with Gaussian distributions, as well as setting up the delta parameters for various numbers of spanwise sections.

Subsection 7.2.2 discusses the correlation between the effects of manufacturing and deterioration on the individual aerofoil section parameters. These need to be taken into account when simulating new delta parameters.

Subsection 7.2.3 also discusses the correlation of effects of manufacturing and deterioration. However, unlike subsection 7.2.2, it discusses the so called cross-parameter effects of manufacturing and deterioration. Essentially, these are the correlations between the effects of a particular parameter and the effects of other parameters.

7.2.1. Modelling of delta parameters

The aim of this section is to present how the manufacturing and deterioration effects are modelled with the use of delta parameters. As discussed in subsection 7.1.2, only the blade-to-blade differences need to be modelled, while the average spanwise behaviour is contained within the base geometry.

Section 5.6.1 showed that most of the blade-to-blade deviations due to manufacturing is due to simple blade-to-blade offsets, however other spanwise behaviours were also observed. These included linear, concave and piecewise linear behaviour.

As described in section 7.1.1, the delta parameters are flushed along their respective averaging sections. Therefore using 2 averaging sections allows the recreation of the linear behaviours. The use of 3 averaging positions can be enough to capture piecewise linear behaviour, provided that there are only two linear behaviours combined together. It could also simulate the concave behaviour to a degree as a piecewise linear approximation.

Within B2P the number of averaging sections needs to be the same for all the parameters. However, different numbers of averaging sections can be used when setting up the delta parameters file, which is the input for B2P. In that case the number of averaging sections in this file would be determined by the parameter that requires the most averaging sections to describe its spanwise behaviour. For that parameter the actual delta parameters in all the averaging sections is calculated. The delta parameters of other parameters can be calculated only in their required number of averaging sections, and those numbers interpolated to obtained the values that would correspond to the actual averaging sections.

This only leaves the issue of simulating the delta parameters in their respective averaging sections. For the case of manufacturing section 5.6.1 showed that the mean and median lines lay on top of each other, which means that the distributions of parameters and their corresponding PDF at the individual spanwise positions are not skewed. The measured set of blades is relatively small, therefore the distribution fitting depends on the individual measurements to a large extent. Note that a set of few values sampled from a normal distribution might not be normally distributed. Because of this, it is only reasonable to assume that the distribution of the values is a Gaussian distribution. If the measured set were larger, other probability distributions could be fitted in an attempt to determine the underlying PDF.

Therefore each delta parameter is being modelled by a dedicated Gaussian PDF. It is possible however that the delta parameters are correlated. This needs to be taken into account when sampling these PDF in order to obtain delta parameters that simulate manufacturing or deterioration. For example, the LE thickness could be correlated with its position. This could occur if the LE were manufactured larger than intended while the overall dimensions of the blade remain the same. On the other hand it might be correlated to chord length, as larger blades could mean larger LE.

For the case of simulating the effect of manufacturing the delta parameters can be sampled directly from the multivariate PDF constructed on the basis of the individual delta parameter PDFs and their covariance matrix. This can be done as the manufacturing does not follow after any previous uncertain process. This is obviously not true for the case of deterioration, as it follows after manufacturing.

7.2.2. Spanwise correlation of manufactured and deteriorated delta parameters

For the case of deterioration any correlations of the deterioration differences with the manufacturing differences need to be taken into account when sampling the delta parameters that simulate deterioration. In other words, it is important to know, and to take into account, whether there is any relationship between the value of a parameter and the amount it deteriorates by.

In order to see how strong these correlations are, they were calculated individually for each of the parameters. The correlations are based on the differences caused by manufacturing and deterioration, as those are the quantities used to simulate these effects. The span was subdivided into 7 different averaging sections. In each of these averaging sections the parameters of the individual blades are averaged. This results in 36 averaged values for each of the 7 averaging sections. This is done for both the manufacturing and deterioration differences. The correlation is then performed pairwise for the different averaging sections of the manufacturing and deterioration differences. This results in a $7x7$ correlation matrix. The correlation is calculated as the Spearman rank correlation coefficient, as per Spearman [58]. The entire procedure is similar to that of Heinze et. al. [21].

Figure 7.4 shows the created rank correlation matrix for axial position. This figure shows the Spearman rank correlation coefficient matrix with the rows and columns corresponding to the deteriorated and manufacturing averaging sections respectively.

The Spearman rank correlation coefficients range from -1 to 1, with the former indicating strong monotonically decreasing trend, and the latter a strong monotonically increasing trend. Values around 0 indicate no correlation. Figure 7.4 shows that all the rank correlation coefficients of axial position are close to 0. This means that the blade-to-blade differences due to manufacturing of the axial position do not have an impact on the blade-to-blade difference of axial position due to deterioration.

This is true for the parameters that have not exhibited much noise in their values in chapter 5. Some of the correlations of the averaged differences of these parameters are strong. The rank correlations are strongest for the case of the TE angle. The rank correlations of this parameter are shown in figure 7.5.

Figure 7.5 shows that the rank correlation coefficients on the diagonal are strong, whereas the off diagonal coefficients are weak. This indicates that the average TE angle blade to blade differences of the manufactured

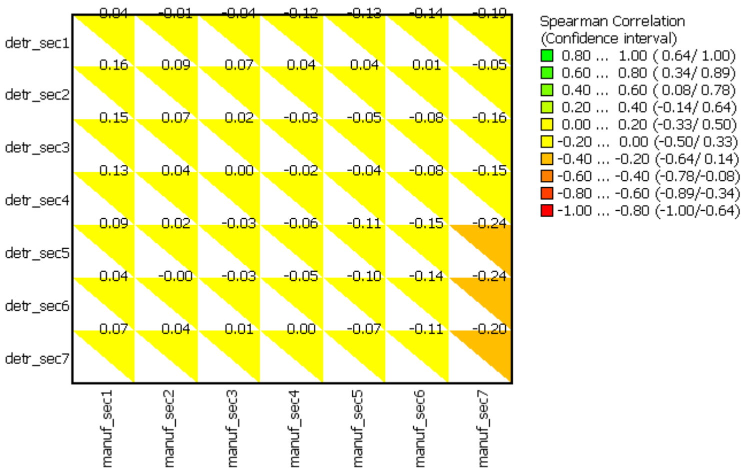

Figure 7.4: Spearman pairwise rank correlation coefficients of section averaged manufacturing and deterioration differences of axial position.

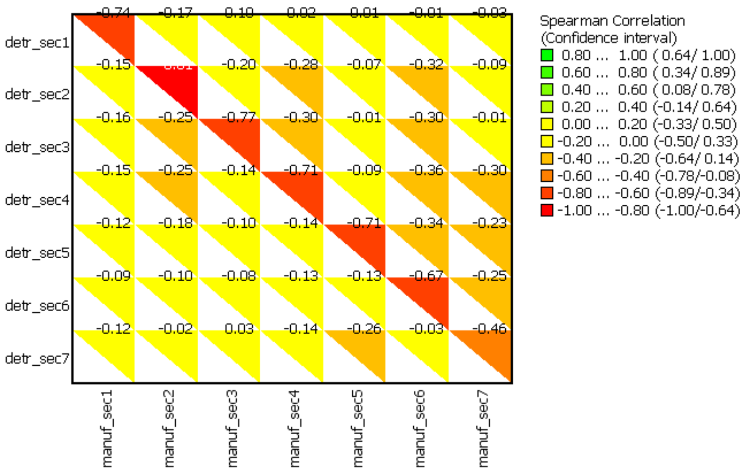

Figure 7.5: Spearman pairwise rank correlation coefficients of section averaged manufacturing and deterioration differences of TE angle.

blades directly affect the average TE angle blade-to-blade differences of the deteriorated blades of the same averaging section. However, they have no impact across different averaging sections.

Therefore, the differences are only correlated locally. This could be the case if all the sections were deteriorating independently, with the deterioration causing significant blade-to-blade differences in the TE angle. However, this does not coincide with the data shown in figure 5.55, where no blade-to-blade behaviour could be isolated.

A more reasonable explanation is therefore that the correlations are caused by different levels of noise, due to the sensitivity of the calculation method. The levels of noise in the actual values of the parameters it can be seen, that the noise in the values of manufactured TE angles is higher that the noise of the deteriorated TE angles. Therefore, when subtracting them to obtain the differences due to manufacturing and deterioration the noise of the former dominates the noise of the latter. Therefore, the correlation between the differences due to manufacturing and deteriroation are a consequence of the different noise levels.

The same conclusions can be drawn for other parameter correlations as well. Note that the other correlation matrices feature rank correlation coefficients that are weaker than the ones seen in figure 7.5. It can therefore be concluded that there is no direct correlation of manufactured parameters to the deteriorated ones.

7.2.3. Cross parameter correlation of manufactured and deteriorated delta parameters

It is not necessary that the parameter value in the manufactured state is the only indicator of how this parameter will deteriorate. It is possible that there could be cross parameter effects. To assess this, the correlations between the aerofoil section parameters are plotted in figure 7.6. This figure shows the pairwise correlations of the manufacturing and deterioration parameter differences, calculated for the individual 7 spanwise averaging sections. The manufacturing differences are the differences of the measured blades to the nominal CAD design, and the deterioration differences are the differences between the manufactured and deteriorated measurement pairs.

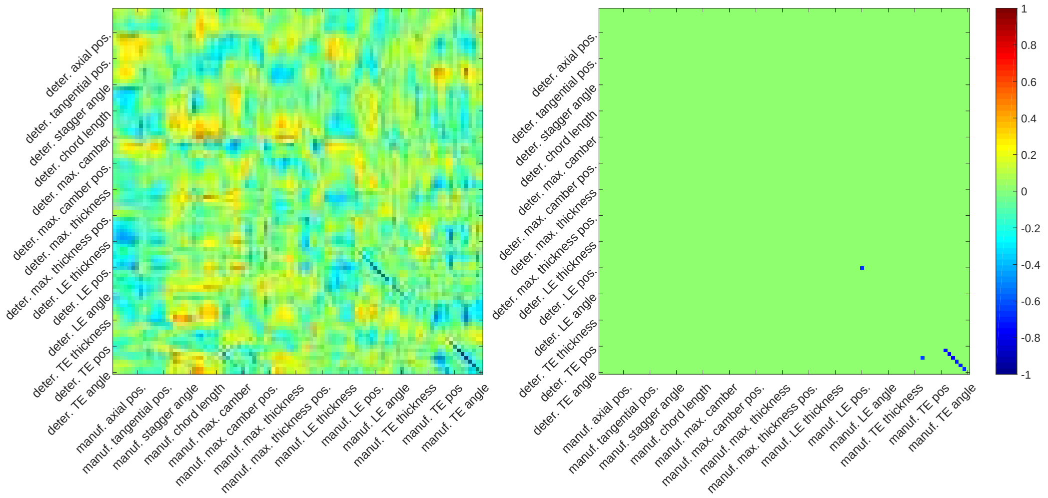

Figure 7.6: Spearman pairwise rank correlation coefficients of section averaged manufacturing and deterioration differences of all parameters.

Figure 7.6 contains two panels. In the left panel the values of the correlations are displayed in color, from red being +1 and blue being −1. In the right panel, the color display is the same, but with all correlation coefficients less than 0.5 and higher than −0.5 replaced with 0. This isolates the correlation coefficients that imply possible correlation. The deteriorated parameters are on the y axis, whereas the manufactured parameters are on the x axis. As the correlations are done for 7 averaging sections, the correlation matrix of manufacturing and deteriorated differences is of the size 7×7. These are located on the diagonal of the left panel, and correspond to the matrices discussed in subsection 7.2.2.

The left panel of figure 7.6 shows that there are regions of higher and lower correlations. For example, between the deteriorated tangential position and manufactured axial position differences. However, the right plot shows that these are all smaller than 0.5. In fact, only a few correlations higher than 0.5 are present, with most belonging to the TE angle. The reson for this has already been discussed in subsection 7.2.2.

The other relatively high correlation coefficients are the between the manufactured and deteriorated position of the LE, and between the TE angle and the TE position. As seen in subsection 5.6, the LE and TE positions, and the TE angle are all parameters with relatively high amounts of noise. Therefore, similarly to the reason of the TE angle correlations, these correlations are due to varying amounts of noise in the manufactured and deteriorated parameters.

7.3. Convergence of blade rebuild

The aim of this section is to show the convergence of the rebuild blades updated with the effects of manufacturing and deterioration. The aim is to achieve good convergence with the minimum amount of delta parameters. The discussion of this section is based on the measured blades, and the convergences are based on attempts to recreate the measured geometries. This is because the effects of deterioration were to small to accurately capture using the current rebuild method, as discussed in section 6.1.3..

The number of parameters obtained during parametrisation is relatively high, which makes simulating all the parameters intractable in a probabilistic simulation. Subsection 7.3.1 shows the reduction of the number of these parameters.

The rebuilt blades are intended to simulate the manufactured and deteriorated blades. Section 7.3.4 shows the convergence of the rebuilt process to the intended shape by comparing the measured parameters to the parameters created by the rebuild process.

Subsection 7.3.3 shows the convergence of the surface deviation maps based on blades rebuilt using the reduced number of uncertain parameters. The parametrisation of the blades is essentially geometry sampling at a few locations. Because of this, information on the blade deviations are inherently lost. Furthermore, the reduction of parameters also inherently causes a loss of information. This section shows how well the surface deviations converge to the measured surface deviations.

Essentially the interest lies with the performance of the evaluated geometry as opposed to its shape. Therefore the performance of rebuilt blades needs to converge to the performance of the measured blades. Subsection 7.3.6 shows the performance convergence of the performance QoI.

Lastly, the manufacturing and deterioration effects are simulated by generating a number of blade geometries which are subsequently evaluated by CFD. As CFD is computationally costly it is advantageous to keep the number of geometries that need to be evaluated to a minimum. Subsection 7.3.7 shows the convergence of the performance QoI with the sample size.

The comparisons made in this section all centre on the comparison of measured and rebuilt geometries. The rebuild process aims to rebuild the measured blades, and requires the delta parameters and the BLM to be supplied. Therefore, for the creation of rebuilt blades, a BLM has been created as a median of the measured blades as described in subsection 7.1.2. The delta parameters between the median and the measured blades were then calculated and used in order to rebuild the original measured geometries. The difference between the measured and the rebuild geometries is the subject of the convergence study.

7.3.1. Reduction of parameters

The aim of this subsection is to determine the rebuild approaches to be used in the subsequent convergence study. As seen in the discussion of the effects of manufacturing and deterioration on the spanwise parameters in chapter 5, most of the blade-to-blade differences are either spanwise constant offsets or linear functions of the span, with some exceptions featuring behaviours that were piecewise linear or could be well approximated by piecewise linear behaviour.

For the constant offsets a rebuild using delta parameters averaged along the entire span is sufficient to capture the blade-to-blade differences. For the linear behaviour two spanwise delta parameters are needed. More are needed for the piecewise linear behaviour.

It is not necessary that the same number of delta parameters are used in order to simulate the blade-to-blade differences. The spanwise behaviour of some parameters can be simulated by many delta parameters, whereas others can be simulated by a single delta parameter.

In this case chapter 5 has shown that the parameters that feature most of the linear behaviour are the axial and tangential positions, stagger angle and the maximum camber. Therefore an approach using two delta parameters to describe the blade-to-blade behaviour of these parameters, but using one delta parameter for all the rest can be used. This approach is called the hybrid approach hereafter.

In order to take into account any piecewise linear behaviour at least 3 delta parameters need to be used. This allows to model any behaviour that can be approximated with three straight lines. However, as seen in figure 5.46, the position of maximum camber has piecewise linear behavour only in a small part of the span, which could be only crudely captured by 3 delta parameters. However, it can be seen that 4 delta parameters should be able to approximate the observed behaviour well.

Therefore the approaches considered for the convergence study will use the delta parameters averaged across the entire span, of 4 different averaging sections and a hybrid approach which models the axial and tangential positions, the stagger angle and the maximum camber with two delta parameters. These approaches are hereafter named the 1 section, 4 section and hybrid approaches. The rebuild is performed by several numbers of delta parameters for equidistant spanwise averaging sections.

7.3.2. Selection of the BLM geometry

As mentioned in section 7.1.2, the BLM is supplied by the user and should resemble the rebuilt blades as much as possible. In the case of rebuilding the measured blades a suitable BLM must be selected based on the magnitude of the differences between the individual blades and the BLM. If the differences are small, the nominal CAD geometry can be used, otherwise a mean or median BLM must be constructed as per 7.1.2.

Figure 7.7 shows the difference between the median BLM and the nominal CAD design of the measured manufactured blades considered. The mean BLM is not shown as it is virtually exactly the same as the median BLM. Figure 7.7 shows the camber distributions in the left panel and the thickness distributions in the right panel. The comparisons are only done for the midspan region, but the findings apply to most other spanwise sections as well.

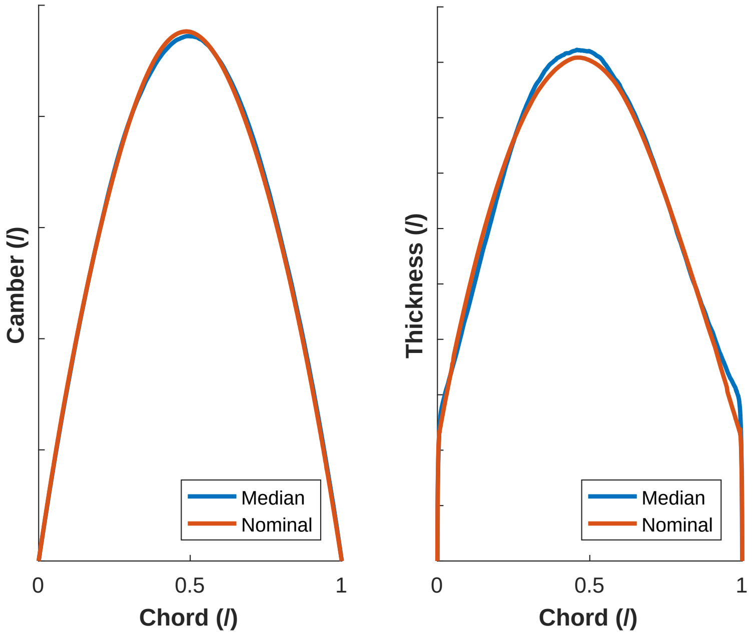

Figure 7.7: Comparison of the median BLM and the nominal CAD design midspan camber (left) and thickness (right) distributions.

In the left panel it can be seen that the nominal and median BLM camber distributions are similar along the majority of the chord. The maximum difference occurs near the mid-chord position, where the peak of the camber distribution is. Here it can be seen that the median BLM camber distribution peaks slightly more towards the TE than the nominal CAD design camber distribution. Furthermore, the median BLM camber distribution has a lower peak.

In the right panel the differences in the thickness distributions can be seen. The differences in these distributions are visible along the majority of the chord, but are most pronounced at the TE, where the median BLM if thicker than the nominal CAD design. Furthermore, the median BLM is also thicker at the mid-chord location, and at the LE. It is thinner in a region between the LE and the peak however.

The differences in the thickness distributions cannot be captured by morphing the nominal distribution, as the differences are local, while the morphing distributes the impact of the delta parameters along the entire chord. Therefore the nominal CAD design cannot be used, and the median BLM model has been chosen as a basis for the rebuild of the measured manufactured blades.

7.3.3. Convergence of surface deviations

The convergence of the different rebuild processes identified in subsection 7.3.1, can initially be analysed by analysing the surface deviation maps. In this case the rebuilt blades are compared to the measured blades they are intended to represent. The surface deviations between the two show the convergence of the rebuild. These deviations are calculated by subtracting the measured surface deviations from the rebuilt surface deviations.

These can be calculated for all 36 measured blades. The information in the 36 separate surface deviation maps per rebuild approach were then summarised in the mean surface deviation map and the map of standard deviations for individual surface points. The former shows the mean differences between the measured and rebuilt blades, while the latter shows the regions that are most variable when rebuilt. Note that the values presented in the plots have been normalised by an appropriate length scale.

Surface deviations of 1 section approach rebuild

Figure 7.8 shows the mean differences between the measured blades and the blades rebuilt using the 1 section approach. The mean deviations are color coded, with the extremes of the spectrum representing ±0.5. Note that the maximum mean deviations of manufactured blades, as shown in section 5.5.2 were ±1, which is twice as much.

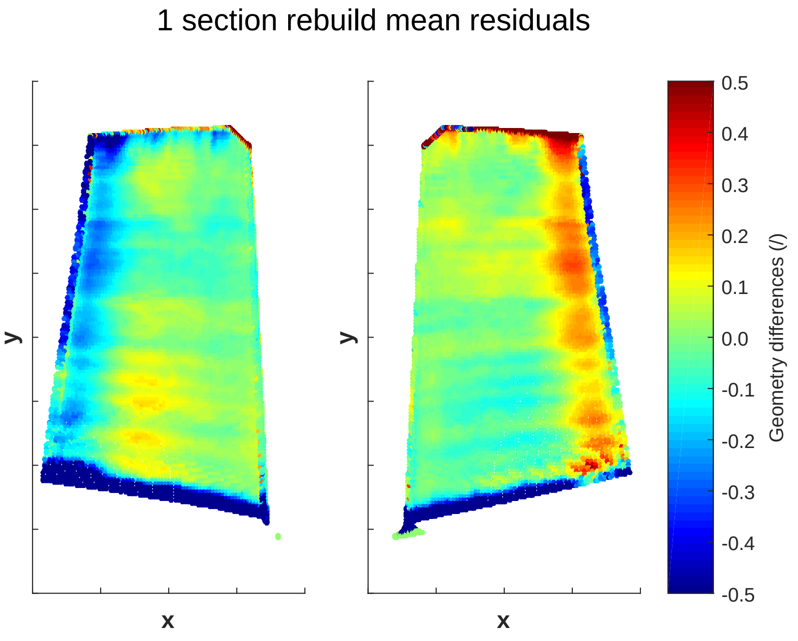

Figure 7.8: The mean differences between the surfaces rebuilt using the 1 section approach and the measured blades for the pressure side (left) and the suction side (right).

It can be seen that the impact of the rebuild is within roughly ±0.15, as most of the blade is coloured in shades between yellow and blue on the spectrum. There are some regions of relatively high differences however, most notably near the TE. Near the TE the differences on the pressure side are negative, while the differences on the suction side are positive. This indicates that the TE the blades are being rebuilt with insufficient curvature. Furthermore, it can be seen that the widths of the regions of high deviations in the x direction are periodic in the y direction. The locations where the widths are smallest are where the aerofoil sections are being rebuilt.

Another feature that can be seen are the streaks of larger deviations along the chord that can be seen both on the pressure and the suction sides. These streaks are located in between the rebuilt sections, and are therefore a remnant of spanwise sampling of the blade geometries.

The regions with negative deviations near the hub are a result of the fillet effects on the blade. As no geometry was sampled in this region the rebuild process cannot account for the fillet effects. The same can be claimed for the region of the tip, as the highest sampled section is below the clipped LE tip geometry. Therefore any special behaviour of the blade above that section cannot be accounted for.

Lastly, it can be seen that the points on the TE are generally coloured blue. This suggests that the rebuild blades are somewhat shorter than the measured blades.

Another aspect of the surface deviation plots is to try and determine which regions are most susceptible to differences as a result of the rebuilding process. This can be assessed by analysing the standard deviations of the differences at each of the geometry surface points. This is shown in figure 7.9.

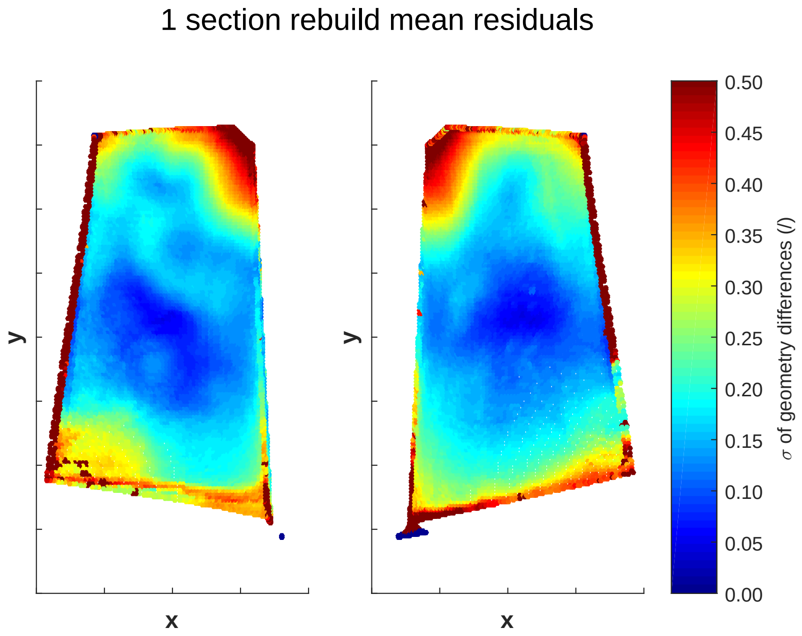

Figure 7.9: The standard deviation of differences between the surfaces rebuilt using the 1 section approach and the measured blades for the pressure side (left) and the suction side (right).

Figure 7.9 shows the standard deviations color coded from blue to red, with blue denoting regions of relatively low standard deviations, and red denoting areas of relatively large standard deviations. This figure shows that the centre of the blade has low standard deviations, which means that the rebuilds are consistent with the measured geometry. However, the tip, hub and LE and TE regions are more uncertain. This could be due to the fact that the mean delta parameters are expected to be similar to the delta parameters of the mid-span sections. This occurs as most of the blade-to-blade behaviour is linear.

Note that the TE is coloured dark red on the left panel of figure 7.9, and coloured red for sections above roughly a third of the span from the hub in the right panel. This shows that the differences at the TE are not very consistent, and therefore it is not necessarily true that the rebuild blades are shorter than the measured blades.

Surface deviations of hybrid approach rebuild

Figure 7.10 shows the mean differences between the measured blades and the blades rebuilt using the hybrid approach. Compared to the deviations of the 1 section approach, as shown in figure 7.8, it can be seen that the hybrid approach converges better to the measured blades. The deviations along the entire span are now roughly within ±0.15.

The chordwise streaks are still visible, however they cannot be completely removed anyway, as they depend on the locations of the spanwise geometry sampling sections. The TE still shows that the rebuilt blades are generally shorter than the measured blades.

Figure 7.11 shows the standard deviations of the differences between the blades rebuilt using the hybrid approach and the measured blades. It can be seen that the standard deviations of the differences are relatively small for virtually the entire blade, compared to the standard deviations of the 1 section approach, as seen in figure 7.9. This means that the rebuild differences are consistent for all the measured geometries being rebuilt.

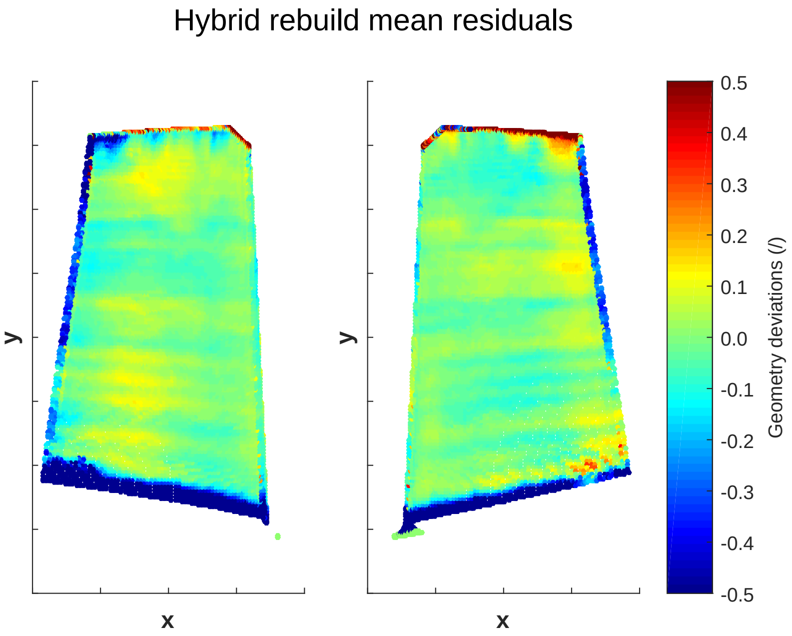

Figure 7.10: The mean differences between the surfaces rebuilt using the hybrid approach and the measured blades for the pressure side (left) and the suction side (right).

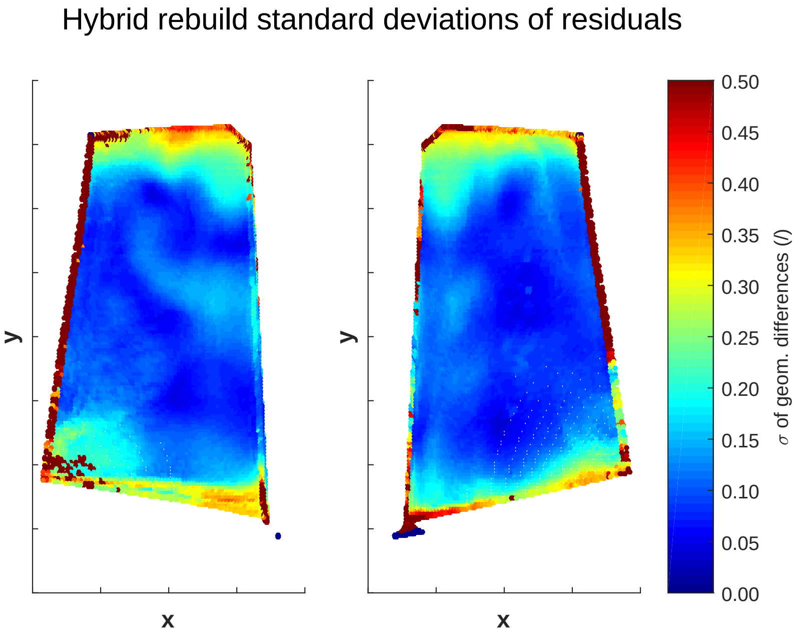

Figure 7.11: The standard deviation of differences between the surfaces rebuilt using the hybrid approach and the measured blades for the pressure side (left) and the suction side (right).

Some areas of higher deviations remain however. These are near the tip, hub and the TE. Furthermore, there are some regions of high deviations on the LE as well, just above the half span. The large TE variability could account for the high mean differences of the TE seen in figure 7.10.

Surface deviations of 4 section approach rebuild

Figure 7.12 shows the mean differences between the measured blades and the blades rebuilt using the 4 section approach. Compared to the deviations of the 1 section approach, as shown in figure 7.8, it can be seen that the hybrid approach converges better to the measured blades in the upper half of the span. However, in the lower half of the span the convergence is worse.

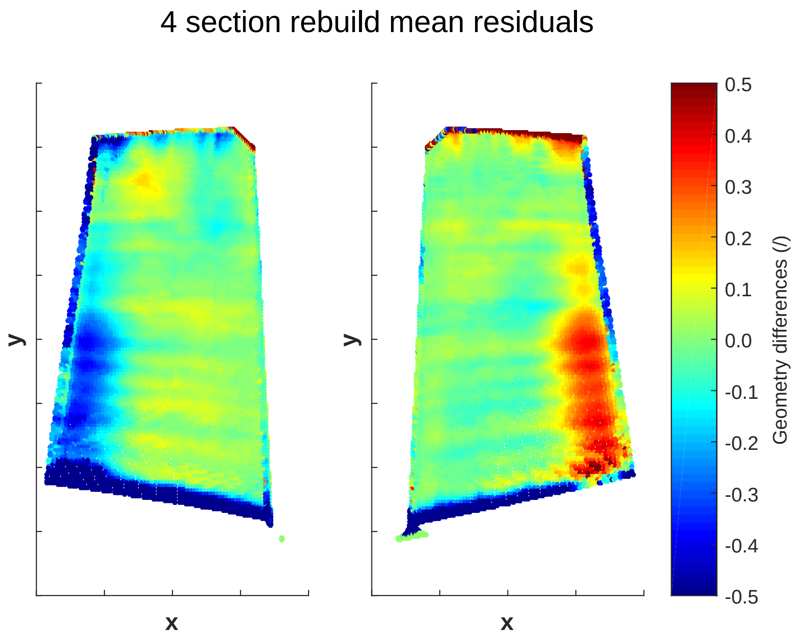

Figure 7.12: The mean differences between the surfaces rebuilt using the 4 section approach and the measured blades for the pressure side (left) and the suction side (right).

The standard deviations of differences in these regions are also relatively large, as seen in figure 7.13. This shows that in these regions the differences between the measured and rebuilt blades are different for individual measured-rebuilt blade pairs. Therefore some blades are rebuilt better than others. However, analysing the individual blade rebuild differences, it can be seen that the large differences in the TE region occur for the vast majority of the measured-rebuilt blade pairs.

These differences could be due to the high levels of noise of the TE angle blade-to-blade differences. It can be seen in figure 5.39 that the blade-to-blade differences of the TE angle are high for the first half of the span, after which they decrease. This coincides with the behaviour of the high surface deviation regions of the 4 section approach rebuilt, seen in figure 7.13.

7.3.4. Convergence of parameters

The aim of this subsection is to demonstrate the convergence of the parameter rebuild using 1 and 4 section and the hybrid approach, identified in subsection 7.3.1.

The measured parameters were obtained by parametrising all the measured blade, and have already been analysed in chapter 5. The rebuilt parameters are obtained by re-parametrising the rebuilt geometries with B2P. This ensures that the parameters in the convergence study are the actual parameters that have been rebuilt.

Essentially the convergence of parameters focuses on the ability of the rebuild process to replicate the blade-to-blade behaviour seen in the chapter 5. Note that the analysis of the parameters in chapter 5 already described the blade-to-blade behaviours after the spanwise constant offsets have been subtracted. This is the level of convergence that a 1 section rebuild can achieve. The successfulness of using more delta parameters can be measured by how much of the remaining blade-to-blade behaviour they manage to capture.

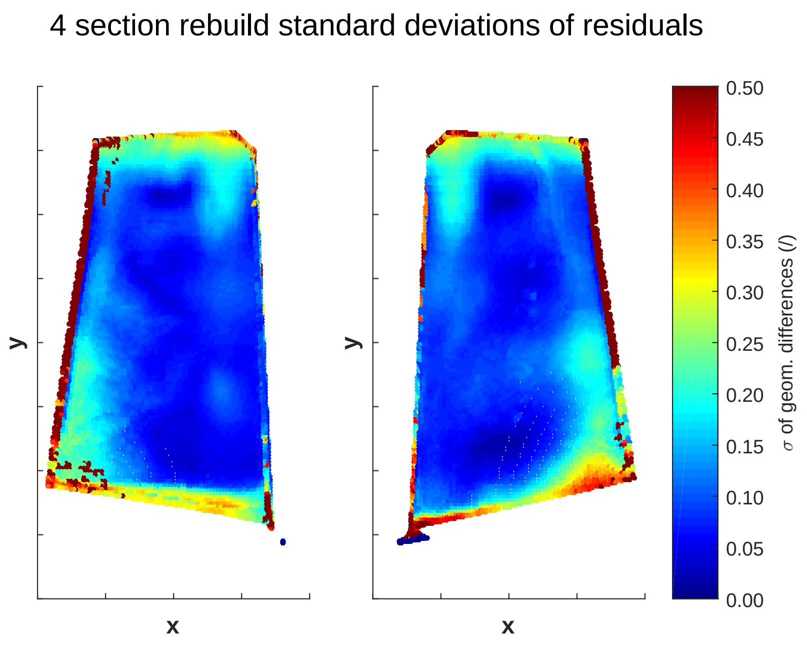

Figure 7.13: The standard deviation of differences between the surfaces rebuilt using the 4 section approach and the measured blades for the pressure side (left) and the suction side (right).

The comparison of the convergence of different rebuilding processes is done using plots similar to those used in chapter 5, and have two panels. The left panel shows the differences between the measured manufactured blades and the nominal design. This is a reference case that gives perspective to the rebuild differences of different rebuild processes. The right panel shows the rebuilt parameters with their corresponding manufactured parameters subtracted, leaving only the blade-to-blade differences lost due to the rebuild.

The comparisons of convergence are given for axial and tangential positions, stagger angle and maximum camber, which showed linear and piecewise linear blade-to-blade behaviour. Other plots do not show significant differences between the proposed rebuilding processes and are omitted here.

Axial position

Figure 7.14 shows the differences of the axial position caused by the rebuild process. In the right panel shows the differences of axial position lost in the rebuild process. It can be seen that the blue lines, corresponding to 1 section rebuild, correspond directly to the blade-to-blade differences seen in figure 5.25.

Furthermore, it can be seen that both the hybrid and the 4 section approaches reduce the loss of blade-to-blade behaviour significantly. This can be seen in the lower half of the span where the scatter of yellow and orange lines is smaller than the scatter of the blue lines. The same is true for all sections above section 30.

However, the hybrid and the 4 section approaches have not been able to capture all of the behaviour. The hybrid approach only uses 2 delta parameters to model the axial position and therefore cannot capture all the behaviour at the tip, which results in the loss of the scatter at the tip. The 4 section approach captures the tip scatter. However, a outlier line belonging to the 4 section group can be seen, showing that in some cases even this approach is not sufficient.

In general the differences between the hybrid and the 4 section approaches are small. Compared to the differences due to manufacturing on the left panel, the amount of blade-to-blade behaviour that is lost both approaches seem to be acceptable in the case of simulating manufacturing blade-to-blade differences.

Furthermore, as there is already a large difference in the number of parameters between the 4 section and the hybrid approach, it is unlikely that increasing the number of parameters beyond 4 will reduce the amount of

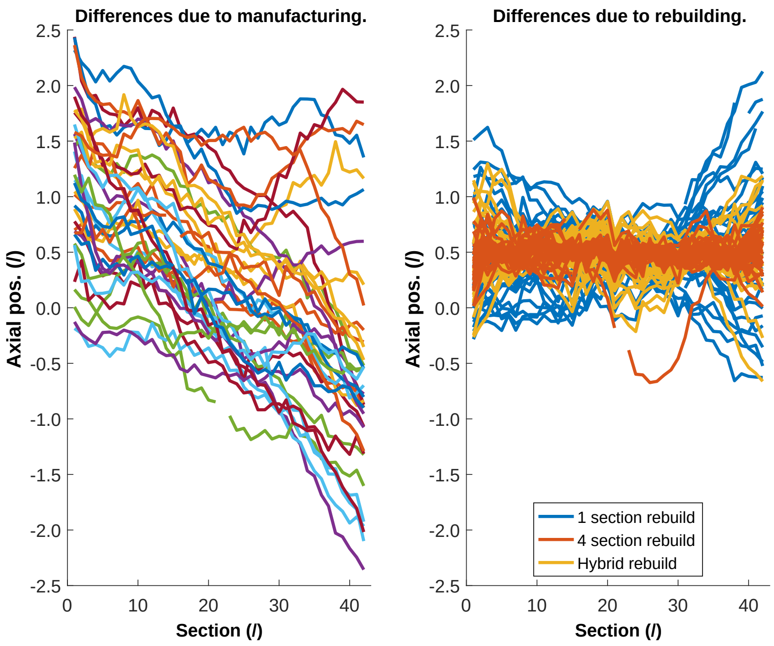

Figure 7.14: The axial position differences due to manufacturing (left), and the axial position blade-to-blade behaviour losses due to the rebuild process (right).

behaviour lost. Therefore the amount of behaviour lost by section 4 is indicative of the limits of the rebuilding process.

Lastly, as the more complex 4 section approach is not decisively better than the hybrid approach it is reasonable to accept the loss of scatter at the tip in exchange for a drastically reduced number of uncertain parameters. Therefore, in the case of the axial position rebuild the hybrid approach is more advantageous.

Tangential position

The loss of blade-to-blade behaviour of tangential position due to rebuild process can be seen in figure 7.15. Note that this figure shows the same differences between the rebuilding process results as figure 7.14, and therefore supports the same conclusions.

It can be seen that both the 4 section and hybrid approaches reduce the amount of behaviour lost significantly. Furthermore the differences between them are rather small. Comparing the behaviour lost due to the rebuild to the differences due to manufacturing, it can be seen that it is relatively small. Lastly, it can be seen that the increase of delta parameters did not significantly reduce the amount of behaviour lost, therefore the hybrid approach is advantageous for the rebuild of tangential positions.

Stagger angle

The loss of blade-to-blade behaviour of stagger angle due to rebuild process can be seen in figure 7.16. Note that this figure shows similar differences between the rebuilding process results as figure 7.14.

It can be seen that both the 4 section and hybrid approaches reduce the amount of behaviour lost in the sections near the tip significantly. However, near the hub the reduction of the behaviour lost is small.

The differences between the result of the 4 section and hybrid approaches are generally small. However, there is a large difference near the hub, where the yellow lines corresponding to the hybrid approach are very scattered. This means that it is unable to accurately simulate the impact of manufacturing in this region. However, it can be seen that in this region all of the proposed rebuilding processes are not capable of accurately simulating the blade-to-blade behaviour observed.

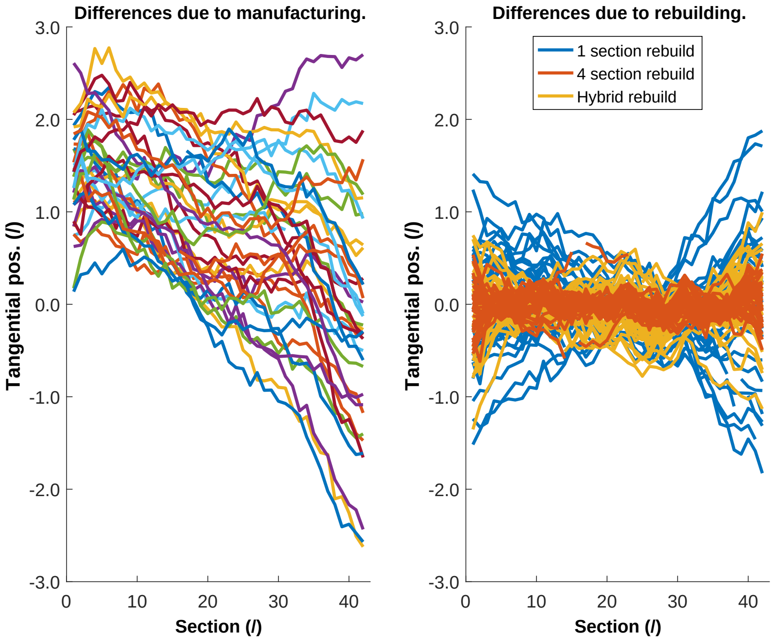

Figure 7.15: The tangential position differences due to manufacturing (left), and the tangential position blade-to-blade behaviour losses due to the rebuild process (right).

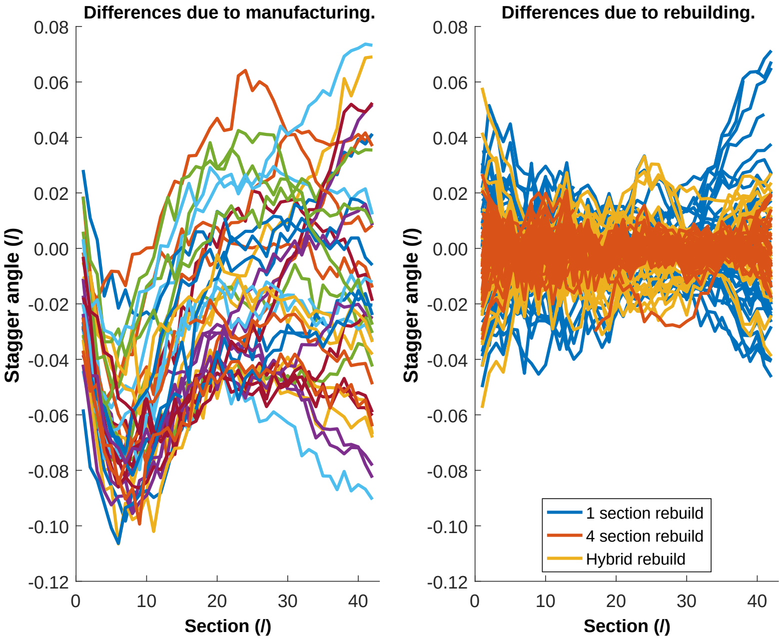

Figure 7.16: The stagger angle differences due to manufacturing (left), and the stagger angle blade-to-blade behaviour losses due to the rebuild process (right).

Maximum camber

The loss of blade-to-blade behaviour of stagger angle due to rebuild process can be seen in figure 7.17. Note that this figure shows similar differences between the rebuilding process results as figure 7.14.

It can be seen that both the 4 section and hybrid approaches reduce the amount of behaviour lost in the region

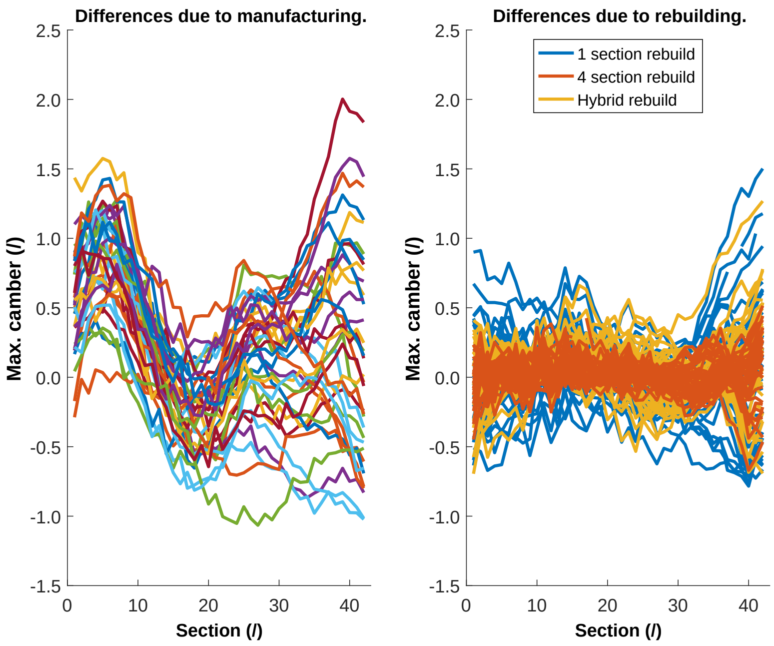

Figure 7.17: The maximum camber differences due to manufacturing (left), and the maximum camber blade-to-blade behaviour losses due to the rebuild process (right).

near the hub and the tip significantly. However, near the tip the hybrid process produced an outlier.

The differences between the result of the 4 section and hybrid approaches are again generally small, apart from the outlier of the hybrid process. However, comparing the behaviour lost to the differences due to manufacturing, it can be seen that that the losses of behaviour at the hub and at the tip are high for all approaches. It is therefore unlikely that any of them is capable of accurately simulating the blade-to-blade differences of the maximum camber due to manufacturing, as seen in the left panel of figure 7.17.

7.3.5. Convergence of parameter distributions

Another aspect of the geometry rebuild are the distributions. The camber and thickness distributions are morphed in the rebuilding process, and therefore their convergence to the expected results should be analysed. This is done by comparing the differences of the distributions caused by manufacturing to the differences between the distributions rebuilt by 1 section, 4 section and hybrid approaches, and the manufactured distributions they are intended to represent. The blade-to-blade behaviour that is lost is calculated by subtracting the measured distributions from the rebuilt ones.

Thickness distributions

The comparison of the thickness distribution manufacturing differences and the blade-to-blade differences lost in the rebuilding using the 1 section approach can be seen in figure 7.18. The comparison of the blade-to-blade differences lost using the hybrid and the 4 section approaches can be seen in figure 7.19. The different spanwise sections are denoted using different colors.

From figures 7.18 and 7.19 it can be seen that the 1 section and hybrid approaches give identical results. This is due to the hybrid approach uses only one delta parameter for each of the morphing points coordinates. In other words, the hybrid and 1 section approaches are equivalent in this case, and use 6 delta parameters to morph all the spanwise thickness distributions. The 4 section approach uses 4 times as many, i.e. 24.

It can be seen that the 4 section approach does lead to less blade-to-blade behaviour being lost. This can be seen as the scatter of the red lines in the right panel of figure 7.19 is smaller than the corresponding scatter in the left panel of this figure. For the sections near the hub this approach does not reduce the amount of

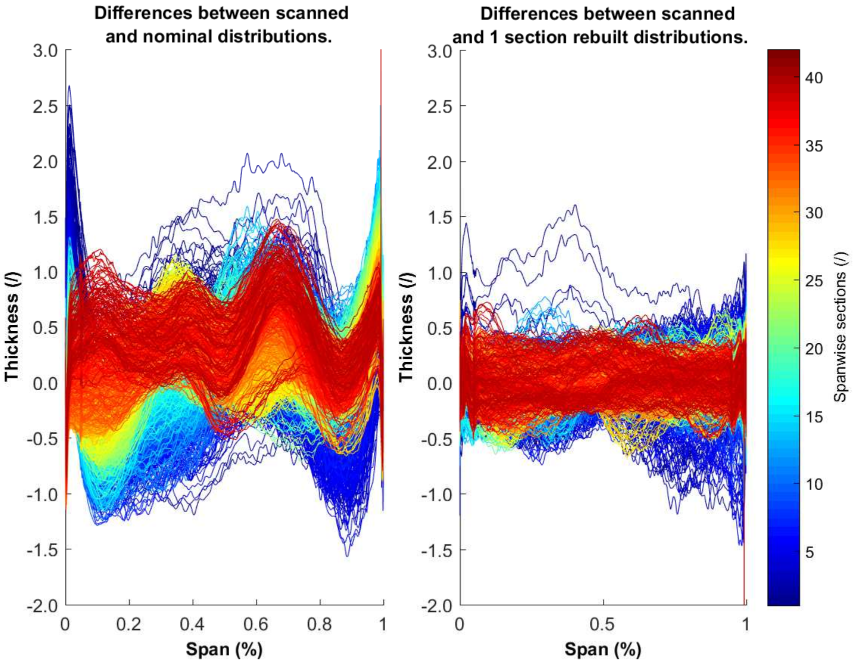

Figure 7.18: The thickness distribution differences due to manufacturing (left), and the thickness distribution blade-to-blade behaviour losses due to the rebuild process (right).

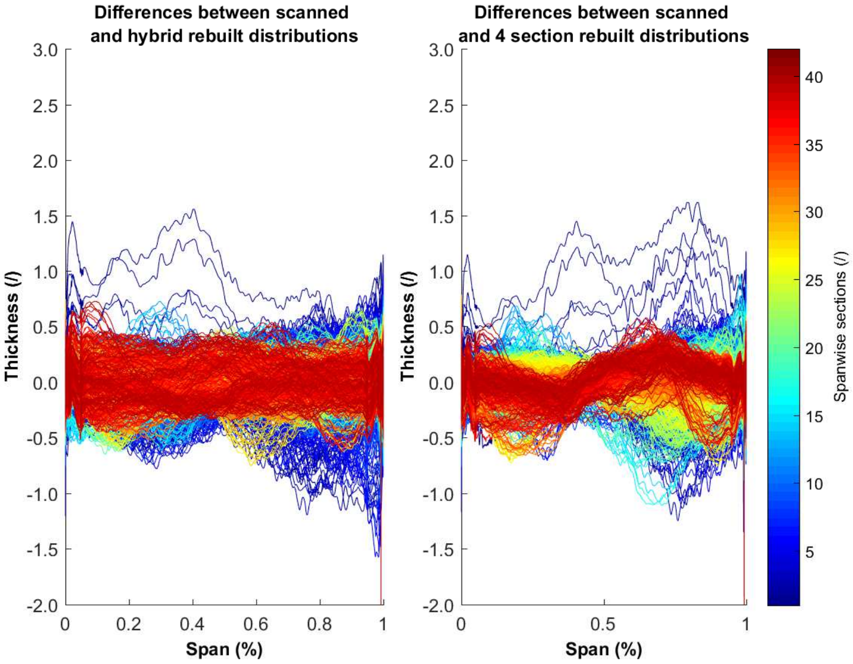

Figure 7.19: The thickness distribution differences due to manufacturing (left), and the thickness distribution blade-to-blade behaviour losses due to the rebuild process (right).

behaviour lost significantly, which can be seen by comparing the blue lines in the left and right panels of figure 7.19. The comparison shows that they are largely the same, with some differences near the TE.

Furthermore, it can be seen that the 1 section and hybrid rebuilds capture the chordwise variation of maximum thickness seen in the left panel of figure 7.18. On the other hand, the 4 section approach retains some chordwise variation from figure 7.18. Because this variation is not prominent in the 1 section rebuild it can be deducted that it is contained within the BLM. This suggests that the chordwise variability is caused by the 4 section approach.

Camber distributions

The comparison of the camber distribution manufacturing differences and the blade-to-blade differences lost in the rebuilding using the 1 section approach can be seen in figure 7.18. The comparison of the blade-to-blade differences lost using the hybrid and the 4 section approaches can be seen in figure 7.19. The different spanwise sections are denoted using different colors.

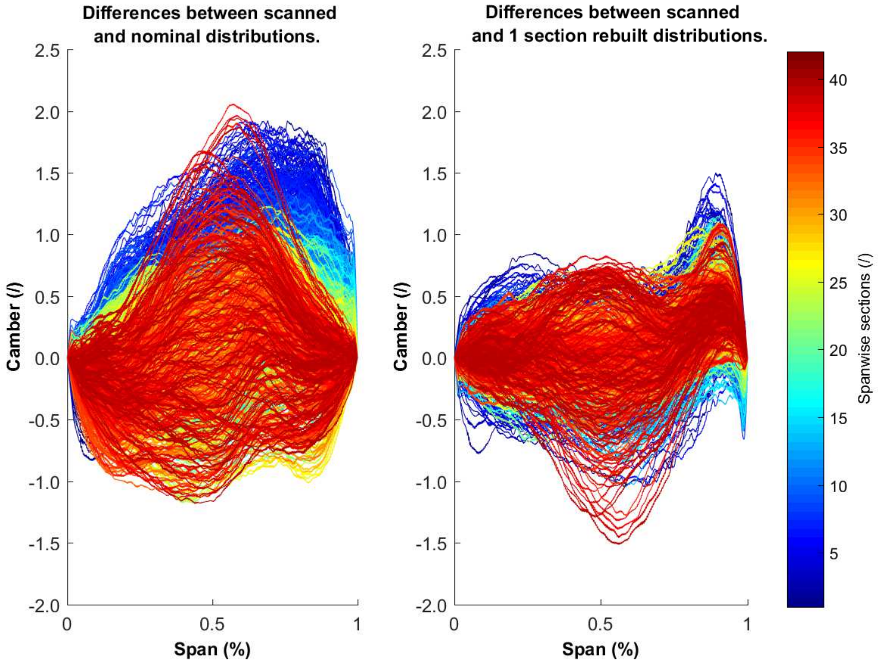

Figure 7.20: The camber distribution differences due to manufacturing (left), and the camber distribution blade-to-blade behaviour losses due to the rebuild process (right).

From figures 7.18 and 7.19 it can clearly be seen that the lowest amount of behaviour lost is using the 4 section approach, with the hybrid approach being a close second. However, for the sections near the hub TE the 4 section approach performs worse. Furthermore, note that the differences of the distributions rebuilt with the hybrid and the 4 section approach, as seen on figure 7.19 are roughly half of the differences caused by manufacturing, as seen in the left panel of figure 7.18. This means that half of the blade-to-blade behaviour has been lost.

7.3.6. Convergence of performance

In the performance analysis of manufactured and deteriorated blades of chapter 6 several QoI have been identified. These QoI are used to determine the convergence of the blade modelling. The figures of the convergence of the QoI show different behaviours, which are illustrated on the examples of the isentropic efficiency η_{is}, turning $\Delta\beta$ and total pressure ratio π_t.

Overall it appears that the hybrid approach performs best, however the performance convergence has only been achieved for $\Delta\beta$. Therefore it is possible that the performance scatter observed in this particular set of measured blades is too small to be accurately modelled with the current methodology. Here methodology stands for the measurement campaign, parameter extraction and rebuild processes.

In order to assess the convergence, the η_{is} differences between the measured and rebuilt manufactured blades, and the nominal blade have been plotted as normalised histograms on top of each other. Fitted Gaussian

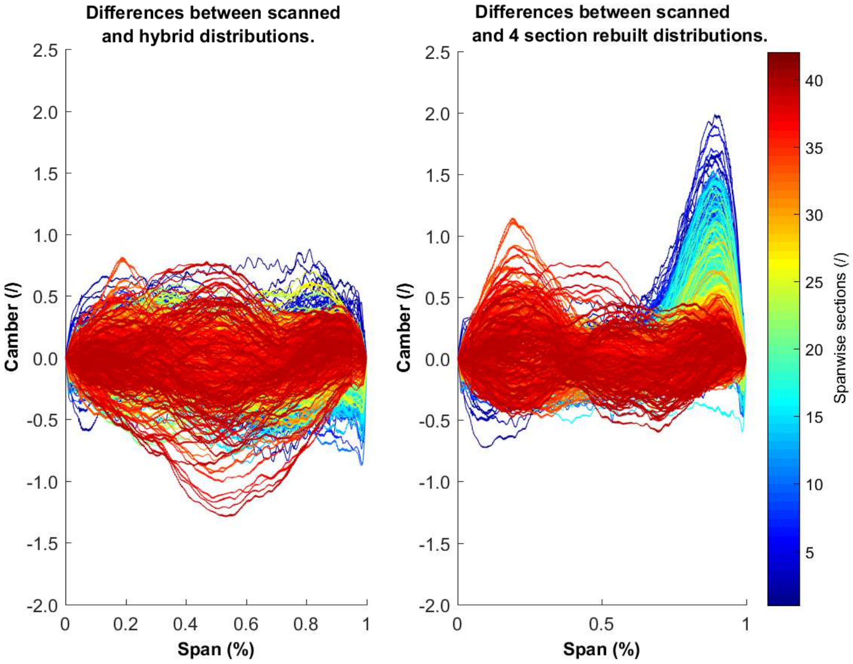

Figure 7.21: The camber distribution differences due to manufacturing (left), and the camber distribution blade-to-blade behaviour losses due to the rebuild process (right).

distributions and their means were also added for ease of comparison. The histograms were normalised such that their area equals 1.

Isentropic efficiency

Figure 7.22 shows the η_{is} convergence of the rebuilt blades. It can be seen that in the case of η_{is} the mean differences due to manufacturing are roughly -0.3. The mean rebuilt differences, at -0.15, are only roughly half of this however. Furthermore, the scatter of the measured manufactured blade η_{is} is much larger than the scatter of the rebuilt blades η_{is}. For the case of η_{is} none of the rebuild approaches has performed sufficiently.

The highest CoI corresponding to the manufactured blades, as shown in figure 6.21, are all spanwise average values, apart from the the spanwise behaviour slope of the LE position, k_{xLE}. Therefore, a better convergence of η_{is} with more spanwise delta parameters is not expected. Figure 7.22 appears to suggest this. In the hybrid approach x_{LE} is modelled using only one delta parameter, as opposed to two, which would allow to capture k_{xLE} in the rebuild as well. However, as figure 5.43 shows that the blade-to-blade differences are in fact mostly noise, it is not expected that adding more delta parameters for x_{LE} will improve the convergence. This is confirmed by the 4 section approach convergence of η_{is}.

Therefore, the cause for not achieving convergence of η_{is} is most likely hidden in the recreation of the thickness and camber distributions. Further studies, which are beyond the scope of this thesis, are required to determine if it is indeed the distribution morphing process that induces these differences.

Turning

Figure 7.23 shows the $\Delta\beta$ convergence of the rebuilt blades. It can be seen that in the case of $\Delta\beta$ the mean differences due to manufacturing are roughly 0.05. The mean offsets of the 1 and 4 section approaches are roughly twice as large, and therefore do not converge well in terms of $\Delta\beta$. However, the hybrid approach does appear to converge well.

Figure 6.21 shows that turning is influenced by many parameters, such as the LE position and length, chord length, stagger angle and maximum camber. However, these are all spanwise averaged values, while the blade-to-blade spanwise behaviour slopes do not appear to have a high influence. It is therefore not surprising

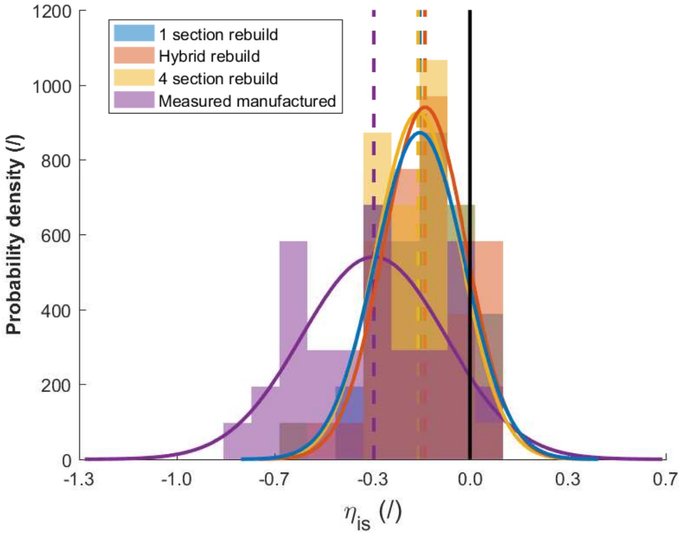

Figure 7.22: The convergence of rebuild manufacturing blade to the measured manufactured blade isentropic efficiencies.

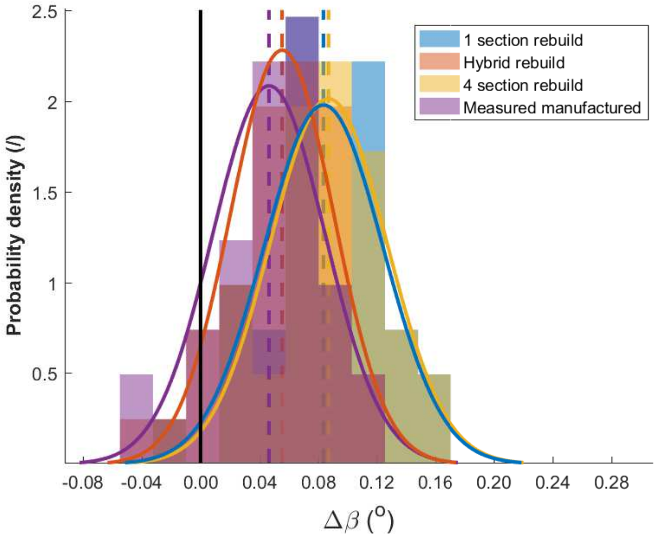

Figure 7.23: The convergence of rebuild manufacturing blade to the measured manufactured blade turning.

that the increase of spanwise parameters does not bring an improved convergence in terms of $\Delta\beta$, however, it is surprising that the hybrid approach performs better than both the 1 section and 4 section approaches.

Total pressure ratio

Figure 7.24 shows the π_t convergence of the rebuilt blades. It can be seen that in the case of π_t the mean differences due to manufacturing are roughly 0.1. The mean offsets of the 1 and 4 section approaches are roughly three times as large, and therefore do not converge well in terms of $\Delta\beta$. Note that the PDF of the 1 and 4 section approaches lie directly one on top of the other, therefore only the 1 section approach PDF is visible.

Similarly, the hybrid approach appears to not converge well, with the mean difference roughly twice the size of those due to manufacturing. The scatter appears to have been rebuilt relatively well in all cases however.

The parameters with the highest CoI are the spanwise mean stagger angle and the TE position, followed by the

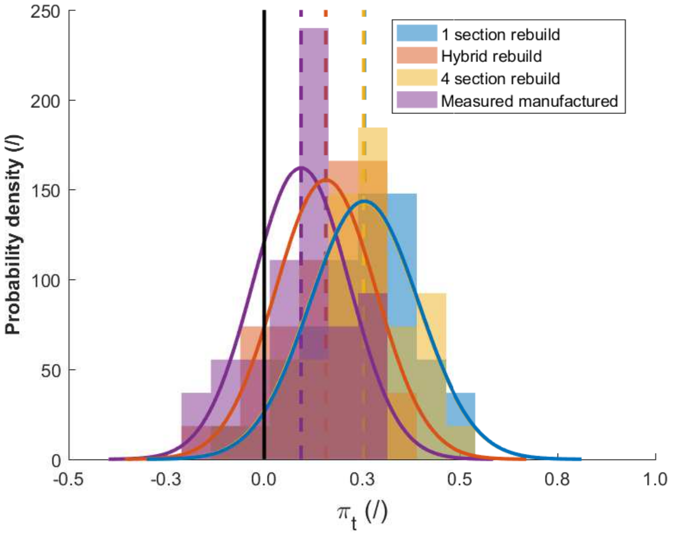

Figure 7.24: The convergence of rebuild manufacturing blade to the measured manufactured blade total pressure ratio.

spanwise behaviour slopes of the chord length and stagger angle. Therefore it seems that in this case the 1 section approach suffers from the lack of the spanwise behaviour information during the rebuild. On the other hand, the 4 section approach appears to be too dependant on the noise of the extracted parameters.

7.3.7. Convergence of sample size

In order to determine how many uncertain geometries need to be evaluated by CFD in order to estimate the performance uncertainty, a dedicated analysis of the impact of sample size on the performance characteristic has been performed. This has been done by simulating a number of geometries and evaluating them with 3D CFD. Based on the convergence of the performance PDFs the sample size for which the level of accuracy is acceptable can be selected.

The performance convergence depends directly on the number and distribution of the geometrical parameters, which in turn depend directly on the shape of the design and the manufacturing capability for that design. Because of this, the sample size convergence analysis should be performed independently when analysing different blade design measurements. Therefore the analysis has been constructed such that the amount of work required for it is minimised.

For the analysis a sample of 50 blade geometries has been created based on the delta parameter distributions described in section 7.2. Bootstrapping was then used to estimate the sampling distributions of the mean and standard deviation of a performance QoI. The performance QoI selected for the analysis is the Loss coefficient, described in chapter 6. The sample size is based on the convergence of the sampling distributions of the mean and the standard deviations of the loss coefficient.

Bootstrapping

Bootstrapping is a method that attempts to estimate the sampling distribution of a particular statistic. A particular statistic might be a mean, standard deviation or other value that summarises information of a sample. Note that the sample is a subset of the population, through which the statistics of the population are inferred. A population will have exact statistics as all the possible elements are included in their calculation. A sample will have uncertain statistics as only a few elements of the population are used in the calculation of its statistics. The PDF that shows which sample statistics are more probable is the sampling distribution of that statistic.

The more elements there are in the sample the more the sampling distribution will hone in on the population statistic. When all the elements of a population are included in the sample, the sampling distribution will have 0 probability density for all values other than the population statistic, where it will be infinite.

This procedure can be applied if the samples are taken directly from the population. However, if the assumption that the sample is representative of the population (independent and identically distributed), the sample can assume the role of the population. The sampling distribution is then based on the subsamples of the original sample. By the assumption made, the sampling distribution of the subsample statistics is also the sampling distribution of the sample statistics.

For this assumption to hold the sampling from the sample must be done with replacement, which means that any element in the sample selected for the subsample, remains in the sample for the next element selection. This ensures that the assumption that the sample is representative of the population is preserved. Essentially, sampling from a PDF can be seen as sampling with replacement, because no value is removed from its domain after being selected. By sampling with replacement the empirical PDF of the sample is therefore preserved. Note that this also means that if the subsample size is the same as the sample size, the sampling distribution will still be distributed around some value. This is because due to replacement some elements will almost certainly be selected more than once.

The parameter that needs to be selected for bootstrapping is the size of the subsample. This parameter needs to be sufficiently high, but is selected arbitrarily.

Convergence of sampling distributions of Loss Coefficient

For the use of bootstrapping the assumption of the sample being representative of the population must be valid. In the current case, the populations are the multidimensional PDF that characterise manufacturing and deterioration. Therefore the assumption is satisfied by construction of the rebuild blades, as the delta parameters are sampled such that they have statistics as close to the population statistics as possible.

Within the bootstrapping framework, a sample of 50 blades was created and evaluated using 3D Hydra to mimic the population. As the populations are normally distributed the statistics of interest in the convergence of sample size are the mean and the standard deviation. Their sampling distributions can be seen in figure 7.25.

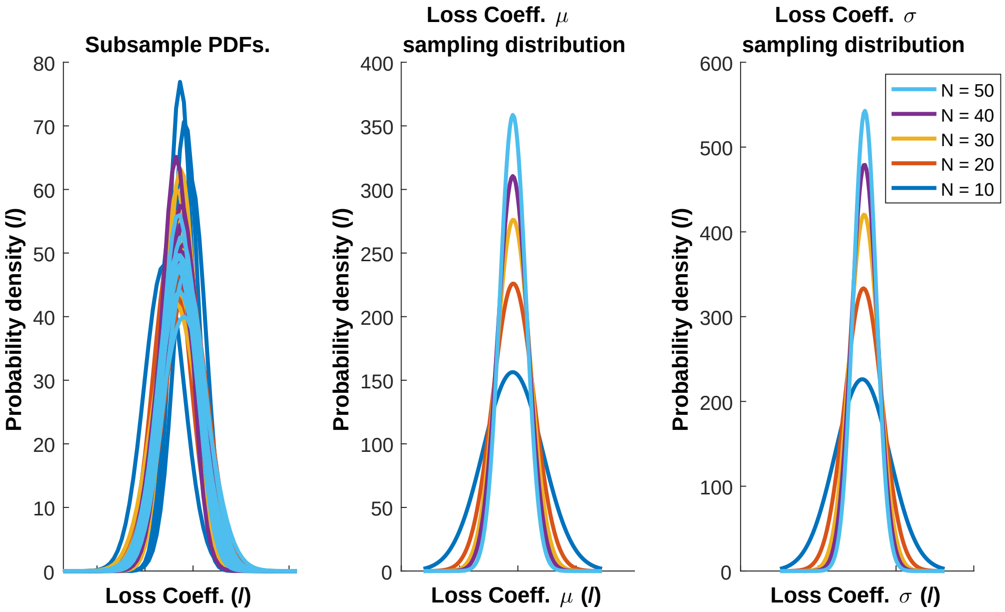

Figure 7.25: PDFs of selected Loss Coefficient subsamples (left), and the sampling distribution of their means (centre) and standard deviations (right).

Figure 7.25 contains three panels. The left panel shows the PDF of the subsamples selected from the sample. This PDF is a Gaussian distribution fitted to the subsample data for ease of visualisation. This can be done as the sample is normally distributed, and therefore the subsample is also expected to be normally distributed.

The left panel shows the scatter of the subsample PDF due to sampling. It can be seen that the lines corresponding to $N = 10$ are the most scattered, while the lines corresponding to $N = 50$ are the most consistent. Furthermore, it can be seen that the lines corresponding to $N = 10$ containt the narrowest PDF. This is because for low sample sizes it is possible that most of the values fall close together, resulting in narrow distributions. Furthermore, because of this the scatter between the PDF is larger.

In order to characterise the behaviour better, the centre and right plots show the sampling distributions of the mean and standard deviations of these PDF respectively. From both of these plots it can be seen, that the sampling distributions become increasingly narrower with the increase in the subsample size. Furthermore, it can be seen that the difference between $N = 10$ and the others is relatively large, whereas thereafter the differences become relatively small. The sampling distributions of $N = 40$ and $N = 50$ are virtually the same, with very little differences between them. Therefore $N = 40$ is selected as the sample size for geometry generation and CFD evaluation.

7.4. Lessons learned in modelling effects of manufacturing and deterioration

The aim of this chapter was to show how to model the influence of manufacturing and deterioration on geometrical NACA like parameters. Section 7.2 showed that deterioration is independent from manufacturing variability for this particular blade design. Furthermore, section 7.2 proposed three approaches in order to create new blades with simulated effects of manufacturing and deterioration.

Section 7.3 showed that of these approaches, the so called hybrid approach performed best. Using this approach, sufficient rebuild convergence has been achieved for surface deviations and most of the rebuild parameters. However, sufficient convergence has not been achieved for the rebuild of the camber and thickness distributions. Furthermore, the performance QoI did not all converge to the expected values.

This suggests that the currently employed process is currently not capable of accurately simulating the manufacturing impact, of the magnitude seen in chapter 5, on blade geometry. The process here includes the measurement campaign, and the parameter extraction and blade rebuild processes.

As the impact of deterioration on geometry is smaller than the impact of manufacturing, it also cannot currently be modelled accurately. The modelling accuracy could be improved by using equipment with smaller resolution in the measurement campaign. However, blades with higher levels of geometry variability could still be modelled by the current process.

An alternate approach for reduction of parameters could be established with the use of Adjoint CFD. Studies on the sensitivities of performance by using both the CoI and the Adjoint technique have already been performed by Schmidt et. al. [56]. An analysis using Adjoint CFD could be performed as a part of the convergence study as well.

Bibliography

[1] Shannon Ackert. Basics of aircraft maintenance reserve development and management. `http://www.aircraftmonitor.com/uploads/1/5/9/9/15993320/basics_of_aircraft_maintenance_reserves___v1.pdf`, 2012. Accessed: 2017-08-20.

[2] Hawaiian Airlines. Hawaiian airlines to purchase up to 20 new boeing 717-200s. `https://newsroom.hawaiianairlines.com/releases/hawaiian-airlines-to-purchase-up-to-20-new-boeing-717-200s`, 1999. Accessed: 2017-08-20.

[3] Roy Andersson, Henrik Eriksson, and Håkan Torstensson. Similarities and differences between tqm, six sigma and lean. *The TQM Magazine*, 18(3):282–296, 2006. doi: 10.1108/09544780610660004.

[4] Ilya Arsenyev, Fabian Duddeck, and Andreas Fischersworring-Bunk. Efficient surrogate-based robust shape optimization for vane clusters. In *16th AIAA/ISSMO Multidisciplinary Analysis and Optimisation Conference*. AIAA, 2015.

[5] T.S. Baines, H.W. Lightfoot, O. Benedettini, and J.M. Kay. The servitization of manufacturing: A review of literature and reflection on future challenges. *Journal of Manufacturing Technology Management*, 20(5): 547–567, 2009. doi: 10.1108/17410380910960984.

[6] Dieter Bestle, Peter Flassig, and Amit Kumar Dutta. Robust design of compressor blades in the presence of manufacturing noise. In *9th European Conference on Turbomachinery*, 2011.

[7] Nadia Bhuiyan and Amit Baghel. An overview of continuous improvement: from the past to the present. *Management Decision*, 43(5):761–771, 2005. doi: 10.1108/00251740510597761.

[8] Jens J. Dahlgaard and Su Mi Dahlgaard-Park. Lean production, six sigma quality, tqm and company culture. *The TQM Magazine*, 18(3):263–281, 2006. doi: 10.1108/09544780610659998.

[9] Antonio D'Ammaro and Francesco Montomoli. Uncertainty quantification and film cooling. *Computers & Fluids*, 71:320–326, 2013.

[10] Ihor S Diakunchak. Performance deterioration in industrial gas turbines. *Journal of Engineering for Gas Turbines and Power*, 114(2):161–168, 1992.

[11] Eric A Dow and Qiqi Wang. Optimal design and tolerancing of compressor blades subject to manufacturing variability. In *16th AIAA Non-Deterministic Approaches Conference*, page 1008, 2014.

[12] J. D. Duffner. The effects of manufacturing variability on turbine vane performance. Master's thesis, Massachusets Institute of Technology, 2008.

[13] General Electric. General electric annual report. Website, 1997. `http://www.annualreports.com/HostedData/AnnualReportArchive/g/NYSE_GE_1997.pdf`.

[14] General Electric. General electric annual report. Website, 2000. `http://www.ge.com/annual00/download/images/GEannual00.pdf`.

[15] Martha Gardner and Gene Wiggs. Design for six sigma: The first 15 years. *Volume 7A: Structures and Dynamics*, 2013. doi: 10.1115/gt2013-95534.

[16] Victor E Garzon and David L Darmofal. Impact of geometric variability on axial compressor performance. In *ASME Turbo Expo 2003, collocated with the 2003 International Joint Power Generation Conference*, pages 1199–1213. ASME, 2003.

[17] GOM. Gom inspect, 08 2017. URL `http://www.gom.com/index.html`.

[18] Martin N Goodhand, Robert J Miller, and Hang W Lung. The impact of geometric variation on compressor two-dimensional incidence range. *Journal of Turbomachinery*, 137(2):021007, 2015.

[19] Michael Gorelik, Jacob Obayomi, Jack Slovisky, Dan Frias, Howie Swanson, John McFarland, Michael Enright, and David Riha. Effect of manufacturing variability on turbine engine performance: a probabilistic study. In *ASME Turbo Expo 2013: Turbine Technical Conference and Exposition*, pages V07AT28A003–V07AT28A003, 2013.

[20] A. Hamed, W. Tabakoff, and D. Singh. Modelling of compressor performance deterioration due to erosion. *International Journal of Rotating Machinery*, 4(4):243–248, 1998.

[21] Kay Heinze, Marcus Meyer, Jens Scharfenstein, Matthias Voigt, and Konrad Vogeler. A parametric model for probabilistic analysis of turbine blades considering real geometric effects. *CEAS Aeronautical Journal*, 5(1):41–51, 2014.

[22] Lars Högner, Alkin Nasuf, Paul Voigt, Matthias Voigt, Konrad Vogeler, Marcus Meyer, Christopher Berridge, and Frederic Goenaga. Analysis of high pressure turbine nozzle guide vanes considering geometric variations. In *ASME Turbo Expo 2016: Turbomachniery Technical Conference and Exposition*, pages V02CT45A025–V02CT45A025. ASME, 2016.

[23] S. Holtzhausen, S. Schreiber, Ch. Schöne, R. Stelzer, K. Heinze, and A. Lange. Highly accurate automated 3d measuring and data conditioning for turbine and compressor blades. In *ASME Turbo Expo 2009: Turbomachniery Technical Conference and Exposition*. ASME, 2009.

[24] A Javed, R Pecnik, and JP Van Buijtenen. Optimization of a centrifugal compressor impeller for robustness to manufacturing uncertainties. *Journal of Engineering for Gas Turbines and Power*, 138(11):112101, 2016.

[25] Yaping Ju and Chuhua Zhang. Robust design optimisation method for centrifugal impellers under surface roughness uncertainties due to blade fouling. *Chinese Journal of Mechanical Engineering*, 29(2):301–314, 2016.

[26] Bengt Klefsjö, Håkan Wiklund, and Rick L. Edgeman. Six sigma seen as a methodology for total quality management. *Measuring Business Excellence*, 5(1):31–35, 2001. doi: 10.1108/13683040110385809.

[27] Daria Kolmakova, Grigorii Popov, Oleg Baturin, and Alexander Krivcov. The quantification of impact of geometrical and operational parameters uncertainty on the numerical simulation results of axial turbine working process. In *ASME Turbo Expo 2015: Turbine Technical Conference and Exposition*, pages V02BT39A009–V02BT39A009. ASME, 2015.

[28] A Kumar, PB Nair, AJ Keane, and S Shahpar. Probabilistic performance analysis of eroded compressor blades. In *ASME 2005 Power Conference*, pages 1175–1183. ASME, 2005.

[29] Apurva Kumar, AJ Keane, PB Nair, and S Shahpar. Robust design of compressor blades against manufacturing variations. In *ASME 2006 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pages 1105–1118. ASME, 2006.

[30] Apurva Kumar, AJ Keane, PB Nair, and S Shaphar. Efficient robust design for manufacturing process capability. In *6th ASMO-UK/ISSMO International Conference on Engineering Design Opimization*, 2006.

[31] Apurva Kumar, Andy J Keane, Prasanth B Nair, and Shahrokh Shahpar. Robust design of compressor fan blades against erosion. *Journal of Mechanical Design*, 128(4):864–873, 2006.

[32] R. Kurz, G. Musgrove, and K. Brun. Experimental evaluation of compressor blade fouling. In *Journal of Engineering for Gas Turbines and Power*, 2017.

[33] Rainer Kurz and Klaus Brun. Fouling mechanisms in axial compressors. *Journal of Engineering for Gas Turbines and Power*, 134(3):032401, 2012.

[34] AN Lakshminarasimha, MP Boyce, and CB Meher-Homji. Modelling and analysis of gas turbine performance deterioration. In *ASME 1992 International Gas Turbine and Aeroengine Congress and Exposition*, pages V004T10A022–V004T10A022. ASME, 1992.

[35] Alexander Lange, Konrad Vogeler, Volker Gümmer, Henner Schrapp, and Carsten Clemen. Introduction of a parameter based compressor blade model for considering measured geometry uncertainties in numerical simulation. In *ASME Turbo Expo 2009: Power for Land, Sea, and Air*, pages 1113–1123. ASME, 2009.

[36] Alexander Lange, Matthias Voigt, Konrad Vogeler, Henner Schrapp, Erik Johann, and Volker Gümmer. Probabilistic cfd simulation of a high-pressure compressor stage taking manufacturing variability into account. In *ASME Turbo Expo 2010: Power for Land, Sea, and Air*, pages 617–628. ASME, 2010.

[37] Alexander Lange, Matthias Voigt, Konrad Vogeler, and Erik Johann. Principal component analysis on 3d scanned compressor blades for probabilistic cfd simulation. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, page 1762, 2012.

[38] Alexander Lange, Matthias Voigt, Konrad Vogeler, Henner Schrapp, Erik Johann, and Volker Gümmer. Impact of manufacturing variability and nonaxisymmetry on high-pressure compressor stage performance. *Journal of Engineering for Gas Turbines and Power*, 134(3):032504, 2012.

[39] Alexander Lange, Matthias Voigt, Konrad Vogeler, Henner Schrapp, Erik Johann, and Volker Gümmer. Impact of manufacturing variability on multistage high-pressure compressor performance. *Journal of Engineering for Gas Turbines and Power*, 134(11):112601, 2012.

[40] J Marx, J Städing, G Reitz, and J Friedrichs. Investigation and analysis of deterioration in high pressure compressors due to operation. *CEAS Aeronautical Journal*, 5(4):515–525, 2014.

[41] C. B Meher-Homji and A. Bromley. Gas turbine axial compressor fouling and washing. In *33rd Turbomachinery Symposium, Houston, TX, September*, pages 20–23, 2004.

[42] Curtis W Moeckel, David L Darmofal, T Robert Kingston, and Robert JG Norton. Toleranced designs of cooled turbine blades through probabilistic thermal analysis of manufacturing variability. In *ASME Turbo Expo 2007: Power for Land Sea and Air*, pages 1179–1191. ASME, 2007.

[43] P. Moinier, J.-D. Müller, and M.B. Giles. Edge-based multigrid and preconditioning for hybrid grids. *AIAA journal*, 40(10):1954–1960, 2002.

[44] Douglas C. Montgomery. *Introduction to Statistical Quality Control*. Wiley, 2008. ISBN 0470169923.

[45] J.F. Montgomery III, T.R. Sewall, and J.J. Batka. Aircraft usage and effects on engine life. *American Society of Mechanical Engineers*, 1983.

[46] F Montomoli, M Massini, and S Salvadori. Geometrical uncertainty in turbomachinery: Tip gap and fillet radius. *Computers & Fluids*, 46(1):362–368, 2011.

[47] Mirko Morini, Michele Pinelli, Pier Ruggero Spina, and Mauro Venturini. Numerical analysis of the effects of nonuniform surface roughness on compressor stage performance. *Journal of Engineering for Gas Turbines and Power*, 133(7):072402, 2011.

[48] Andy Neely. The servitization of manufacturing: an analysis of global trends. In *14th European Operations Management Association Conference, Ankara, Turkey*, pages 1–10, 2007.

[49] Rolls-Royce plc. *The Jet Engine*. Rolls-Royce plc, 1996.

[50] Rolls-Royce plc. `http://www.rolls-royce.com/media/press-releases/yr-2015/pr-17c-06-15-rolls-royce-trent-1000-engines-worth-$500m-to-power-ethiopian-airlines.aspx`, 2015. Accessed: 2017-08-20.

[51] Ricardo F. Powell. Impact of parametric aerothermal variability on aircraft engine operating cost, 2002.

[52] R Qui, Yaping Ju, Yazhou Wang, and Chuhua Zhang. Flow analysis and uncertainty quantification of a 2d compressor cascade with dirty blades. *ASME Paper No. GT2016-56915*, 2016.

[53] Gerald Reitz and Jens Friedrichs. Procedure for analyzing, manipulating and meshing of compressor blades to simulate their flow. *International Journal of Gas Turbine, Propulsion and Power Systems*, 8(1), 2016.

[54] A. Savitzky and M.J.E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964.

[55] Jens Scharfenstein, Kay Heinze, Matthias Voigt, Konrad Vogeler, and Marcus Meyer. Probabilistic cfd analysis of high pressure turbine blades considering real geometric effects. In *ASME Turbo Expo 2013: Turbine Technical Conference and Exposition*, pages V06BT43A005–V06BT43A005, 2013.

[56] Robin Schmidt, Matthias Voigt, Konrad Vogeler, and Marcus Meyer. Comparison of two methods for sensitivity analysis of compressor blades. In *ASME Turbo Expo 2016: Turbomachniery Technical Conference and Exposition*, pages V07BT29A003–V07BT29A003. ASME, 2016.

[57] S. Shahpar and L. Lapworth. Padram: Parametric design and rapid meshing system for turbomachinery optimisation. In *American Society of Mechanical Engineers, International Gas Turbine Institute, Turbo Expo (Publication) IGTI*, volume 6, 01 2003.

[58] C. Spearman. The proof and measurement of association between two rings. *American Journal of Psychology*, 15:72–101, 1904.

[59] Stefan Spieler, Stephan Staudacher, Roland Fiola, Peter Sahm, and Matthias Weißschuh. Probabilistic engine performance scatter and deterioration modeling. *Journal of engineering for gas turbines and power*, 130(4):042507, 2008.

[60] Geoff Tennant. *Six Sigma: SPC and TQM in manufacturing and services*. Gower Publishing, Ltd., 2001. ISBN 9780566083747.

[61] Kirill A. Vinogradov, Roman A. Didenko, Kseniya V. Otryahina, Karelin Dmitry V., Gennady V. Kretinin, and Shmotin Yury N. Robust optimisation of the hpt cooling blade in conjugate heat transfer computations. In *ASME Turbo Expo 2015: Turbomachniery Technical Conference and Exposition*, pages V02CT45A003–V02CT45A003. ASME, 2016.

[62] Paul Voigt, Lars Hoegner, and Thomas Backhaus. *Blade2Parameter: Geometry Reconstruction User Guide*. Technische Universität Dresden, Institut für Strömungsmechanik, Helmholtzstr. 10, 01069 Dresden, april 2017 edition, 4 2017.

[63] Aviation Week. Business airplane tables. `http://aviationweek.com/site-files/aviationweek.com/files/uploads/2015/05/Business%20Airplane%20Tables_May_2015_revised.pdf`, 2015. Accessed: 2017-08-20.

[64] Anthony V Zaita, Greg Buley, and Glen Karlsons. Performance deterioration modeling in aircraft gas turbine engines. In *ASME 1997 International Gas Turbine and Aeroengine Congress and Exhibition*, pages V004T16A007–V004T16A007. ASME, 1997.