Imperial College London
Department of Aeronautical Engineering

A Data-Driven Uncertainty Quantification Method for Scarce Data and Rare Events

Richard Benedikt Heinrich Ahlfeld

October 2017

Supervised by Dr Francesco Montomoli
and Prof Enrico Scalas

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Aeronautical Engineering of Imperial College London
and the Diploma of Imperial College London

Abstract

The efficient integration of manufacturing, experimental or operational data into physical simulations using uncertainty quantification techniques is becoming increasingly important for the engineering industry. There is a genuine interest to use uncertainty quantification methods to improve product efficiency, reliability, and safety. However, temporary methods still have various limitations, which have prevented a wider industrial application. This thesis analyses existing problems in industry and contributes new solutions in five areas: data scarcity, the curse of dimensionality, rare events, discontinuous models and epistemic model-form uncertainty.

The major novelty of this work is a new data-driven Polynomial Chaos framework called SAMBA that can be used to improve engineering designs by accounting for uncertainties in simulations more efficiently and accurately. SAMBA provides a single solution to many industrial problems. It is particularly useful for applications where only scarce statistical data is available and to efficiently quantify the likelihood of rare events with disastrous consequences, the so-called Black Swan. Other benefits are the simple reconstruction of adaptive and anisotropic sparse grid quadrature rules to alleviate the curse of dimensionality, a simpler combination of arbitrary distributions and random data within a single method, and higher accuracy for data sets that do not follow a definable distribution. To deal with special cases, two extensions have been added to SAMBA's framework: first, a modification that allows the use of arbitrary points within Padé approximations for models containing discontinuities and second, a general multi-fidelity framework to counteract model-form uncertainty in simulations.

The thesis concludes with the application of SAMBA to various state-of-the-art models in Formula 1, turbomachinery, and space flight (one example analyses structural dynamics uncertainty of NASA's Space Launch System). Particularly innovative are two applications: one showing how to account for Black Swans in gas turbine simulations and one how to deal with scarce data in Formula 1.

Contents

1 Introduction: Uncertainty in Engineering Simulations 27

1.1 The Science of Uncertainty Quantification (UQ) 28

1.1.1 Uncertainty Propagation 29

1.1.2 Advantages for Industry 31

1.2 Challenges for Industrial Application 33

1.2.1 Limited Data and Statistical Assumptions 34

1.2.2 Limited Time and the Curse of Dimensionality 36

1.2.3 Rare Events and Black Swans 38

1.2.4 Models with Discontinuities 39

1.2.5 Epistemic Model-Form Uncertainty 40

1.3 Thesis Statement and Contributions 42

1.4 Thesis Content 45

2 New Algorithm for Data-Driven Polynomial Chaos 46

2.1 Need for Data-Driven Polynomial Chaos 46

2.2 Theory of Polynomial Chaos Expansions 47

2.2.1 Gaussian Quadrature 49

2.3 New Algorithm for Data-Driven Polynomial Chaos 50

2.3.1 Ill-Conditioned Moment Matrix 54

2.4 Demonstration for Different Inputs 54

2.4.1 Probability Distributions 54

2.4.2 Random Samples 54

2.4.3 Propagation of Multi-Modal Data 56

2.5 Validation: CFD Simulation of a Diffuser 57

3 SAMBA: Adaptive and Anisotropic Sparse Grids 61

3.1 SAMBA's Algorithm 62

3.1.1 Sensitivity Analysis using Sobol Indices 64

3.1.2 Simple Example and Implementation Advice 66

3.2 Validation of SAMBA 68

3.2.1 Effect of Sample Size 70

3.2.2 Error Compared to PDF Fitting 72

3.3 Rare Event Simulation using SAMBA 72
3.3.1 Fat-Tailed Probability Distributions 74
3.3.2 Tempered Alpha Stable Distributions 75
3.3.3 Fitting Alpha-Stable Distributions 77
3.3.4 Validation of SAMBA for Rare Events 78
3.4 Applicability and Limitations 80
3.4.1 Integration Errors 80
3.4.2 Scarce Data Errors 81
3.4.3 Convergence 81
3.4.4 Higher Order Coefficient Errors 82

4 A Global Discontinuous Model Based on Arbitrary Points 86
4.1 Matrix Notation 87
4.2 A Stabilised and Automated Least Squares Padé Approximation 89
4.2.1 New Least Squares Algorithm 90
4.2.2 Stabilisation via Optimal Quadrature Sub-sampling . . 92
4.2.3 Automation Using Statistical Learning 93
4.3 Validation 95

5 Multi-Fidelity to Reduce Model-Form Uncertainty 98
5.1 Multi-Fidelity Polynomial Chaos 99
5.1.1 Nested optimal SAMBA grids 100
5.1.2 Weighted Regression 102
5.2 Demonstration for Turbulent Flow over Periodic Hills 104
5.2.1 Direct Numerical Simulation 106
5.2.2 RANS Simulations 108
5.2.3 RANS and DNS Compatibility 109
5.2.4 Multi-fidelity Results of RANS and DNS 114

6 Engineering Applications 116
6.1 Race Car Aerodynamics 117
6.1.1 Wing Tip Vortex 117
6.1.2 F1 Race Car Front Wing 119
6.2 Turbomachinery 124
6.2.1 Manufacturing Errors in Low-Pressure Gas Turbine Blades 125
6.2.2 Tip Leakage in a Multistage Gas Turbine 129
6.3 Structural Dynamics 137
6.3.1 NASA's Space Launch System 137

6.4 Rare Events in Aircraft Engine Simulations 144
6.4.1 Transients in Gas Turbine Cycle Missions 146
6.4.2 Hot Gas Ingestion 149
6.5 Discontinuities in Aircraft Engine Simulations 154
6.5.1 Film Cooling and Shock Interaction 154

7 Conclusion **166**
7.1 Summary of Contributions 167
7.2 Future Research 168

List of Figures

1.1 Information about a technical problem can be displayed as a matrix of knowledge with four quadrants. Mathematical uncertainty quantification takes place in the third quadrant. 28
1.2 Non-intrusive uncertainty propagation techniques can quantify the effect of input variability on model predictions without modifying a model. 29
1.3 Non-intrusive uncertainty quantification methods are trivial to parallelise, because individual model runs/samples are never connected. 31
1.4 Uncertainty quantification methods have many advantages to support engineering design in addition to providing more rigorous error bars: predicting natural variability, finding important parameters and making designs more robust under variable operating conditions. 32
1.5 Illustration of three scenarios of variance estimation: 1) Overestimation 2) Underestimation 3) Correct Estimation. The black dots represent available data. The red dot is a yet unknown data point. 34
1.6 This figure shows the effect of different probability distributions on the temperature variance field of a heat exchanger CFD simulation. Even though the range of the distributions is similar, the Gaussian PDF creates significantly higher variance. 35
1.7 Patterns in random processes can be misleading. The stochastic process depicted was only observed until the vertical line evoking the impression that it is periodic. 36
1.8 Illustration of the curse of dimensionality. The number of model evaluations rises quickly from one to three dimensions. 37
1.9 The visual difference between a Gaussian and a fat-tailed distribution is small. Only zooming into the tail region reveals that rare events are significantly more likely for fat-tailed distributions. 38

1.10 The left graph shows film-cooling/shock interactions in a transonic CFD simulation. The right graph shows a discontinuous response surface for a local flow parameter that is affected by a shock in the CFD domain. In order to accurately resolve the discontinuity in the response surface, special uncertainty quantification methods are needed. 40
1.11 The top graph show the streamlines of flow over periodic hills for a DNS (left) and a RANS (right) solver at $Re = 700$. Below are the corresponding PDFs for the reattachment location calculated using DNS and RANS. The large difference between the PDFs illustrates the effect of epistemic model-form uncertainty on UQ results. 41

2.1 Optimal Gaussian collocation points calculated with data-driven PC for three Askey scheme PDFs: standard Normal $N(0,1)$, Uniform $U[0,1]$, Exponential $exp(3)$ and three non-Askey scheme PDFs: fat-tailed Student-t $t(3)$, Weibull $Wb(1,1.5)$ and generalised Extreme Value Distribution $gev(1)$. 55
2.2 Optimal Gaussian collocation points calculated with data-driven PC for various mixed and multi-modal histograms. The results show that the method adapts well to multi-modal data. The collocation points favour points of high likelihood. 55
2.3 This graph compares the results of Monte Carlo Simulation with data-driven PCE for a multi-modal input histogram. The results in (b) and (c) show that multi-modality is not a problem as long as the input/output mapping is continuous and sufficiently smooth. 56
2.4 This figure shows the mean velocity field of the diffuser flow on a coarser mesh than used for CFD. The taper angle α is highlighted at the bottom side of the channel. 58
2.5 The taper angle α of the diffuser varies stochastically in the CFD simulation. This histogram shows the various values the taper angle can randomly assume. 58
2.6 Various parametric PDFs are fitted to the available random samples of the diffuser taper angle. The different PDFs show that even though 100 samples is not a small data set, there is a large difference between the various distributions. 60

3.1 Illustration of how Smolyak's sparse grid algorithm alleviates the curse of dimensionality. The number of collocation points for a level 2 Smolyak grid is orders of magnitude lower than for a 3^{rd} order full tensor grid. 64
3.2 This figure illustrates the high level of generality and adaptability that is obtained by using SAMBA for quadrature grid generation. The graphs show various symmetric and asymmetric data-driven sparse grids for parametric distributions and random data histograms. 65
3.3 Response surfaces of the four multi-dimensional non-linear test functions from Table 3.1 including level 2 data-driven sparse grids for one Gaussian and one Weibull input distribution. . . . 69
3.4 The four graphs show the relative error of the predicted mean including 100% confidence intervals for full tensor Gaussian and Smolyak's method compared to MCM for increasing sample size and for the first 2 and 5 input distributions from Table 3.2. It can be seen that the error grows with the number of inputs and decreases with the number of random samples available. The confidence intervals are created by repeating the uncertainty quantification process with different random samples. 71
3.5 The two figures compare SAMBA with Polynomial Chaos based on Gaussian PDF fitting for the first five input distributions from Table 3.2 and for growing sample size. It can be seen that SAMBA is more accurate for scarce data and large data. 73
3.6 The left graph shows a time series with one extreme outlier. If this outlier leads to detrimental consequences it is called a Black Swan. The right graph shows a fat-tailed probability distribution that can be used to model rare, extreme outliers statistically. 73
3.7 The left graph shows several alpha-stable distributions for different shape factors. The right graph shows a discontinuously truncated tail of a heavy-tailed distribution. 75
3.8 For tempered alpha stable distributions the length of the fat tail can be controlled through a cut-off parameter. This graph shows tails of varying truncation. 77
3.9 Example of a tempered alpha stable distribution fitted to fat-tailed Student-t samples. 78

3.10 Optimal collocation points for the Gaussian distribution (left) and Student-t distribution (right). The bottom graphs illustrate that optimal points can also be calculated for parts of distributions. 79
3.11 Output variation for a Gaussian input (left) and a heavy-tailed input distribution (right). 80
3.12 The figures show the relative error for the standard deviation calculated with SAMBA compared to 10^8 Monte Carlo samples. The four test functions are taken from Table 3.1. The graphs show that there is an accuracy limit for increasing order and that this accuracy limit is dependent on the number of samples available. 83
3.13 Magnitude of the relative error in the sparse Fourier coefficients based on data compared to analytic tensor product coefficients. It can be seen that the error is negligible for orders lower than five. Higher levels than 10, however, contain significant errors. 84

4.1 Response surfaces of a hyperbolic tangent function recreated from a low number of data points. The surface at the top was determined using the classic Padé approximation based on 16 Gaussian quadrature points. The surface in the middle is based on the new Padé approximation developed in this work based on 18 optimally sub-sampled points. 96
4.2 Demonstration of how an automatic trade-off between the leave-one-out error and the surface smoothness parameter is determined. 97

5.1 A multi-fidelity nested Gaussian sparse grid comprised of five DNS and 29 RANS solutions. 101
5.2 The multi-fidelity grid used for the weighted regression comprised of nine DNS and 81 RANS solutions. 103
5.3 (a) Computational domain used for the CFD simulation. (b) Illustration of how the hill shape is changed parametrically to create the assumed geometric uncertainty. 106
5.4 Instantaneous DNS visualisation of the vorticity modulus (for $||W|| = 5Uc/h$) coloured by the stream-wise velocity. 106
5.5 DNS validation for $Re = 700$ using horizontal velocity profiles. 107
5.6 Comparison of horizontal velocity profiles for different turbulence closure models. 109

5.7 Horizontal velocity streamlines visualisation comparing DNS (a,b,c,d) and RANS simulations (e,f,g,h). 110
5.8 (a) Velocity profile trend validation of k-ω SST with DNS at the location $x/h = 6$. (b) Skin friction trend validation of k-ω SST with DNS along the x-axis. 111
5.9 Probability distributions obtained from pure DNS and RANS data: (a,b) minimum velocity in bubble, (c,d) bubble height at $x/h = 2$, (e,f) bubble height at $x/h = 4$ and (g,h) reattachment location. 113
5.10 Response surfaces for the maximum back-flow velocity in bubble (left) and the reattachment point location (right). 114
5.11 Multi-fidelity PDFs (c-f) obtained for minimum velocity inside the bubble (left) and reattachment location (right) compared with pure RANS (a,b) and pure DNS (g,h) results. 115

6.1 Meshed 3D NACA0012 airfoil (left) and the computational domain used for the CFD simulation (right). 118
6.2 RANS streamlines illustrating the wing tip vortex. 118
6.3 Random samples used to create a stochastic positioning of the NACA0012 airfoil. 119
6.4 Variability of lift and drag given the uncertainty in the positioning of the NACA0012 airfoil. 120
6.5 F1 front wing CAD model and CFD simulation domain. 121
6.6 Illustration of the imposed rotational and translation uncertainty in the car body and front wing. 121
6.7 Random sample histograms and the directions of how they are assumed to affect the positioning of car model components. . . 122
6.8 Variability of lift, drag and lateral force acting on the car for the assumed position uncertainties. 123
6.9 Normalised measurements of geometric manufacturing deviations for various blade regions. 126
6.10 Schematic of the H2 airfoil, not in scale, with 6 sections and 2 points varied according to 6 different measurement sets and two PDFs given in Figure 6.9 and isentropic Mach number. . . 127
6.11 Histogram of pressure losses p_L of H2 profile for the 6 manufacturing uncertainty histograms and 2 assumed PDFs. 128
6.12 Schematic illustration of the experimental turbine, from Mazzoni [84]. 130

6.13 Representation of the experimental turbine showing the arrangement for shroud sealing, from Mazzoni [84]. 131
6.14 TBLOCK computational domain, from Mazzoni [84]. 131
6.15 Measured and predicted total pressure coefficient downstream of stator 2 & 3, from Mazzoni [84]. 132
6.16 Probability distribution used as input PDF for the clearance gap including optimal Gaussian collocation points. 133
6.17 The graphs show the effect of the assumed clearance gap uncertainty on the stagnation pressure for the three rotor and stator stages. 134
6.18 The graphs show the effect of the assumed clearance gap uncertainty on the yaw angle for the three rotor and stator stages. 135
6.19 Ares I-X vehicle schematic, Finite Element Model and measurement instrumentation positions, taken from Horta [59]. 138
6.20 Modal analysis of the FEM model of the Ares I-X, taken from Horta [59]. 138
6.21 Illustration of the concept of FEM model updating: the modal response of the model is updated through parameter changes until it agrees with the experimental modal response. 139
6.22 The left graph shows the model and experimental frequency response. The right graph shows the difference between them, calculated with the correlation coefficient used in the predictor corrector method. 140
6.23 Illustration of the simplified problem (called 'toy problem') used to test and demonstrate the efficiency of new UQ methods. . . 141
6.24 Result of the global sensitivity analysis for 64 parameters. It can be seen that the number of parameters that is important for model updating is smaller than the total number. 142
6.25 Improved calibration within the lower [0,200Hz] frequency domain using only active/sensitive parameters instead of all parameters. 143
6.26 Example data set for which data-driven Polynomial Chaos is more reliable than PDF fitting for rare event simulation. The Gaussian PDF smoothes out the rare event in the tail region. . 145
6.27 Empirical data showing that aircraft accidents occur with a heavy-tailed distribution, from Tseng [142]. 146

6.28 a.) Physical domain b.) Computational model c.) Temperature measurement point d.) Temperatures curves at different rotor cavity locations 147
6.29 Illustration of how the PDF 'moves' through the truncated time domain to simulate timing uncertainty. 148
6.30 SAMBA and MCM comparison for the variability of the transient temperature based on a.) Gaussian and b.) Student-t input distribution. 149
6.31 Illustration of a.) hot gas ingestion mechanism b.) the two gaps varied in the simulation. 150
6.32 Schematic overview of how SAMBA is applied to the hot gas ingestion problem. The graph on the left shows the used histograms and the created sparse grid. The graph in the middle shows the CFD domain and mesh. The graph on the right shows the output distribution obtained when using Gaussian input PDFs. 151
6.33 SAMBA grids obtained for a.) two Gaussian and b.) two KoBoL input distributions. 152
6.34 Contour plots obtained with a.) SAMBA and b.) Monte Carlo meta-model. 152
6.35 Output PDF for engine temperature when using fat-tailed input distributions. It can be seen that the critical temperature T_{crit} is clearly exceeded. 154
6.36 (a) Control volume (b) Flow features 156
6.37 CFD grid details in coolant/shock interaction region. 158
6.38 Comparison between CFD (upper) and experimental data (lower) for the deterministic reference configuration case. The numerical simulation greatly underestimates the lateral spreading of the cooling flow. 159
6.39 Film cooling effectiveness η_{aw} along the centreline of Figure 6.36 159
6.40 Flow streamlines and interaction with the shock. 160
6.41 The tornado vortex after the shock. 160
6.42 Temperature maps and streamlines between hole and shock for different x locations as shown in Figure 6.38. 161

6.43 The graphs illustrate how the shock moves through the many stochastic response surfaces of the thermal efficiency. Each of the nine points represents a CFD simulation. Until x/D=4.75 the response surfaces are smooth. Then, a shock (illustrated by two red lines) moves through the domain until x/D=4.85. . . . 162
6.44 Choice of the automatic response surface as a trade-off between leave-one-out error and smoothness indicator. 163
6.45 Three examples of automatically fitted Padé response surfaces before (1), at (2) and after (3) the shock location. The bottom row shows three PDFs to the x-locations before, at and after the shock. 164
6.46 Mean and standard deviation of the film cooling effectiveness (left) and sensitivity analysis of input parameters using Sobol Indices (right) 165

List of Tables

2.1 The Askey scheme: a table that shows which polynomials provide optimal Polynomial Chaos basis functions for a given probability distribution. 50
2.2 Relative error for mean and variance of data-driven PC compared to PCE based on fitted parametric distributions. 60

3.1 Multivariate test functions f used for 5 and 10 input random variables. 68
3.2 Probability distributions used for the input random variables ξ_i during testing with their parameters (e.g. μ, σ, a, b) and the used truncation interval $[l, r]$. 69
3.3 Relative error of standard deviation for five different input distributions compared to MC reference solution for the four test function in Table 3.1 and first five random variables of Table 3.2. 70
3.4 Relative error of standard deviation for ten different input distributions compared to MC reference solution for the four test function in Table 3.1 and first ten random variables of Table 3.2. 70
3.5 Test cases used in Constantine [26] to demonstrate errors in higher order Fourier coefficients. 82

5.1 Correlation test of DNS and RANS for separation bubble parameters. 111

6.1 Mean and standard deviation of lift and drag under the influence of the imposed random variables (in percent and normalised by the deterministic reference values). 119
6.2 Mean and standard deviation of the lift, drag and lateral force under the influence of the imposed random variables (in percent and normalised by the deterministic reference values). 123
6.3 Cascade geometrical parameters as defined in Montomoli [94] . 128

6.4 Sobol sensitivity indices showing the effect that manufacturing uncertainty has on the pressure losses dependent on region. It can be seen that the middle to end section of the suction side is most important. 129
6.5 Probability of critical hot gas ingestion 153

Declaration of Originality

I hereby confirm that I am the sole author of this thesis and the work described in it and that any ideas or information derived from the published and unpublished work of others have been acknowledged through standard referencing practices of the discipline.

Richard B.H. Ahlfeld
Imperial College London
30 September 2017

Copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Acknowledgements

Developing general methods for combining engineering data and simulations requires a supply of sophisticated models that can only be developed by experts in a respective field. Therefore, many people, companies and organisations contributed to this work by providing either simulations, data or industrial use cases. I would like to acknowledge and thank the following people and organisations for their contribution:

Francesco Montomoli, my primary Ph.D. supervisor, for his help with turbomachinery simulations. This work has drawn significantly on his experience in turbomachinery design and simulation. He also made it possible for me to visit both Stanford and NASA, as well as founding the start-up UQuant based on this work. On a personal level, I also want to thank him for his unlimited support and creative ideas - it made writing this Ph.D. a highly enjoyable task.

Enrico Scalas, my secondary Ph.D. supervisor, for giving me the idea to investigate the connection between statistical moments and arbitrary random data.

Luis Crespo and Sean Kenny for inviting me to work with NASA Langley in Virginia and for supervising my work on uncertainty in structural dynamic models of space launch systems.

Mauro Carnevale and Simone Salvadori for providing transonic CFD simulations to model film-cooling and shock interactions.

Cosimo Mazzoni and Budimir Rosic for providing CFD simulations and experimental data of a multistage gas turbine.

Sylvain Laizet for running more than 20 high-quality immersed boundary Direct Numerical Simulations and the use of his computing time on the UK supercomputer.

Gianluca Iaccarino for his help and insights on multi-fidelity uncertainty quantification and the invitation to Stanford.

Gianluca Geraci for setting up the software framework to run SAMBA sim-

ulations on the Certainty Cluster in Stanford.

Shahrokh Shahpar for providing fluid-structure interaction models of Rolls-Royce engines to perform rare event simulations.

Badia Belkouchi for her analytical work and survey on sparse Smolyak quadrature rules within the context of her Master thesis that helped to develop the SAMBA algorithm.

Audrey Gaymann for her help with transonic RANS simulations within the context of multi-fidelity simulations.

Fabio Ciampoli for his support of this research from within McLaren Racing and SpaceX.

Heng Xiao for his collaboration on optimisation using physics-informed machine learning.

Rolls-Royce Derby, McLaren Racing and CRIEPI for testing and providing feedback on the uncertainty quantification software package 'UQuant'.

Airbus, in particular, Sanjiv Sharma, for supporting uncertainty quantification research within the context of wing load analysis and landing gear design.

The Stanford CTR Certainty Cluster for providing access to their supercomputer.

The Knowledge Transfer Network and Innovate UK for their support and Matt Butchers for promoting this research.

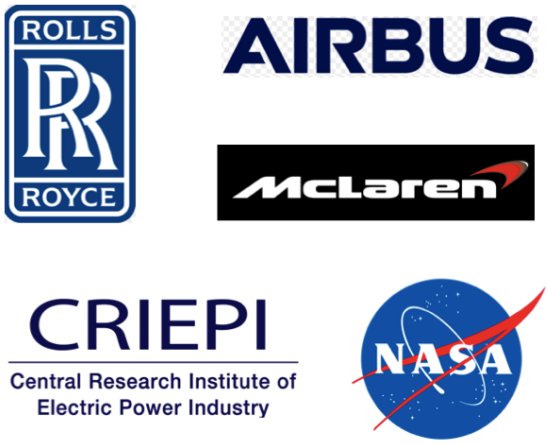
Imperial College London and the EPSRC research grant that allowed me to freely pursue my own ideas during the last three years.

Anna-Luise Wagner for providing feedback on all the research papers and presentations that lead to this work.

My family for their support.

Industrial Partners

The methods in this work were developed with the help of the following industrial and academic partners:



Publications

The work presented in this dissertation is supported by a range of publications in various peer-reviewed journals and conferences:

Journal Publications

Ahlfeld, R., Ciampoli, F. Pietropaoli, M. & Montomoli F, (Forthcoming) "Data-Driven Uncertainty Quantification for Formula 1: Diffuser, Wing Tip and Front Wing Variations." Under Review in *Journal of Automobile Engineering*

Ahlfeld, R., Laizet, S., Geraci, G., Iaccarino, G., & Montomoli, F. (Forthcoming). "Multi-Fidelity Uncertainty Quantification for Geometric Manufacturing Errors in Turbulent Flow." Under Review in *Journal for Flow Turbulence & Combustion*

Ahlfeld, R., Carnevale, M., Salvadori, S. & Montomoli, F. (July 2017). "Autonomous Uncertainty Quantification for Discontinuous Models Using Multivariate Padé Approximations." *Journal of Turbomachinery*

Ahlfeld, R., Montomoli, F., Scalas, E., & Shahpar, S. (January 2017). "Uncertainty Quantification for Fat-Tailed Probability Distributions in Aircraft Engine Simulations." *Journal of Propulsion and Power*

Ahlfeld, R., & Montomoli, F. (June 2016). "A Single Formulation for Uncertainty Propagation in Turbomachinery: SAMBA PC." *Journal of Turbomachinery*

Ahlfeld, R., Belkouchi, B., & Montomoli, F. (March 2016). "SAMBA: Sparse Approximation of Moment-Based Arbitrary Polynomial Chaos." *Journal of Computational Physics*

Mazzoni, C. M., **Ahlfeld**, R., Rosic, B., & Montomoli, F. (January 2017). "Uncertainty Quantification of Leakages in a Multistage Simulation and Comparison with Experiments." *Journal of Fluids Engineering*

Pietropaoli, M., **Ahlfeld**, R., Montomoli, F., Ciani, A., & D'Ercole, M. (June 2016). "Design for Additive Manufacturing: Internal Channel Optimization." *Journal of Turbomachinery*

Conference Proceedings

Ahlfeld, R., & Montomoli, F. (June 2017).“An Autonomous Uncertainty Quantification Method for the Digital Age: Transonic Flow Simulations Using Multivariate Padé Approximations.” *American Society of Mechanical Engineers Turbo Expo*

Ahlfeld, R., Laizet, S., Geraci, G., Iaccarino, G., & Montomoli, F. (October 2016).“Multi-Fidelity Uncertainty Quantification Using RANS and DNS.” *Proceedings of the CTR Stanford Summer Program*

Ahlfeld, R., & Montomoli, F. (June 2016).“A Single Formulation for Uncertainty Propagation in Turbomachinery: SAMBA PC.” *American Society of Mechanical Engineers Turbo Expo*

Mazzoni, C. M., **Ahlfeld**, R., Rosic, B., & Montomoli, F. (April 2016).“Uncertainty Quantification of Leakages in a Multistage Simulation and Comparison with Experiments.” *International Symposium on Transport Phenomena and Dynamics of Rotating Machinery*

Pietropaoli, M., **Ahlfeld**, R., Montomoli, F., Ciani, A., & D’Ercole, M. (June 2016).“Design for Additive Manufacturing: Internal Channel Optimization. *American Society of Mechanical Engineers Turbo Expo*

Salvadori, S., Carnevale, M., **Ahlfeld**, R. & Montomoli, M. (May 2016) “Stochastic Variation of the Aero-Thermal Flow Field in a Cooled High-Pressure Transonic Vane Configuration.” *European Turbomachinery Conference*

Invited Talks and Seminars

Ahlfeld, & R. Montomoli, F. (September 2017), "Combining Machine Learning and Simulations: The Future of Engineering Design." *International Talent and Entrepreneurship Week in Ningbo, China*

Ahlfeld, R. & Montomoli, F. (March 2017), "Robust Model Calibration of Advanced Manufacturing Robots under Uncertainty." *UK Knowledge Transfer Summit on Uncertainty in Industry in Bristol, UK*

Montomoli, F. & **Ahlfeld**, R. (October 2017), "Combining Data Science and Engineering Simulations: Scarce Data and Rare Event in Aerospace Engineering." *6TH CEAS Air & Space Conference Aerospace Europe in Bucharest, Romania*

Montomoli, F. & **Ahlfeld** (June 2017) Uncertainty Quantification in Turbomachinery, *ASME Turbo Expo Workshop in Charlotte, Virginia, USA*

Montomoli, F. & **Ahlfeld**, R. (October 2016), "Data-driven Uncertainty Quantification for Formula 1." *ERCOFTAC UM&Q in Industrial Analysis and Design in Ghent, Belgium*

Montomoli, F. & **Ahlfeld** (June 2016) Uncertainty Quantification in Turbomachinery, *ASME Expo Workshop 2017 in Seoul, South Korea*

Prizes and Recognitions

John & Frances Jones Prize for the greatest all-around contribution to Imperial College London in 2016 among all postgraduate students.

Elaine Austin Centenary Memorial Prize for the greatest contribution to a scientific field for supervising Badia Belkouchi's Master thesis on SAMBA.

Climate Launchpad UK Finalist for using UQuant to reduce CO2 emissions within highly uncertain electric grids in Africa.

Venture Catalyst Challenge Finalist for being one of the five best deep-tech ideas coming out of Imperial College London in 2016.

Reynolds Competition Finalist for integrating rare events into aircraft engine simulations.

Deep Science Ventures Fellowship to apply UQuant to GPS localisation uncertainty.

EPSRC Doctoral Prize Fellowship to integrate unsupervised machine learning into the uncertainty quantification software UQuant.

HelloTomorrow Challenge Winner for UQuant as one of the top seven new start-ups with unicorn potential in industry 4.0 in the world.

Preface

This dissertation summarises and generalises my work on Uncertainty Quantification (UQ) in engineering over the last three years. Most of the examples and methods have been published already, but I have also included bits and pieces that I had no time to publish, yet.

I was fortunate enough to stumble across the idea of data-driven Polynomial Chaos quite early on in my Ph.D. The credit goes to Professor Enrico Scalas, who, during a discussion in December 2014, mentioned that he had been wondering for a while whether there was not a beneficial connection between statistical moments and Polynomial Chaos expansions that could be exploited to propagate fat-tailed probability distributions. This early inspiration allowed me to develop a comprehensive theory based on statistical moments and apply the resulting methods to many different fields.

The readership I had in mind when writing this thesis are engineers from the aerospace, automotive, and civil engineering industry. In my experience, such engineers are quite proficient in the use of advanced computational modelling techniques and also familiar with higher mathematics. However, they lack a detailed knowledge of statistics and are consequently a bit wary of probabilistic methods. They have not been introduced to the merits of systematic uncertainty quantification techniques. I hope that this thesis and its various examples will help to demonstrate the benefits of it.

Finally, I would like to remark that I have never really liked the terminology 'uncertainty quantification' for this field of research. My experiences at conferences and with reviewers have shown that this terminology generally just seems to convey the idea that uncertainty quantification tries to solve physical problems using statistics, or is some magical number crunching to quantify the unknown. The field investigates the effect of natural variability on computer models, seeks efficient ways for parameter sensitivity studies and finds new strategies to optimise technical designs for a broader range of conditions that take into account nature's variability. In my opinion, it would, therefore, be more appropriate to call the field 'data analytics for engineering' or 'risk quantification'. As a matter of fact, the term 'robust design' seems to have become more popular recently and it is usually the concept of robust design that first convinces people of the merits of UQ.

This thesis will give the reader a detailed and all-around overview of how data-driven uncertainty analytics can be helpful to address various challenges arising from uncertain measurement data and modelling conditions in engineering simulations. I have also added detailed explanations and simple ex-

amples so that this work can serve as a tutorial for readers who wish to find out more about the state-of-the-art and advantages of using rigorous uncertainty quantification methods. Parts of this work were carried out during research visits to the Stanford Center for Turbulence Research in California and NASA Langley Research Center in Virginia and in collaboration with industrial partners, such as McLaren Racing, Rolls-Royce, Airbus, and CRIEPI.

1 Introduction: Uncertainty in Engineering Simulations

There are known knowns; there are things we know we know. We also know there are known unknowns; that is to say, we know there are some things we do not know. But there are also unknown unknowns,
– there are things we do not know we don't know.
D. Rumsfeld

This thesis introduces several new data-driven uncertainty quantification methods for engineering simulations. Therefore, it makes sense to begin by defining what is meant by uncertainty quantification: the application of statistical methods to predictive models in order to find the likelihood of simulated outcomes when taking into account the variability/uncertainty of parameters. The quote at the beginning of this thesis is often used to differentiate between the various problems that are caused by uncertainty. It can be depicted as a **matrix of knowledge** with four quadrants [91] as shown in Figure 1.1.

Typical engineering work is done in the first quadrant of **known knowns**. New engineering products are designed using well-understood physical models and simulation methods. The deficits of these models and simulations are generally known. These shortcomings fall into the second square, the sector of **known unknowns**. Understanding the effect of known unknowns is what requires experience in modern-day engineering. For example, engineers know that manufacturing tools and computational simulations can only be trusted within certain error margins. Therefore, safety factors are used to ensure reliability [43]. The third sector, **unknown knowns**, usually refers to cases for which data exists, but it is either not rigorously used or misinterpreted. This is the area most often addressed by scientific uncertainty quantification methods. For most industrial cases, data is available to estimate the volatility in measurements or during operation. However, it is not systematically evaluated. It is still not very common to rigorously quantify and document all known variabilities in measurements or models. Various governmental research agencies like the British National Physical Laboratory

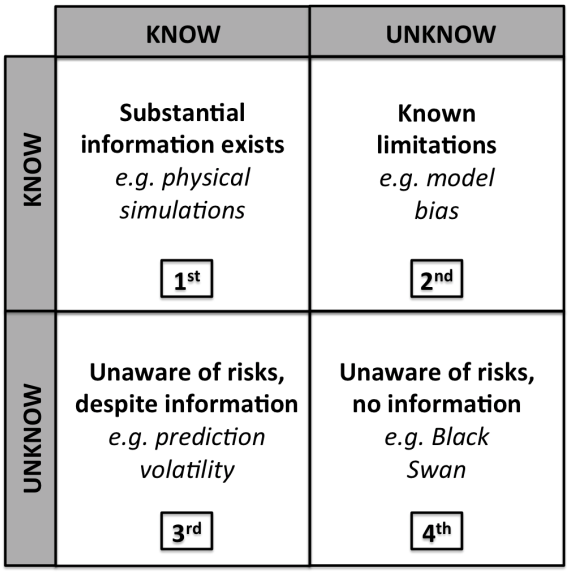


Figure 1.1: Information about a technical problem can be displayed as a matrix of knowledge with four quadrants. Mathematical uncertainty quantification takes place in the third quadrant.

therefore strongly advocate best practice guidelines to deal with uncertainty [27, 45]. The treatment of the fourth sector, **unknown unknowns**, is most difficult. However, there are methods to address problems in this sector, like for example the p-box approach [132]. The general idea is to explore the parameter space in a systematic way to search for possible complications and quantify their risk. Particularly important is the search for rare events with disastrous consequences, so-called Black Swans.

This work proposes methods for all three sectors of uncertainty. Most of the work is devoted to the third sector: quantifying the volatility of engineering predictions by efficiently incorporating available experimental or operational data. However, one new method to find Black Swans and one quite general mathematical approach to remove model-form uncertainty are also introduced.

1.1 The Science of Uncertainty Quantification (UQ)

Uncertainty quantification is a multi-disciplinary field that combines methods from statistics and scientific computing (numerical methods). From a mathematical point of view, one differentiates between two types of uncertainty: **aleatory** and **epistemic** [104]. Uncertainty is called aleatory, if it can be quantified using statistical methods. For example, the stochastic variation when repeating an experiment creates aleatory uncertainty. Aleatory uncer-

tainty quantification studies are used to understand the reliability or failure probability of a design. Epistemic uncertainty is due to insufficient knowledge about the physical behaviour of a system. For example, all models in Computational Fluid Dynamics (CFD) used to simplify the highly complex behaviour of flow turbulence create epistemic uncertainty.

1.1.1 Uncertainty Propagation

A large proportion of the research done in uncertainty quantification is on efficient uncertainty propagation methods. With such methods, fluid or structural problems are treated using partial differential equations whose domains and coefficients are random parameters or processes. For instance, wind loads on bridges can be described by temporal stochastic processes [114] that affect a structural Finite Element Model (FEM). Geometric manufacturing variations of a turbine blade can be described by a series of spatial random variables that affect a Computational Fluid Dynamics (CFD) simulation [13]. The exploration of the reliability of a simulation model using probabilistic methods is referred to as **non-intrusive** uncertainty quantification in the literature. The term non-intrusive means that the model is not modified but used as a **black box**. The basic procedure of a non-intrusive uncertainty quantification process is demonstrated in Figure 1.2. The first step is to describe uncertain-

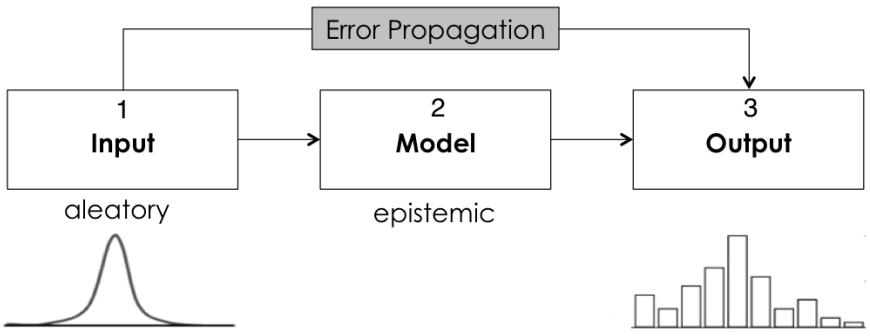


Figure 1.2: Non-intrusive uncertainty propagation techniques can quantify the effect of input variability on model predictions without modifying a model.

ties in the design space with a probabilistic framework. For example, model parameters can be replaced with random variables [78], or a domain region can be defined as a random field [35, 33]. In the second step, the computational solver is run repeatedly for random samples drawn from the probabilistic input distributions. The third and last step is to perform a statistical analysis

of the output quantities to determine the likely range of operating conditions.

Non-intrusive uncertainty quantification methods are sometimes frowned upon by engineering specialists. To them, the reduction of uncertainty is only feasible by improving a physical model. However, there are many processes in industrial manufacturing or in nature that produce **irreducible** uncertainties. For example, the thrust produced by a combustion chamber is highly stochastic. This is referred to as firing noise. If irreducible uncertainties are present, UQ methods can be used to quantify their effect on model output quantities and estimate the reliability of a product. For instance, temperature variations can be used as a proxy for turbine blade life prediction. In addition, UQ methods can be used to identify which components of a model cause a deviation from experiments for problems with a high number of parameters. This can help engineers to establish which parts of a problem need further exploration [104]. Naturally, there are also uncertainties that are **reducible** by improving modelling and measurement techniques. For example, geometric manufacturing errors can be reduced by using a finer and more advanced manufacturing technique.

Only non-intrusive UQ methods will be used in this thesis. They are more suitable for industrial applications at the current state-of-the-art. For one, most companies use complex multi-physics simulation software. Intrusive methods cannot easily be integrated into industrial software packages because they require the basic solvers to be modified [110]. Non-intrusive methods may be computationally more expensive than intrusive methods, but they are trivial to **parallelise**. Once the first step of Figure 1.2 has been performed and the available random data has been analysed, the physical model has to be computed in particular configurations to determine the effect of the input data. These configurations are computationally not connected to each other so that they can be submitted to different processors, as illustrated in Figure 1.3. Predominantly **spectral methods**, like Polynomial Chaos, are used for uncertainty propagation in this thesis. An extensive description of spectral methods can be found in Le Maıtre [70].

At the current state-of-the-art, most industrial simulations are highly complex and non-linear. For such models, small variations of uncertain parameters can lead to large variability in model outputs. If such behaviour is discovered during a UQ study, it is often an indication that the prediction of a simulation cannot be trusted. Performing uncertainty quantification studies is therefore of great importance to validate the reliability of a computational prediction.

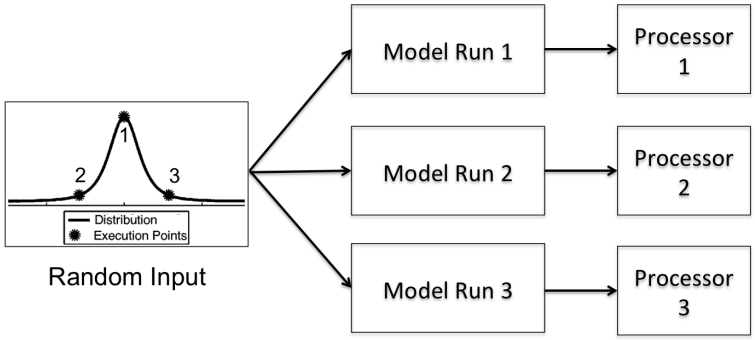


Figure 1.3: Non-intrusive uncertainty quantification methods are trivial to parallelise, because individual model runs/samples are never connected.

1.1.2 Advantages for Industry

Virtual design through modelling and simulation is extremely attractive for the engineering industry. Physical simulations help to improve product design, efficiency and competitiveness. They reduce the number of required real-life tests, allow to design and test products with minimal hardware prototyping, enhance speed to market and help to avoid costly production mistakes. With the rise of data science, virtual designs are transformed into more than mere idealised designs: they become **Digital Twins**, digital replica that combine data and simulations to mirror their real-life counterparts [50]. Uncertainty quantification methods can be used to extend the capabilities of virtual prototypes even further:

- It is widely known that UQ methods can be used to calculate error bars for experiments and computations. It is less known, however, that UQ can help to identify which assumptions, parameters or model components cause **discrepancies between simulations and experiments**.

- UQ methods can be applied to incorporate the effects of **manufacturing errors** and **in-service degradation** into physical simulations. By incorporating manufacturing and usage data product volatility can be described. In addition, the physical model can be used to predict the probability of failure of a manufactured product. Reliability predictions obtained through UQ are more accurate than pure statistical predictions because they complement the statistical data with high-fidelity physical information.

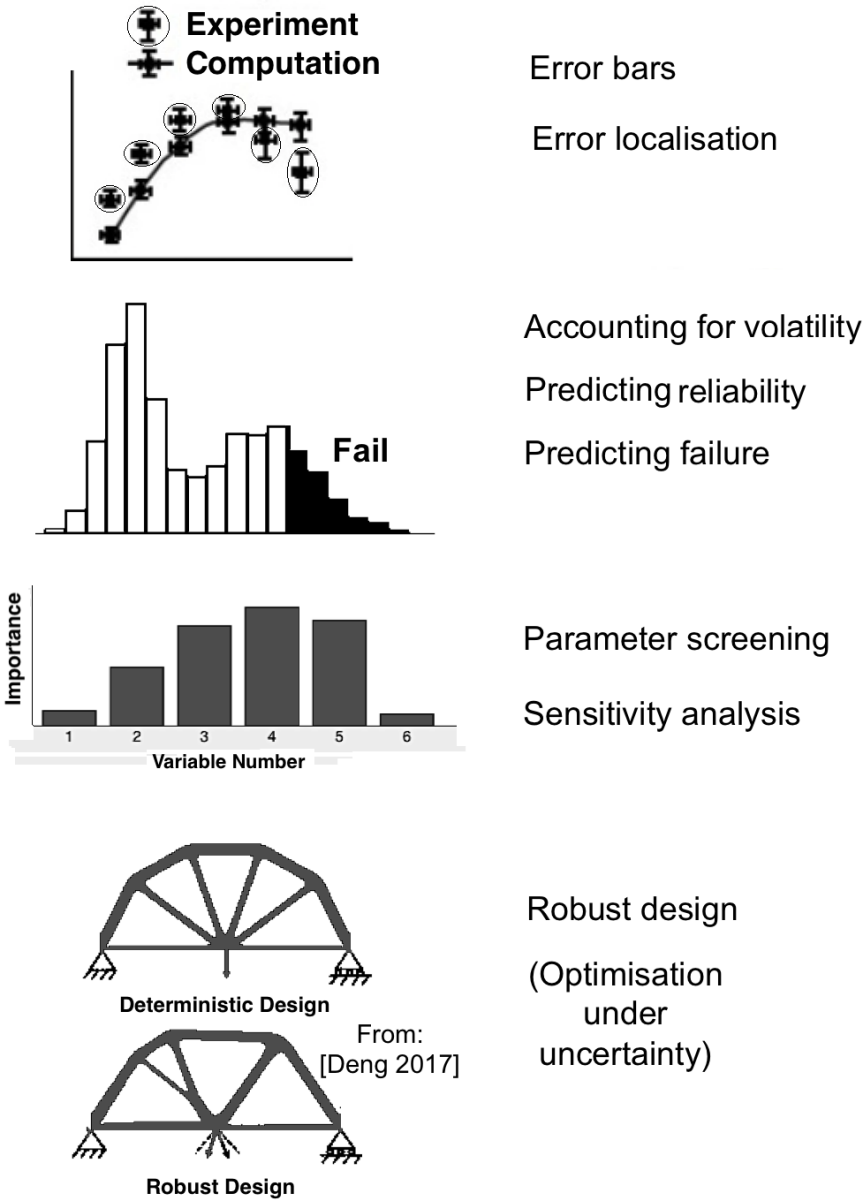


Figure 1.4: Uncertainty quantification methods have many advantages to support engineering design in addition to providing more rigorous error bars: predicting natural variability, finding important parameters and making designs more robust under variable operating conditions.

- UQ methods can also be used to identify areas of designs simulations that are most important or for which the safety margins are too low/high. This is called **parameter screening**, or **sensitivity analysis** [138]. Generally, lowering unnecessary safety margins helps to lower costs without losing efficiency. Increasing safety margins that are too low helps to improve safety.

- Model optimisation taking uncertainty into account is called **robust design**. The robust product is optimal for the entire range of conditions that are not known.

A summary of these advantages is illustrated in Figure 1.4.

Uncertainty quantification methods are closely related to **supervised machine learning** [65, 97]. Probabilistic reasoning is the main tool for the machine to make sense of real-world data. Thus, one way to look at UQ methods is as hybrid approaches that complement supervised machine learning with results from advanced physical simulations. Similar to machine learning, the probabilistic reasoning is often divided into two prevailing approaches **Bayesian** or **Frequentist**. Which one is preferable is a matter of opinion and still a heated debate in statistics [97]. It is generally possible to obtain the same results with both methods. The main difference is that Bayesians consider probability as an indication of the plausibility of a proposition or a situation. Hence, probability (or the priors) can be freely assigned. In the frequentist interpretation, this is not allowed. Probabilities are dictated by occurrences, like measurement outcomes or product performance statistics. Within this work, no Bayesian methods will be used. In the author's experience, Frequentist reasoning is more appealing to engineers who can better understand probabilities that are inferred from data. However, this is by no means intended as a dismissal of Bayesian methods, which are powerful mathematical tools with many advantages.

1.2 Challenges for Industrial Application

While the mathematical theory of uncertainty quantification has been developed greatly in the last decade, many methodologies are still difficult to apply to actual engineering cases. An example for a problem that is often overlooked is that many UQ codes require that the simulation model is executed in optimal collocation points as shown in Figure 1.3. For industrial applications, however, this can be problematic. Design studies are rarely conducted in the

optimal points of a UQ study. Similarly, computer simulations cannot always be performed in such optimal points, either due to convergence issues with the solver or simply because certain configurations were already performed by the simulation engineer, and it makes economically no sense to perform more time-consuming simulations. Already generated data should always be recycled for an uncertainty quantification study.

The main goal of this work is to develop a new theory of uncertainty quantification that better fits industrial needs. Therefore, core issues in industry were identified in collaboration with industrial partners and the British Knowledge Transfer Network's (KTN) focus group on uncertainty. These problems are: **limited data and statistical assumptions** and **limited time and the curse of dimensionality**, estimating the likelihood of **rare events and Black Swans**, dealing with models that contain **discontinuities** and mitigating the effect of **epistemic model-form uncertainty** on the results of a UQ study.

1.2.1 Limited Data and Statistical Assumptions

Obtaining **measurement data** is generally costly. In the aviation industry, it is common to repeat measurements five times. While this is enough to estimate mean and variance as shown in Figure 1.5, it is generally not enough to calculate statistical failure probabilities or perform a rare event analysis. Results illustrating the connection between data availability and convergence behaviour are presented in Chapter 3. Limited measurement data generally

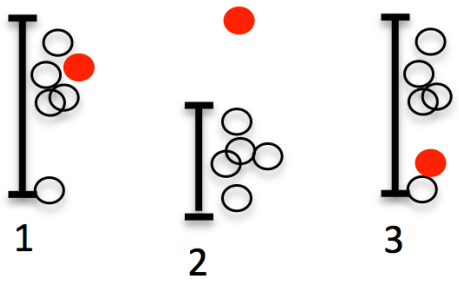


Figure 1.5: Illustration of three scenarios of variance estimation: 1) Overestimation 2) Underestimation 3) Correct Estimation. The black dots represent available data. The red dot is a yet unknown data point.

mean that the engineer in charge needs to make assumptions about the more general behaviour of the system. This can be problematic in statistics.

The majority of methods in the field of uncertainty quantification assume that a sufficiently high number of samples is available to correctly infer the

probability distributions underlying the input random variables. However, for industrial cases, there is often not enough data available due to limited time and cost issues. The construction of a parametric PDF for real cases consequently involves a lot of assumptions and subjectivity. The errors caused by making the wrong assumptions for input probability distribution are often underestimated. Figure 1.6 shows the impact that the choice of distribution has on the results of a UQ study performed for the 2D CFD simulation of a heat exchanger. The CFD simulation was performed with QuickerSim, a Matlab integrated CFD code[1]. Both the Gaussian and the Uniform distribution are

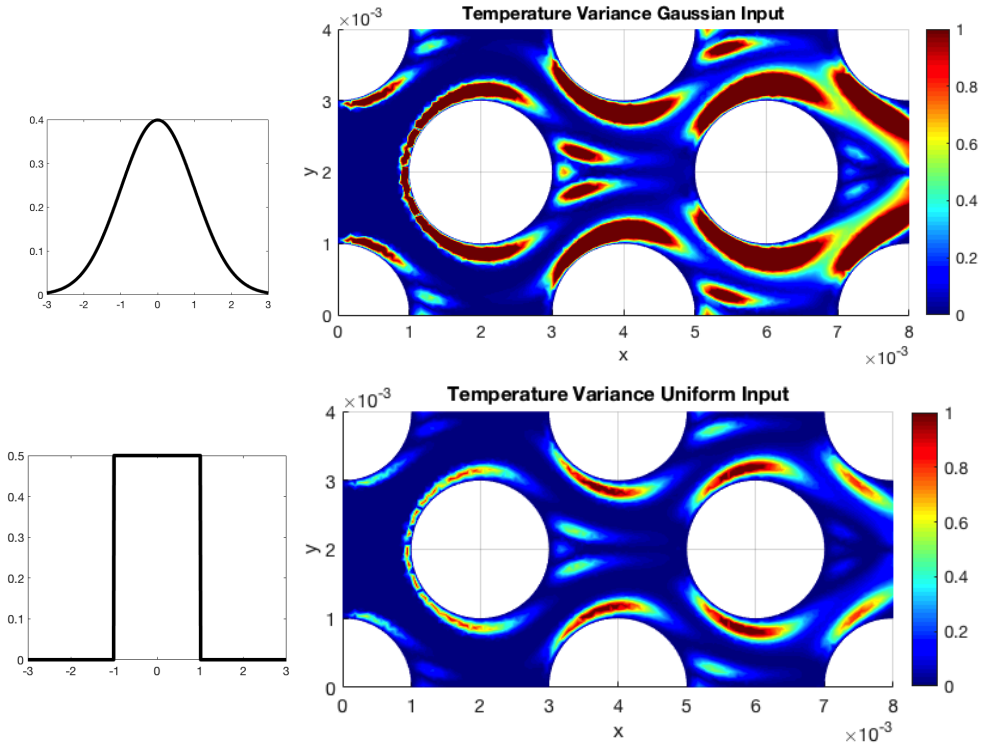


Figure 1.6: This figure shows the effect of different probability distributions on the temperature variance field of a heat exchanger CFD simulation. Even though the range of the distributions is similar, the Gaussian PDF creates significantly higher variance.

acceptable assumptions for scarce data. The variance in the temperature distribution is however quite different.

Another problem with the assumption of a specific distribution is that the convergence behaviour of Polynomial Chaos methods[2] is strongly dependent

[1] See https://quickersim.com/cfdtoolbox

[2] Polynomial Chaos methods are efficient spectral uncertainty quantification techniques. They are explained in Section 2.2.

on the correct choice of distribution [107]. The reason why an incorrect choice of distribution inflicts a particularly large error for Polynomial Chaos methods is that PCE defines a closed solution space based on the inputs. If the inputs are wrong, the actual solution will be outside the defined space. This conceptual error cannot be corrected by increasing the order. The data-driven Polynomial Chaos method that will be introduced in Chapter 2 never makes wrong assumptions about the solution space and therefore avoids such errors and converges always to the correct solution.

The human mind is trained to recognise patterns. In the field of statistics, this can be misleading. Let's take for example the random process in Figure 1.7 and assume that only the data until the division line has been observed. The process seems to be periodic. When looking at the actual realisation after the knowledge line, however, one can see that this was a false assumption. En-

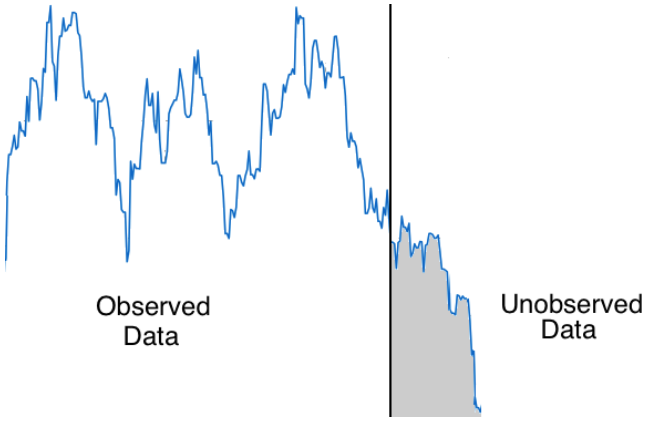


Figure 1.7: Patterns in random processes can be misleading. The stochastic process depicted was only observed until the vertical line evoking the impression that it is periodic.

gineers are better trained at recognising patterns than they are in dealing with stochastic processes. There is therefore a need in UQ for reliable **automatic statistics**.

The new data-driven Polynomial Chaos algorithm introduced in Chapter 2 provides a single solution to the problems of how to deal with scarce data, avoid wrong statistical assumptions as well as performing statistics automatically.

1.2.2 Limited Time and the Curse of Dimensionality

The maximum **acceptable time frame** for a complete UQ analysis in the industry is about two weeks. Ideally, the results of a simulation should be available within a few working days. Receiving reliable UQ results on a standard

workstation this fast is only possible with a very efficient and intelligently automated uncertainty quantification method and a reasonably fast model. However, available computing resources, in particular, access to high-performance computing, is becoming more and more common throughout many industries. Uncertainty quantification methods will greatly benefit from this development in the future because parallelising a UQ analysis is trivial. Individual simulations are always independent and can be submitted to different machines. Nevertheless, at the current state of the art, efficiency is still of paramount importance for UQ. Most models are computationally too expensive to allow thousands of simulation runs.

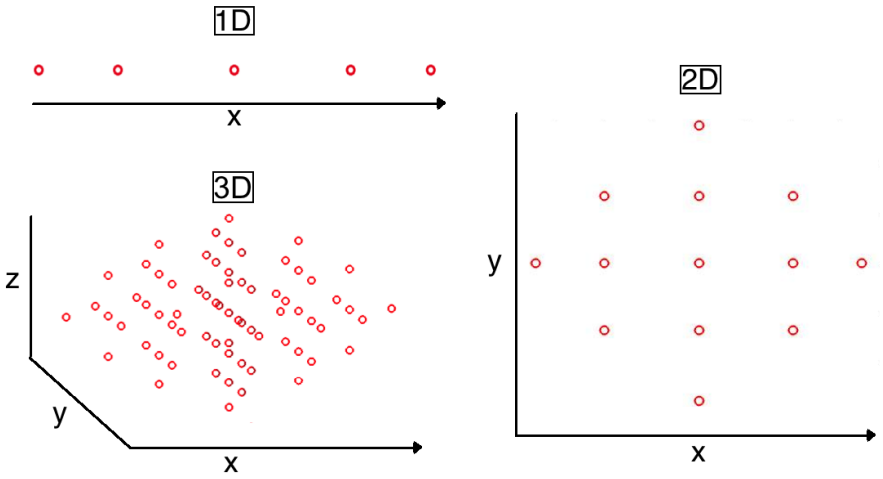


Figure 1.8: Illustration of the curse of dimensionality. The number of model evaluations rises quickly from one to three dimensions.

One of the major problems in uncertainty quantification is the so-called **curse of dimensionality** [40]. It is a problem that UQ shares with many other mathematical disciplines that need to analyse high-dimensional spaces, like numerical analysis, optimisation or machine learning. Although the applications may vary, the core problem is always the same: the amount of data needed to obtain reliable results grows exponentially with the number of dimensions With three samples per dimension, the number of samples required for twenty dimensions is already higher than 3.4 Billion samples. For up to three dimensions this grows can still be visualised. Figure 1.8 shows the number of collocation points that are needed to integrate a Gaussian distribution for one to three dimensions x, y and z.

The Polynomial Chaos methods investigated in this work are no different. For this reason, Smolyak's rule the most widely known rule to alleviate to curse

of dimensionality is modified in Chapter 3 for data-driven Polynomial Chaos. The result is a highly general new sparse quadrature generator that can be used throughout any discipline that is affected by the curse of dimensionality.

1.2.3 Rare Events and Black Swans

Uncertainty quantification can be used to perform risk analysis for a physical system using its virtual simulation model. Most popular in engineering is risk analysis using Gaussian distributions. However, they can give a false sense of security. The assumption of normality removes the possibility to account for rare events from the model. Events that are further away than five or more standard deviations from the mean are extremely unlikely for the normal distributions - more unlikely than they often are in reality. If such rare events have catastrophic consequences, they are called a Black Swan, as suggested by Taleb [140]. Black Swans have so far been studied mostly in mathematical finance to explain catastrophic market crashes. However, Montomoli [91] showed the importance of **fat-tailed distributions** for the prediction of rare events in engineering applications for a turbine model a few years ago. The reason for that can be seen in Figure 1.9. When considered with the mere eye, fat-tailed probability distributions look almost identical to Gaussian distributions. However, when considered closer, as in the right graph, it can be seen that while unlikely rare outliers are possible. While physical laws and

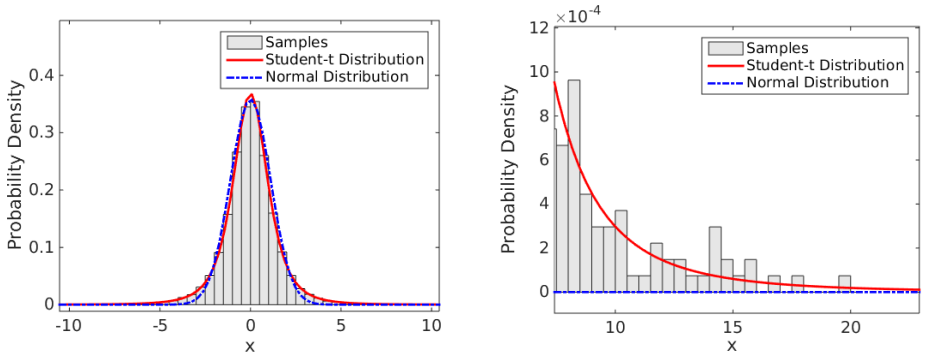


Figure 1.9: The visual difference between a Gaussian and a fat-tailed distribution is small. Only zooming into the tail region reveals that rare events are significantly more likely for fat-tailed distributions.

even aleatory input conditions can be discovered and incorporated into the simulation process, some unknown factors inevitably remain outside the scope of simulation. For this reason, the fourth quadrant of the matrix of knowledge in Figure 1.1 is considered as the most unexplored region and most of the pub-

lished literature deals with the third quadrant only. Especially for rare event analysis there is still a lack of efficient uncertainty propagation methods. In Chapter 3.3 a new technique will be introduced to explore the fourth quadrant with the goal to limit the effect of variations in the input distributions on the output and design systems which are less prone to be affected by rare events. The method used employs the algorithm of the data-driven Polynomial Chaos method to find optimal polynomials for fat-tailed distributions. This provides a fast and efficient way to increase the probability of tail events within a more **conservative risk analysis** that still **preserves mean characteristics** unlike other methods suggested for rare event simulation. Usually, the tails are treated separately from the bulk of the distribution, see [125].

1.2.4 Models with Discontinuities

Some computational models contain instabilities, bifurcations or sharp gradients across a certain region of their computational domain [130]. An important example from supersonic aircraft engineering or turbomachinery is **compressible flow simulation**, where discontinuities can be encountered if uncertainties of local flow quantities in the shock region are to be quantified. Overall, uncertainty quantification for discontinuous models lacks behind methods for continuous problems in efficiency [53]. One reason is that global Polynomial Chaos expansions inherently assume that the random domain is smooth. When continuous Polynomial Chaos expansions have to deal with a discontinuity across the parameter space, the expansion will either be insufficiently accurate at low orders or suffer from the **Gibbs phenomenon** at high orders.

The left graph in Figure 1.10 shows a system of shocks in transonic flow within a gas turbine. The right graph shows a connected model response surface where a shock divides the domain. While this response surface is well-resolved, it is often difficult to distinguish between strong gradients and discontinuities [128] when only sparse data points are available. Unfortunately, this is the most common situation in engineering. Usually, only scarce data points from test simulations in manually determined locations are available. To address this problem, a new global discontinuous non-intrusive Polynomial Chaos method is developed in Chapter 4. The basic idea of the method is to use a ratio of two Polynomial Chaos expansions. The poles in the denominator expansion can then be used to 'filter' out the discontinuity. While this approach was developed years ago [19], this work adds an important novelty: an algorithmic modification that allows determining a from **arbitrary simu-**

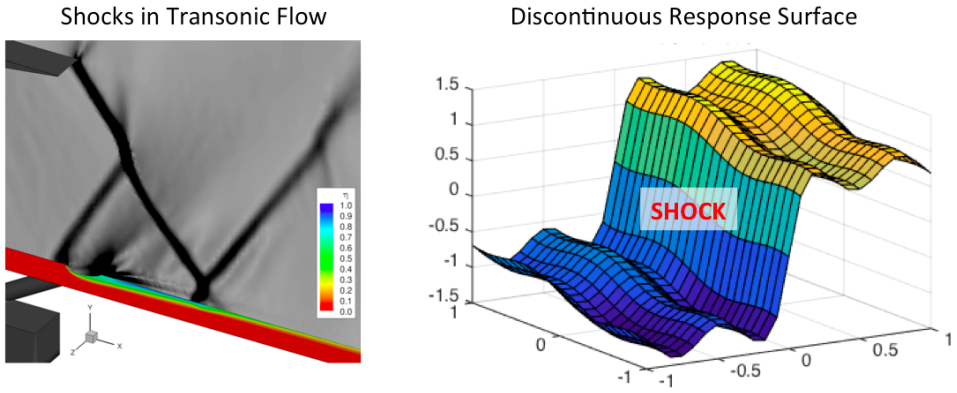


Figure 1.10: The left graph shows film-cooling/shock interactions in a transonic CFD simulation. The right graph shows a discontinuous response surface for a local flow parameter that is affected by a shock in the CFD domain. In order to accurately resolve the discontinuity in the response surface, special uncertainty quantification methods are needed.

lation points. Previously, only predefined mathematically optimal Gaussian quadrature points could be used. The modification makes the method more attractive for industry because it provides a method that requires very few points and allows to recycle already performed test simulations.

1.2.5 Epistemic Model-Form Uncertainty

All models are mere approximations of reality. Some degree of epistemic model-form uncertainty is always present. The problem is that purely aleatory uncertainty propagation studies often ignore the effect that model-form uncertainty has on the results. Similar to the errors in purely physical simulations, also the probabilities derived from a model that contains epistemic uncertainty are biased. For complex non-linear models, the effect of model-form uncertainty on the shape of a PDF can be very difficult to estimate.

To give an example, Figure 1.11 shows two probability distributions that show the variability of the location where a flow reattaches after a hill. Both results were calculated using CFD codes. However, it can be seen how the shape and range of the PDFs are quite different. For the high-fidelity and more accurate Direct Numerical Simulation (DNS) on the left-hand side, the mean flow reattaches at about $x/h = 4$ and the distribution is skewed to the left with a standard deviation of about $x/h = 1$. For the much cheaper

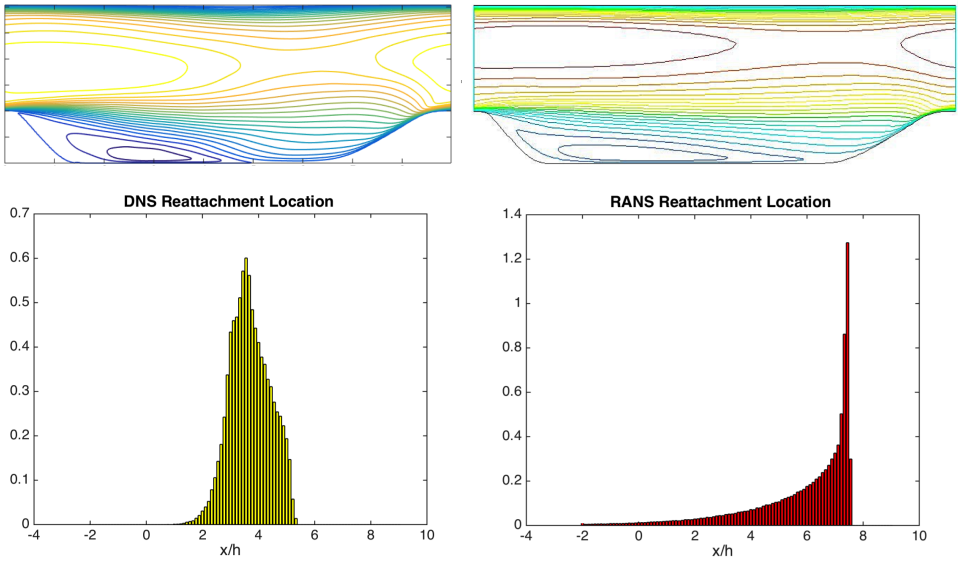


Figure 1.11: The top graph show the streamlines of flow over periodic hills for a DNS (left) and a RANS (right) solver at $Re = 700$. Below are the corresponding PDFs for the reattachment location calculated using DNS and RANS. The large difference between the PDFs illustrates the effect of epistemic model-form uncertainty on UQ results.

but also significantly less accurate Reynolds Averaged Navier-Stokes (RANS) simulation on the right-hand side, the mean flow reattaches at $x/h = 6$ and the distribution is highly skewed to the right with a standard deviation of about $x/h = 3$. This example demonstrates the significant effect of the model-form uncertainty on a stochastic prediction obtained through UQ.

The study performed for Figure 1.11 is continued in Chapter 5. It investigates the effect that RANS turbulence modelling errors have on an uncertainty quantification study by comparing UQ results from RANS simulations with pure DNS (the most accurate results obtainable in CFD) to analyse the effect of model-form uncertainty. The DNS UQ study is quite unusual, because a tremendous computational cost is already needed to perform a single DNS: even a simple DNS simulation, like the flow over periodic hills at the very low Reynolds number of $Re = 700$, takes about three days on a supercomputer with 5000 cores. The performed study features ten such DNS simulations at $Re = 700$.

In addition to merely analysing the effect of uncertainty, also two generic **multi-fidelity** strategies are explored to counteract model-form uncertainty by using a combination of higher and lower fidelity models. The advantage

of multi-fidelity methods for UQ is that due to the generality of non-intrusive multi-fidelity methods they can be applied to any kind of physical industrial problem for which higher and lower order models exist. No highly advanced physical knowledge is needed by the UQ engineer to improve the actual modelling. For CFD, the idea is that a combination of RANS and DNS provides a trade-off between the computational speed of RANS and the higher accuracy of DNS.

1.3 Thesis Statement and Contributions

The most important contribution of this work is a new, numerically more efficient **data-driven Polynomial Chaos** algorithm and a data-driven sparse Gaussian quadrature rule generator called **SAMBA PC** (short for Sparse Approximation of Moment-Based Arbitrary Polynomial Chaos). Within this work, SAMBA PC will also be referred to as **SAMBA** only. SAMBA was designed to better deal with real-sized industrial problems in a non-intrusive and highly efficient way. To date, it is the only data-driven, multi-dimensional UQ method that can propagate available random data without making erroneous statistical assumptions. Its main advantages can be summarised as follows:

1. It can provide more accurate statistical predictions for output quantities than probability distribution fitting, if only **scarce data** (or **large non-parametric data**) sets are available.

2. It is **easy to implement** as well as to **parallelise** because it is formulated in matrix notation. This notation is also **more comprehensive** in providing all necessary information for numerical quadrature and sensitivity analysis within a single framework, - unlike arbitrary Polynomial Chaos.

3. It can handle **exotic distributions** for which no closed-form solution is available.

4. It can automatically and optimally deal with **all PDFs, histograms** and even **multi-modal distribution** (as long as they are determinate in the Hamburger sense)

5. It can efficiently deal with a high number of **mixed distributions**, for example, to create a single optimal quadrature rule that comprises both PDFs and non-parametric random data.

6. It can be used for efficient **global sensitivity analysis** that does not require gradient information for multidimensional and complex problems.

7. It is **non-intrusive**. There is no need to modify any computational model.

In addition to a new data-driven Polynomial Chaos algorithm and the first multi-dimensional data-driven Polynomial Chaos algorithm SAMBA, this work provides three novel contributions to the scientific field of uncertainty quantification:

- The first derivation in the open literature of how to create Polynomial Chaos Expansions for **fat-tailed distributions** to perform efficient **rare event simulation** for computationally expensive models.

- The first least squares discontinuous Polynomial Chaos method. In this work, the mathematical formulas for global discontinuous response surface approximations (multivariate Padé) are modified so that a **a single discontinuous global model** can be determined from **arbitrary collocation points**. Previously, such an approximation was only possible by using a fairly high amount of mathematically value-restricted optimal Gaussian quadrature points. This, however, restricts the application of global approximations for industrial purposes, where designs are performed for manually chosen points.

- A UQ study based on DNS was performed to investigate how **RANS and DNS** can be combined within a more general, model-independent **multi-fidelity framework** to remove turbulence modelling errors from robust optimisation or uncertainty quantification problems.

The majority of examples in this work use CFD simulations. A comprehensive work detailing only uncertainty quantification for CFD was written by Najm [99]. However, various examples also feature structural mechanics simulations. A well-written overview of methods in structural mechanics as well as CFD is [129]. A general introduction to advanced uncertainty quantification methods can be found in [136].

The main emphasis of this work is on Polynomial Chaos methods - the most popular method in the scientific community at the moment. The main new method developed in this work is a multi-dimensional data-driven Polynomial Chaos method. Various other sampling methods, like **Meta-Model Monte**

Carlo and **Importance Sampling** are used in various examples for comparison. The aim of this work is to improve the applicability of uncertainty quantification methods for industrial applications. The work was undertaken in collaboration with industrial and governmental partners from multiple disciplines: Formula 1, turbomachinery and rocket structural mechanics of the NASA Space Launch System.

1.4 Thesis Content

For the aforementioned reasons the following building blocks are needed for a comprehensive UQ framework for industry: a method to deal with scarce experimental data to avoid statistical errors, a method capable of dealing with many input parameters to mitigate the curse of dimensionality, an efficient method to perform rare event analysis for complex computational models, a method to deal with high gradients and shocks, and a method to reduce high model-form uncertainty. Each aspect, although inherently interconnected, will be covered in a separate chapter for the sake of clarity.

Chapter 2 introduces a new algorithm for data-driven Polynomial Chaos and firstly demonstrates its superiority for problems where only scarce data is available. **Chapter 3** introduces a sparse quadrature rule based on data-driven Polynomial Chaos called SAMBA that can be applied to solve any multi-dimensional integration problem. **Chapter 3.3** shows how the presented data-driven PC algorithm can be adapted to create an efficient new framework for rare event simulation. **Chapter 4** introduces a new least squares method to propagate uncertainty through models containing discontinuities that works with arbitrary collocation points. **Chapter 5** describes how to reduce epistemic uncertainty with a completely non-intrusive multi-fidelity framework suitable for industrial simulations. **Chapter 6** demonstrates data-driven Polynomial Chaos and SAMBA for practical engineering problems from the fields of Formula 1, turbomachinery and rocket design. In addition, examples from the field of turbomachinery are provided to demonstrate the new approaches for rare event simulation and uncertainty quantification for models containing shocks.

2 New Algorithm for Data-Driven Polynomial Chaos

Polynomial Chaos methods are efficient uncertainty propagation techniques to determine the effect of aleatory input uncertainties on complex computational models [41, 61]. They are used widely because of their computational cost, which is orders of magnitude lower than for sampling based approaches. This speed-up is mathematically achieved by expanding continuous input Probability Density Functions (PDF) into a basis of optimal orthogonal polynomials. The coefficients of these optimal polynomial basis functions are then determined in two steps. First, the model is run as a black box (non-intrusive) for various parameter values. Second, the coefficients are determined using either linear regression or quadrature based approaches on tensor-product or sparse grids. A comparison of these two Polynomial Chaos approaches can be found in Eldred [40].

2.1 Need for Data-Driven Polynomial Chaos

All available Polynomial Chaos methods assume that a sufficiently high number of samples is available to correctly infer the underlying parametric input PDFs. The problem with this approach is that in reality the given information is usually scarce and the convergence behaviour of spectral expansion methods is strongly dependent on the correct choice of input distribution [107]. If information is scarce, choosing the wrong distribution can inflict a large bias onto the solution, which cannot be reduced by increasing the order.

In this way, engineering applications have created a growing demand for the extension of Polynomial Chaos techniques to more general input distributions [107]. The gPC was extended to arbitrary input distributions by splitting the random space into piecewise elements and propagating them locally using the Askey scheme [144, 145, 146, 117]. The splitting of the random space allows to treat discontinuous input distributions [76], but it comes at the cost of increased computational effort, especially for multiple input variables. Approaches to finding global polynomial expansions for arbitrary distributions

have been developed based on Gram-Schmidt orthogonalisation [152, 154]. Unfortunately, the convergence of gPC and Gram-Schmidt approaches both rely on the accurate availability of a parametric input PDF. More recently, it was therefore suggested to base the Polynomial Chaos Expansion (PCE) on raw statistical moments by Witteveen [154].

Oladyshkin and Nowak observed that statistical moments provided a very general approach to propagate data without requiring the determination of a deterministic PDF. They found that every set of random data, as well as any continuous or discrete PDF, could be described using the moments without making any assumptions about the shape or existence of a suitable probability distribution. Oladyshkin [107] then promoted the concept in the geosciences and successfully applied it to identify uncertainties in carbon dioxide storage in geological formations [106, 109, 5] and also for robust design [108]. Moreover, Oladyshkin and Nowak presented a derivation of the optimal orthogonal polynomials from the moments. In the cases they reported, the convergence of the moment-based expansion was always significantly better than for any other polynomial expansion [107] based on fitted parametric PDFs. In order to differentiate Polynomial Chaos-based expansions based on statistical moments from so-called 'moment methods' as described in [14], Oladyshkin called his method **arbitrary Polynomial Chaos (aPC)**. In the following section, a new version of aPC will be developed. This version will be referred to as **Data-Driven Polynomial Chaos** as it was designed particularly for industrial problems with scarce or limited data.

2.2 Theory of Polynomial Chaos Expansions

Polynomial Chaos is an expansion approach which formulates a stochastic model output $f(\bar{\xi})$, dependent on a vector of N_U independent stochastic input random variables $\bar{\xi} = \xi_1, \xi_2, ..., \xi_{N_U}$ with space of events Ω, σ-algebra Λ and probability measure Γ (formally, a stochastic process in the probability space $(\Omega, \Lambda, \Gamma)$) as a linear combination of N_P stochastic multivariate orthogonal polynomials $\Psi_i(\bar{\xi})$ and deterministic coefficients α_i

$$f\left(\xi_1, \xi_2, ..., \xi_{N_U}\right) \approx \sum_{k=1}^{N_P} \alpha_k \cdot \Psi_k\left(\xi_1, \xi_2, ..., \xi_{N_U}\right). \tag{2.1}$$

According to Ernst [42], a Polynomial Chaos expansion exists and converges for a set of random variables as long as each basic random variable has finite statistical moments of all orders. The speed of convergence is mostly

determined by the non-linearity of the mapping between input and output variables the expansion is applied to. For more non-linear mappings, higher orders become necessary. For a total order expansion, the number of linear combinations terms N_P for a polynomial expansion of order p is truncated to:

$$N_P = \frac{(N_U + p)!}{N_U! p!}. \tag{2.2}$$

The multivariate orthogonal polynomials Ψ_k are calculated as products of univariate orthogonal or orthonormal polynomials $\psi_j^{(i)}$

$$\Psi_k\left(\xi_1, \xi_2, ..., \xi_N\right) = \prod_{i=1}^{N_U} \psi^{(i)}_{I_k^{(i)}}\left(\xi_i\right), \ k \in \overline{1, N_P}. \tag{2.3}$$

such that $\sum_{i=1}^{N_U} I_k^{(i)} \leq N_P \ \forall i$. The superscript index in brackets $\psi^{(i)}$ refers to the input random variable, and the subscript index to the order ψ_j of the orthogonal polynomial in the corresponding univariate basis. The rows of the index matrix $I_k^{(i)}$ contain the information, which order of each univariate input polynomial basis contributes to a particular global multivariate polynomial. For example, the 7^{th} row of a three dimensional index matrix corresponds to the polynomial product:

$$\Psi_7 = \psi_0^{(1)} \cdot \psi_2^{(2)} \cdot \psi_0^{(3)} \ \Leftrightarrow \ I_7^{(1,..,3)} = \begin{bmatrix} 0 & 2 & 0 \end{bmatrix}. \tag{2.4}$$

The univariate orthogonal or orthonormal polynomials $\psi_j^{(i)}$ with $j \in 0, .., p$ and $\psi_0^{(i)} = 1$ need to be developed individually for each input random variable ξ_i for $i \in 1, .., N_U$ and have to fulfil the orthogonality condition

$$\int_{\xi \in \Omega} \psi_m^{(i)}\left(\xi\right) \psi_n^{(i)}\left(\xi\right) d\Gamma = \delta_{mn} \ \forall m, n \in \overline{0, p}. \tag{2.5}$$

The n^{th} polynomial for the input distribution i can be defined through its $n+1$ polynomial coefficients $p_{n,j}^{(i)}$ as

$$\psi_n^{(i)}\left(\xi\right) = \sum_{j=0}^{n} p_{n,j}^{(i)} \xi^j, \ n = \overline{0, p}. \tag{2.6}$$

The use of custom basis functions is necessary to achieve the exponential convergence behaviour. For example, the Hermite polynomials form an optimal orthogonal basis for the normal distribution.

The statistics of $Y(\bar{\xi})$ can be evaluated through the coefficients α_k. The

mean and the variance of $Y(\bar{\xi})$ are

$$\mu_Y = \alpha_1 \quad \sigma_Y^2 = \sum_{k=2}^{N_P} \alpha_k^2 \left\langle \Psi_k^2 \right\rangle. \tag{2.7}$$

Similar formulas for skewness and kurtosis can be found in [24]. The coefficients α_k can be found through the integral

$$\alpha_k = \frac{1}{\left\langle \Psi_k^2 \right\rangle} \int_{\xi \in \Omega} Y\left(\bar{\xi}\right) \Psi_k\left(\bar{\xi}\right) d\Gamma\left(\bar{\xi}\right), \tag{2.8}$$

for $k = 1, .., N_p$, which can be solved using Galerkin projection, collocation or numerical integration. As the Ψ_k only consist of orthogonal polynomials, they can be accurately integrated by using a Gaussian quadrature rule as described in the next section. Although recent advances made by Navarro [101] have shown that a multivariate PCE of polynomials Ψ_k can be derived directly from dependent input variables using the Gram Schmidt method, the method described above is only valid in general, if the random variables are fully independent. Linear correlations between the input variables can be removed in advance, for example with a Karhunen Loève expansion [159].

2.2.1 Gaussian Quadrature

The concept of Gaussian quadrature is to find $p+1$ optimal Gaussian collocation points ξ_i and weights ω_i, such that their sum yields an exact integral for polynomials of degree $2p+1$ or less

$$\int_a^b f\left(\bar{\xi}\right) d\Gamma\left(\bar{\xi}\right) = \sum_{i=0}^{p} \omega_i f\left(\xi_i\right). \tag{2.9}$$

Multivariate quadrature formulas for multiple input random variables can be formed from one dimensional quadrature rules. Only tensor product quadrature and Smolyak's method are used in this thesis, but any other hyperbolic sparse quadrature is of course also possible. For a set of N_U random variables ξ_i with probability measure Γ the full tensor product quadrature formula is

$$\underbrace{\int_0^1 ... \int_0^1}_{N_U} f\left(\xi_1, ..., \xi_{N_U}\right) d\Gamma \approx \tag{2.10}$$

$$\approx \sum_{j_1=0}^{p_1} ... \sum_{j_{N_U}=0}^{p_{N_U}} f\left(\xi_{j_1}, ..., \xi_{j_{N_U}}\right) \left(\omega_{j_1} \otimes ... \otimes \omega_{j_{N_U}}\right).$$

In the literature, quadrature formulas exist in many variations to cover all types of parametric input distributions [72]. The approaches range from tabulated formulas [1] to moment matching equations [133] and the use of orthogonal polynomials [48]. With regard Gaussian-type quadrature rules based on orthogonal polynomials and their optimal collocation points, analytic formulas are readily available for several parametric continuous probability distributions. For example, the Hermite polynomials can be used to find optimal collocation points for the normal distribution; the Legendre polynomials for the Uniform distribution and the Laguerre polynomials for the exponential distribution [1]. The list of parametric distributions, for which pre-calculated optimal expansions exist, includes, among others also the Beta and Gamma distribution and is referred to as 'generalised Polynomial Chaos' (gPC) based on the Askey scheme [6, 158, 159]. A few examples are shown in Table 2.1.

Distribution	Interval	PDF	$w(x)$	Polynomial
Uniform	$[-1, 1]$	0.5	1	Legendre
Gaussian	$[-\infty, \infty]$	$\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$	$e^{-\frac{x^2}{2}}$	Hermite
Exponential	$[0, \infty)$	e^{-x}	e^{-x}	Laguerre

Table 2.1: The Askey scheme: a table that shows which polynomials provide optimal Polynomial Chaos basis functions for a given probability distribution.

Another advantage of NIPC approaches using Gaussian quadrature is that the optimal polynomials derived for uncertainty propagation can be reused to obtain the expansion coefficients. In the following, it will be explained how optimal Gaussian quadrature rules can be found based on the moments using matrix relations derived by Mysovskih [98] and then be sparsified using Smolyak's rule.

2.3 New Algorithm for Data-Driven Polynomial Chaos

One important novelty of this work is a new algorithm for arbitrary Polynomial Chaos. It allows the calculation of the optimal collocation points and weights, needed for Gaussian quadrature, based on random data only, without making statistical assumptions. All the quantities are calculated directly from the input data using only matrix operations performed on the Hankel matrix of moments. The matrix relations were derived by Mysovskih [98]

for Gaussian quadrature with arbitrary weights and a summary was written by Gautschi, which can be found in Golub and Welsch [51]. The algorithm developed in this work is generally more comprehensive than the previously suggested algorithm of aPC. It provides the user directly with the optimal Gaussian quadrature points and weights, which are always needed to evaluate the statistics of the posterior distributions. Due to the fact that it is often difficult to differentiate between the various Polynomial Chaos methods and their different methodologies in the field of UQ, the approach developed in this thesis will be referred to as data-driven Polynomial Chaos.

The crucial mathematical component of data-driven Polynomial Chaos is the use of statistical moments. These moments are a quantitative measure to describe the shape of a set of random samples or a probability distribution. The zeroth moment represents the integral of the PDF, which is always one by definition, the first moment the mean, the second the variance, the third the skewness, the fourth the kurtosis and so on. For a continuous random variable $\xi \in \Omega$ with PDF $w(\xi)$, the k^{th} raw moment μ_k can be determined by integrating

$$\mu_k = \int_{\xi \in \Omega} \xi^k w\left(\xi\right) d\xi. \tag{2.11}$$

For a discrete random variable $\xi \in \Omega^h$ with discrete PDF $w(\xi)$, the k^{th} raw moment μ_k is

$$\mu_k = \sum_{\xi \in \Omega^h} \xi^k w\left(\xi\right) \tag{2.12}$$

and for set of N samples (random draws or random measurement data) $\zeta_1, ..., \zeta_N$, the k^{th} raw moment μ_k can be calculated with

$$\mu_k = \frac{1}{N} \sum_{i=1}^{N} \zeta_i^k. \tag{2.13}$$

In general, a convergent PCE can be found for any arbitrary probability distribution or set of random data if its Hankel matrix of moments (defined in Eq. 2.15) is determinate in the Hamburger sense [42]. A given set of N samples is determinate in the Hamburger sense, if and only if all corresponding quadratic forms are strictly positive, that is if

$$\det\left(M\right) = \begin{vmatrix} \mu_0 & \mu_1 & \cdots & \mu_p \\ \mu_1 & \mu_2 & & \mu_{p+1} \\ \vdots & & \ddots & \\ \mu_p & \mu_{p+1} & & \mu_{2p} \end{vmatrix} > 0, \tag{2.14}$$

where the entries μ_k for k from 0 to $2p$ are the statistical moments. Thus, the number of moments needed to determine a p^{th} order expansion is $2p+1$.

Due to the generality with regard to the input that is achieved by basing PCEs on the moments, one can consider continuous, discrete PDFs or even data sets for which no probability distribution is defined, without requiring a change of methodology. The idea to base PCEs on the moments has been mentioned in several previous works [154, 107, 42, 137], but was most prominently elaborated for UQ by Oladyshkin and Nowak [107]. The methodology described by Oladyshkin and Nowak is tailored to a small number (less than five) of uncertain input random variables. The algorithm described in the following dates back to matrix relations derived by Mysovskih [98] described in Golub and Welsch [51] and is more suitable for problems with more input uncertainties. The algorithm can be divided into two steps:

1. the three-term recurrence for the optimal orthogonal polynomials is found from the Hankel matrix of moments

2. the corresponding optimal Gaussian quadrature points and weights are calculated from the three-term recurrence.

These two steps have to be performed individually for each input variable. In this way, the method allows a direct calculation of the optimal collocation points and weights needed for Gaussian quadrature rules from a given set of moments. The Hankel matrix of the moments, required to perform the first step, is defined as

$$M = \begin{bmatrix} \mu_0 & \mu_1 & \cdots & \mu_p \\ \mu_1 & \mu_2 & & \mu_{p+1} \\ \vdots & & \ddots & \\ \mu_p & \mu_{p+1} & & \mu_{2p} \end{bmatrix}. \tag{2.15}$$

Because the Hankel matrix is positive definite its Cholesky decomposition of $M = R^T R$ can be computed, so that

$$R = \begin{bmatrix} r_{11} & r_{12} & \cdots & r_{1,p+1} \\ & r_{22} & \cdots & r_{2,p+1} \\ & & \ddots & \vdots \\ & & & r_{p+1,p+1} \end{bmatrix}. \tag{2.16}$$

Mysovskih Theorem states that the entries of the inverse matrix R^{-1} of R

$$R^{-1} = \begin{bmatrix} s_{11} & s_{12} & \cdots & s_{1,p+1} \\ & s_{22} & \cdots & s_{2,p+1} \\ & & \ddots & \vdots \\ & & & s_{p+1,p+1} \end{bmatrix} \tag{2.17}$$

form an orthogonal system of polynomials ψ_j according to the relation

$$\psi_j = s_{0j}\xi^0 + s_{1j}\xi^1 + ... + s_{jj}\xi^j \quad for \quad j = 0, ..., p. \tag{2.18}$$

To avoid the inversion of the matrix, Rutishauser [124] derived explicit analytic formulas to obtain the polynomial coefficients of the orthogonal polynomials s_{ij} from the Cholesky matrix entries r_{ij}. These relationships can be used to determine the coefficients a_j and b_j of the three-term recurrence, according to which each orthogonal polynomial $\psi_j(\xi)$ for $j = 1, ..., p$ satisfies

$$\xi\psi_{j-1}(\xi) = b_{j-1}\psi_{j-2}(\xi) + a_j\psi_{j-1}(\xi) + b_j\psi_j(\xi). \tag{2.19}$$

The coefficients a_j and b_j in terms of r_{ij} are

$$a_j = \frac{r_{j,j+1}}{r_{j,j}} - \frac{r_{j-1,j}}{r_{j-1,j-1}} \qquad b_j = \frac{r_{j+1,j+1}}{r_{j,j}} \tag{2.20}$$

with $r_{0,0} = 1$ and $r_{0,1} = 0$. The knowledge of the three-term recurrence relation allows the calculation of the optimal collocation points and weights for any orthogonal polynomial [51]. They can be calculated by describing the three-term recurrence relation through a symmetric tri-diagonal Jacobi matrix J, which is again positive definite:

$$J = \begin{bmatrix} a_1 & b_1 & & & & \\ b_1 & a_2 & b_2 & & \emptyset & \\ & b_2 & a_3 & b_3 & & \\ & & \ddots & \ddots & \ddots & \\ & \emptyset & & b_{p-2} & a_{p-1} & b_{p-1} \\ & & & & b_{p-1} & a_p \end{bmatrix}. \tag{2.21}$$

The eigenvalues of J are the roots of the polynomial of order p and the weights are found by

$$\omega_i = v_{1,i}^2, \tag{2.22}$$

where $v_{1,i}$ is the first component of the normalized eigenvector corresponding

to the i^{th} eigenvalue.

2.3.1 Ill-Conditioned Moment Matrix

For some random data sets, the moment matrix can become too ill-conditioned to perform a Cholesky decomposition without preconditioning. This is generally never the case for parametric probability distributions unless they are fat-tailed and the tail effects are considered for a long range. There are various solutions to solve this problem. The easiest solution is to apply probability density matching. A tutorial and recent advances can be found in Che [21].

2.4 Demonstration for Different Inputs

Using the method described above, the optimal collocation points for arbitrary continuous distributions and almost arbitrary random samples can be calculated easily.

2.4.1 Probability Distributions

Figure 2.1 shows several examples of various distributions both contained and not contained in the Askey scheme. It can be seen by comparing the graphs in Figure 2.1 that the Gaussian collocation points change their position according to the input distributions. Unfortunately, they also change their position according to the polynomial expansion order p, which is why Gaussian quadrature rules are generally not nested. This means that lower order expansions cannot be reused for higher orders. On the upside, optimal Gaussian quadrature rules are more accurate than other quadratures using weights and collocation points. They are most commonly used based on tabulated formulas. However, the use of tabulated formulas restricts Gaussian quadrature methods to a limited set of continuous probability distributions. Moreover, it becomes tedious to combine individual rules to obtain integrals for multi-dimensional integrals with several different input distributions. The examples in Figure 2.1 demonstrate how all Polynomial Chaos approaches can be calculated using the data-driven PC algorithm.

2.4.2 Random Samples

Not only parametric distributions but also histograms can be propagated using the moments without a change in methodology. Figure 2.2 shows that a moment-based arbitrary PCE can even be calculated for multi-modal, or

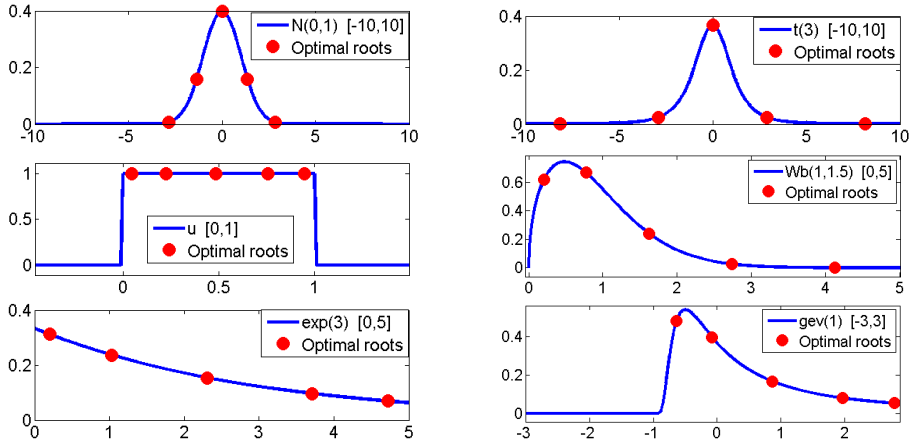


Figure 2.1: Optimal Gaussian collocation points calculated with data-driven PC for three Askey scheme PDFs: standard Normal $N(0,1)$, Uniform $U[0,1]$, Exponential $exp(3)$ and three non-Askey scheme PDFs: fat-tailed Student-t $t(3)$, Weibull $Wb(1,1.5)$ and generalised Extreme Value Distribution $gev(1)$.

mixed probability distributions, as shown by the collocation points (dots). It

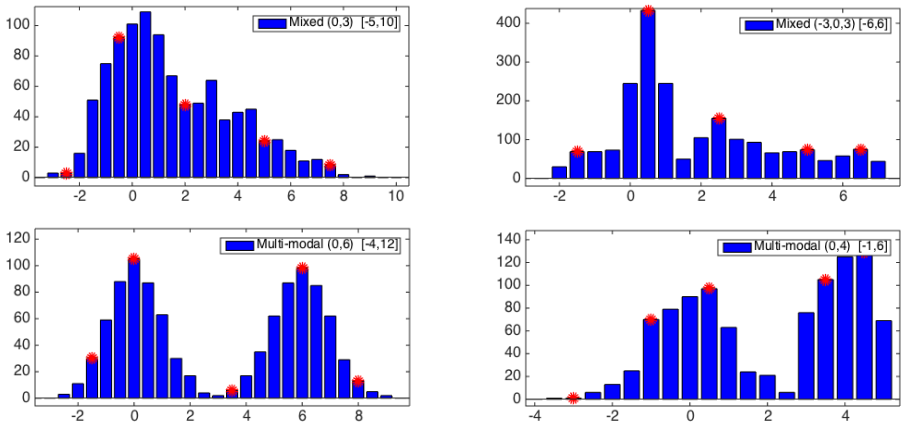


Figure 2.2: Optimal Gaussian collocation points calculated with data-driven PC for various mixed and multi-modal histograms. The results show that the method adapts well to multi-modal data. The collocation points favour points of high likelihood.

is visible how the moment-based approach favours regions in the input histograms of high probability. The input data can then be propagated through the computational model with exponential convergence if the input-output mapping is analytic. In general, the rate of convergence improves with increasing smoothness of the input-output mapping. Finally, the moments of

the posterior distribution can be trivially evaluated using quadrature, since the data-driven PC algorithm provides the optimal quadrature weights. If, however, the shape of the posterior PDF is desired, it can only be obtained by sampling the PCE. A rough estimate of its shape can be obtained using the input samples, but for more accurate illustrations of the posterior PDF, a refined random number generator based on the samples is required.

2.4.3 Propagation of Multi-Modal Data

The graph at the bottom-left corner of Figure 2.2 shows two main modes: one with its peak at zero and one with its peak at six. This could be interpreted as two means. However, the data-driven Polynomial Chaos expansion is based on a single sequence of moments. Therefore, the question may arise, if it can correctly propagate this multi-modal input data.

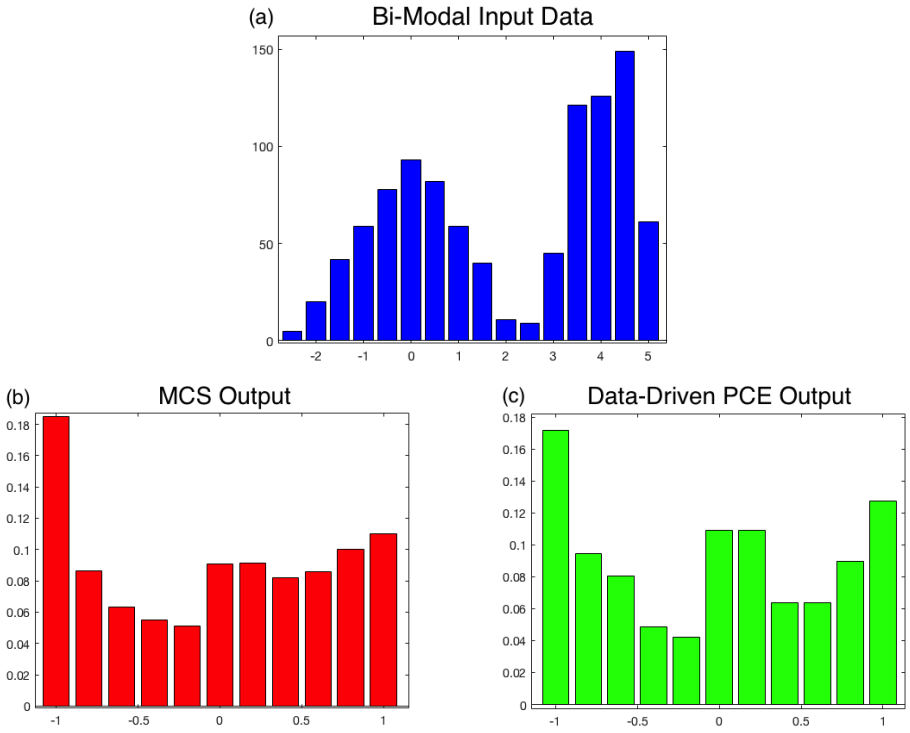


Figure 2.3: This graph compares the results of Monte Carlo Simulation with data-driven PCE for a multi-modal input histogram. The results in (b) and (c) show that multi-modality is not a problem as long as the input/output mapping is continuous and sufficiently smooth.

To answer this question, 1000 random samples are drawn from a bi-modal distribution that is created by drawing 500 samples from the normal distribu-

tion $N(0,1)$ and 500 samples from $N(4,1/3)$. This multi-modal distribution is shown in Figure 2.3a. The samples are propagated with the non-linear function $y(x) = sin(x)$. The PDF resulting from the PCE approach is compared with the PDF that is obtained from 10^7 direct Monte Carlo samples drawn from the same distribution. Figures 2.3b,c show that, although the input distribution is bi-modal and the PCE is based on 10^4 samples less than the Monte Carl simulation, the resulting probability distributions are quite similar. This can be explained by the fact that the PCE creates a meta-model of the response surface and the two modes are then found by bootstrapping the available random samples through the polynomial expansion. Therefore, multi-modality is no problem for data-driven Polynomial Chaos, unless the input to output mapping is discontinuous (a solution for this is suggested in Chapter 4) or not enough samples are available to see that there are two modes.

2.5 Validation: CFD Simulation of a Diffuser

In this example, it is shown how the fitting of parametric PDFs to small data sets can lead to a large bias and that the direct use of the available data is more accurate if the underlying distribution of a small sample set is unknown. This is done by propagating uncertainty through the CFD simulation of a diffuser, highlighting the impact of different PDF fittings on the output. From the results, it is concluded that the direct propagation of the experimental data set is preferable to the fit of parametric distributions if data is scarce. Thus, the suggested method offers an alternative to the maximum entropy theorem to handle scarce data.

The used CFD simulation of a 2D diffuser is shown in Figure 2.4 and was carried out with the open-source packages OpenFoam. To ensure accuracy all PC approaches were chosen fifth order accurate, although convergence was already observed for the third order expansion. Diffusers are common components present in various engineering applications: cars, gas turbines or pipes. Their efficiency is strongly dependent on the occurrence and size of the separation, which in turn is sensitive to the geometry. A particularly important parameter of the geometry is the diffuser angle. For the given test case, the separation behaviour was analysed by Chandavari et al. [18]. They set up the CFD simulations based on a test case, for which experimental validation data were obtained by Dippold and Georgiadis [29]. Chandavari et al. varied the taped angle deterministically and evaluated the outcome with respect to the occurrence of separation. They found that separation occurs for ta-

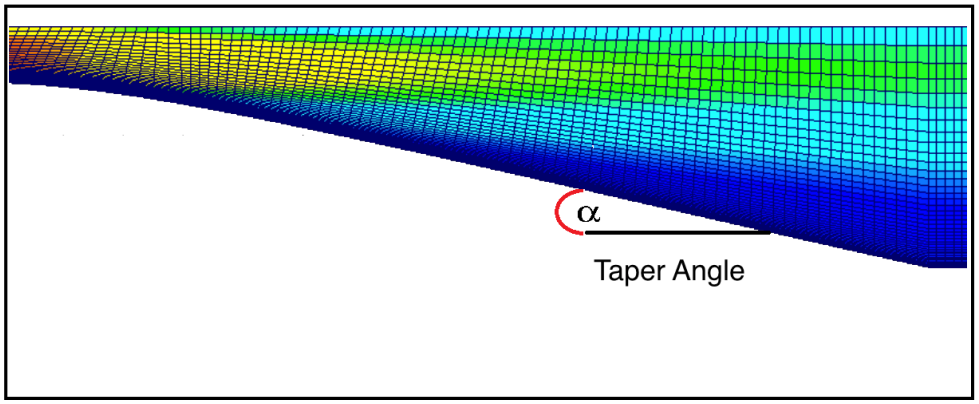


Figure 2.4: This figure shows the mean velocity field of the diffuser flow on a coarser mesh than used for CFD. The taper angle α is highlighted at the bottom side of the channel.

per angles larger than 7 degrees of the given configuration. In the following, the same CFD simulations described by Chandavari is performed assuming a stochastic taper angle variation. All geometrical measures, details about the solver, turbulence model and boundary conditions are described in full in [18]. The results demonstrate that the direct usage of the data significantly reduces the risk of large errors if no information about the probability distribution is available.

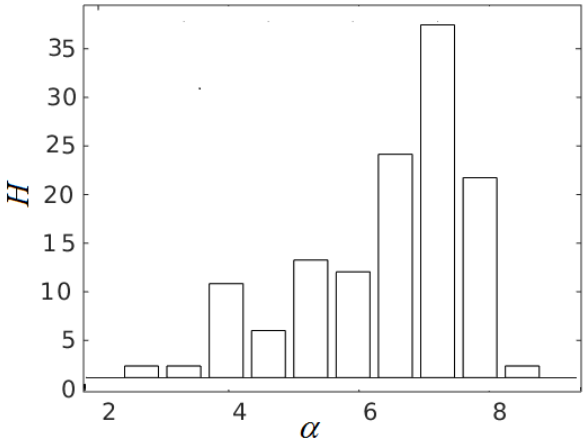


Figure 2.5: The taper angle α of the diffuser varies stochastically in the CFD simulation. This histogram shows the various values the taper angle can randomly assume.

The taper angle α is assumed to be affected by manufacturing variation, for which a set of 100 angle measurements between 2 and 8 degrees is drawn

from a $\beta(12, 2)$ distribution. The histogram of chosen measurement samples is shown in Figure 2.5 and is representative of what is commonly obtained by optical or mechanical measurements of geometrical errors: it shows clustering around a mean and a certain skewness to the left side. This skewness might disappear for larger sample sizes, or it might stay. The modeller cannot know without performing more measurements or having similar data from a previous or similar test case at hand. If neither is possible, data-driven PC becomes the best choice of method. The CFD case described in the following demonstrates this advantage of using data-driven PC.

1. Data-driven PC propagation of 100 samples
2. Parametric Beta Distribution as reference solution $\beta(12, 2)$
3. PDF: fitted Gaussian Distribution: $N(0.841, 0.096)$
4. PDF: fitted Gamma Distribution: $\gamma(71.051, b = 0.012)$
5. PDF: fitted Beta Distribution: $\beta(12.898, 2.457)$
6. PDF: fitted Weibull Distribution: $W(0.880, 11.958)$

The results in Table 2.2 show that data-driven PC is very accurate, even though it does not require any knowledge about the true distribution and is based on only 100 samples. All PDFs propagated through the CFD code are:

Graphically the fits are shown in Figure 2.6. Table 2.2 shows the relative error of mean and standard deviation compared to the moments of the assumed solution: a Beta distribution $\beta(12, 2)$. The results of Table 2.2 can be summarised as follows: The mean is estimated most accurately using data-driven PC and the error for the standard deviation is less than 5%. The Gaussian and Gamma distribution drastically misrepresent the solution, the Beta and Weibull deliver good results being skewed. However, the error in the standard deviation is about 3% even for these two skewed PDFs. The relative errors are calculated with respect to the assumption that the parametric Beta distribution $\beta(12, 2)$ from which the 100 samples were drawn provides the correct solution. data-driven PC is the only method that does not require any PDF fitting and can directly model the chosen 100 samples.

For real data, it would be difficult to know from only 100 samples whether the final distribution is actually skewed or not and the best assumption is to use the available data. This shows a typical situation in engineering and only the direct use of the moments does not limit the model to a specific PDF

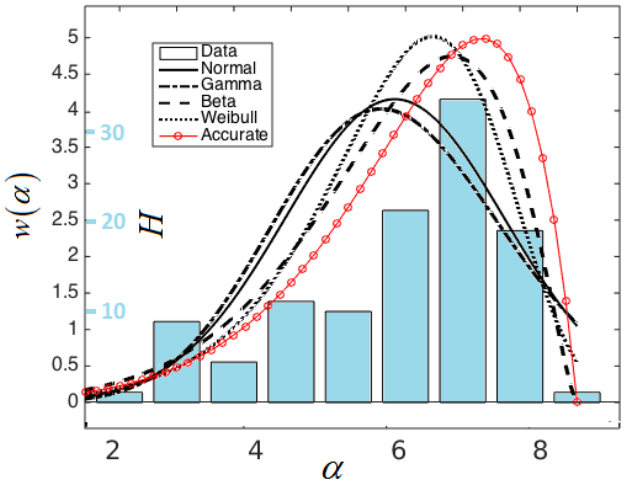


Figure 2.6: Various parametric PDFs are fitted to the available random samples of the diffuser taper angle. The different PDFs show that even though 100 samples is not a small data set, there is a large difference between the various distributions.

assumption with creating potentially large bias errors. The data in Table 2.2 show the possibility to avoid a bias caused by PDF fitting. Even the best fit of the most suitable model can eventually just reproduce the same moments as the purely moment-based expansion [106]. The results of a data-driven expansion will naturally converge towards the true solution for increasing number of input samples. Further examples investigating higher dimensional problems and a range of analytical test functions are analysed in the next section.

Fitted Distribution	$\epsilon_{rel}(\mu)$	$\epsilon_{rel}(\sigma)$
Normal	0.89%	72.39%
Gamma	0.46%	154.52%
Weibull	0.14%	2.13%
Beta	0.18%	3.05%
Data-driven PC (level 1)	**0.09%**	**4.94%**

Table 2.2: Relative error for mean and variance of data-driven PC compared to PCE based on fitted parametric distributions.

3 SAMBA: Adaptive and Anisotropic Sparse Grids

Previous works on data-driven Polynomial Chaos considered only one or very few stochastic inputs [107]. Industrial applications, however, require the consideration of between $10-50$ input variables. This number of variables is quite challenging for all Polynomial Chaos methods. As the dimension of a problem increases, the computational cost grows exponentially and the rate of convergence becomes slower. This is generally referred to as the 'curse of dimensionality' [40]. To give an example: a tensor grid (standard) NIPC approach of third order requires $\sim 10^6$ model evaluations for ten variables and $\sim 10^{30}$ for fifty variables. This shows how ten variables are already prohibitively expensive for engineering simulations if a tensor grid is used in multiple dimensions.

Prior to this thesis, it was unknown how the curse of dimensionality would affect a data-driven Polynomial Chaos expansion. It was unknown if the numerical approximation errors caused by obtaining polynomial chaos expansion through scarce data would cause a problem in higher dimensions. This issue is analysed in this section and it is shown that the error is reasonably small compared to the error one would otherwise make by fitting an inappropriate parametric distribution.

A sparse quadrature rule called Smolyak's algorithm is adapted to alleviate the curse of dimensionality and to extend data-driven Polynomial Chaos efficiently to higher dimensions. Smolyak was the first to propose sparse quadrature to reduce the number of model evaluations in Gaussian quadrature [135]. The central principle of Smolyak's algorithm is to select a small number of the most important elements of a full tensor quadrature to represent the multidimensional function. The idea behind this is that most often the low-order interactions between the input variables are the most important ones. Similar to the full tensor Gaussian quadrature being improved with higher polynomial levels, the Smolyak quadrature can be improved by increasing a parameter called the level [63]. Most works, like [40, 7], use the Smolyak formulas based on polynomial interpolation at the extrema of the Chebyshev polynomials (Clenshaw-Curtis nodes). Similarly, Smolyak formulas based on Gaussian

nodes in the literature are often only explained based on a single univariate rule. Conventionally Smolyak's method is only used to sparsify a single univariate quadrature rule to multiple dimensions. Therefore, it is another novelty of this work that Smolyak is used to define an anisotropic quadrature rule.

In this chapter, Smolyak is adjusted so that it can perform the complete procedure from almost arbitrary random data (or arbitrary distribution) to sparse multivariate quadrature rule. This generality means the described quadrature technique can be used outside the field of UQ. It simplifies the calculation of any anisotropic and adaptive Gaussian quadrature rule for multi-dimensional integrals (anisotropic meaning different weights in each direction and adaptive meaning individual polynomial expansion orders for the input variables), as it removes the need to refer to tabulated integration formulas.

3.1 SAMBA's Algorithm

In the following, Smolyak's conventional quadrature rule is adapted for multiple univariate optimal Gaussian quadrature rules all based on moments. The N_u sequences of one-dimensional quadrature rules $\{U^{i_j}\}_{j=1,...,N_u}$ are denoted as

$$U^{i_j} = \sum_{k=1}^{m_{i_j}} f(\xi_k^{i_j})\omega_k^{i_j}, \tag{3.1}$$

where m_{i_j} $j \in \{1, ..., N_u\}$ is the maximum order, individually chosen for a quadrature, in case an adaptive quadrature is wanted (adaptive meaning that individual orders for the input distributions can be realised). For the Smolyak quadrature based on moments, the set of nodes corresponds to the eigenvalues of the matrices J. The Smolyak quadrature based on different input distributions is calculated with the formula

$$A\left(N_u + l, N_u\right) = \sum_{\substack{l+1 \\ \leq |i| \leq \\ l+N_u}} (l-1)^{l+N_u-|i|} \begin{pmatrix} N_u - 1 \\ l + N_u - |i| \end{pmatrix} \otimes_{k=1}^{N_U} U^{i_k} \tag{3.2}$$

where l is the level. It is used to control the accuracy of the result, similar to the order in full Gaussian quadrature. For the same number of uncertain inputs, increasing the level l may make for a better accuracy, but it comes with an exponential growth of the number of points. The term $|i|$ is the norm of the vector $i = \{i_1, ..., i_{N_u}\}$ which also stands for the sum of a row j of the

index matrix $I_j^{(k)}$:

$$|i| = \sum_{k=1}^{N_u} i_k = \sum_{k=1}^{N_u} I_j^{(k)}. \tag{3.3}$$

Equation (5.2) is a linear combination of the tensor products, where only a small number of nodes is used, but the interpolation properties for $N_u = 1$ are preserved for higher dimensions [148]. The array of sparse grid points $H(N_u + l, N_u)$ can be obtained through the following formula:

$$H\left(N_u + l, N_U\right) = \bigcup_{l+1 \leq |i| \leq l+N_u} (\xi^{i_1} \times \cdots \times \xi^{i_{N_u}}), \tag{3.4}$$

where the ξ^{i_j} are $m_{i_j} + 1$ points used by the quadrature U^{i_j} for $j \in \{1, ..., N_u\}$. The Fourier coefficients α_k can be calculated either by sparse integration of

$$\alpha_k = \int_\Omega f\left(\xi_1, ..., \xi_{N_u}\right) \Psi_k\left(\xi_1, ..., \xi_{N_u}\right) d\Gamma \tag{3.5}$$

or by linear regression

$$\hat{\alpha} = \arg\min \left[\alpha^T \psi\left(\xi\right) - Y\right]^2, \tag{3.6}$$

where $Y = \left[f\left(\xi_0\right), ..., f\left(\xi_{N_{sp}}\right)\right]$ and $\Psi\left(\xi\right) = \left[\psi_0\left(\xi\right), ..., \psi_{N_{sp}}\left(\xi\right)\right]$ with N_{sp} as the number of sparse points. The minimum of Equation (3.6) can be solved by defining

$$A = \begin{bmatrix} \psi_1\left(\xi_1\right) & \cdots & \psi_1\left(\xi_{N_u}\right) \\ \vdots & \ddots & \vdots \\ \psi_{N_{sp}}\left(\xi_1\right) & & \psi_{N_{sp}}\left(\xi_{N_u}\right) \end{bmatrix} \tag{3.7}$$

and solving

$$\left(A^T A\right) \hat{\alpha} = A^T Y. \tag{3.8}$$

The regression approach becomes increasingly useful for higher levels.

The PC Fourier coefficients are useful in case the shape of the posterior distribution is of interest. If only the moments of the posterior distribution $E\left[f^k\right]$ are required, they can be obtained more easily by using the sparse quadrature formulas directly on the formula for the k^{th} moment. For a model f they can be calculated by sparsely integrating

$$E\left[f^k\right] = \underbrace{\int_0^1 ... \int_0^1}_{N_U} \left(f\left(\xi_1, ..., \xi_{N_U}\right) - E\left[f^{k-1}\right]\right)^k d\Gamma. \tag{3.9}$$

Using the sparse collocation points η_i and weights θ_i this becomes

$$E\left[f_k\right] = \sum_{i=1}^{N_{Sp}} \left(f\left(\eta_i\right) - E\left[f^{k-1}\right] \right)^k \theta_i. \tag{3.10}$$

A comparison of the increasing computational effort of the lowest level Smolyak quadrature with full tensor Gaussian quadrature of third order is shown in Figure 3.1. Several examples of anisotropic grids, obtained from the moments

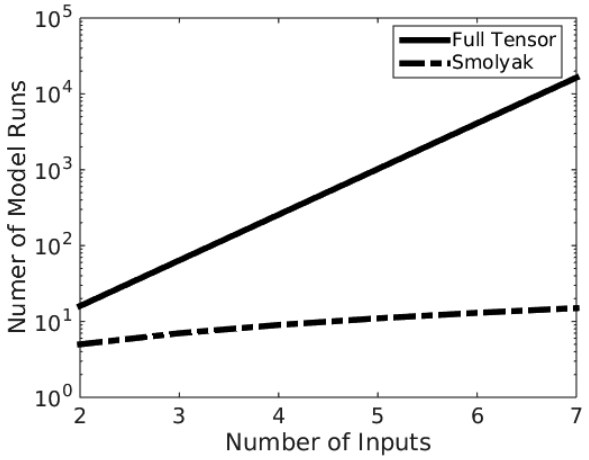


Figure 3.1: Illustration of how Smolyak's sparse grid algorithm alleviates the curse of dimensionality. The number of collocation points for a level 2 Smolyak grid is orders of magnitude lower than for a 3^{rd} order full tensor grid.

of one Gaussian and one Weibull probability distribution, are shown in Figure 3.2.

3.1.1 Sensitivity Analysis using Sobol Indices

Polynomial Chaos provides a powerful and computationally efficient tool to calculate global sensitivity indices [138, 28] based on Sobol Indices. As the PCE formulation is very similar to the Sobol decomposition, all necessary information to perform sensitivity analysis is given with the PCE Fourier coefficients. Thus, the computational cost to find the effect of input parameters onto the model output reduces to that of estimating the PCE coefficients [138].

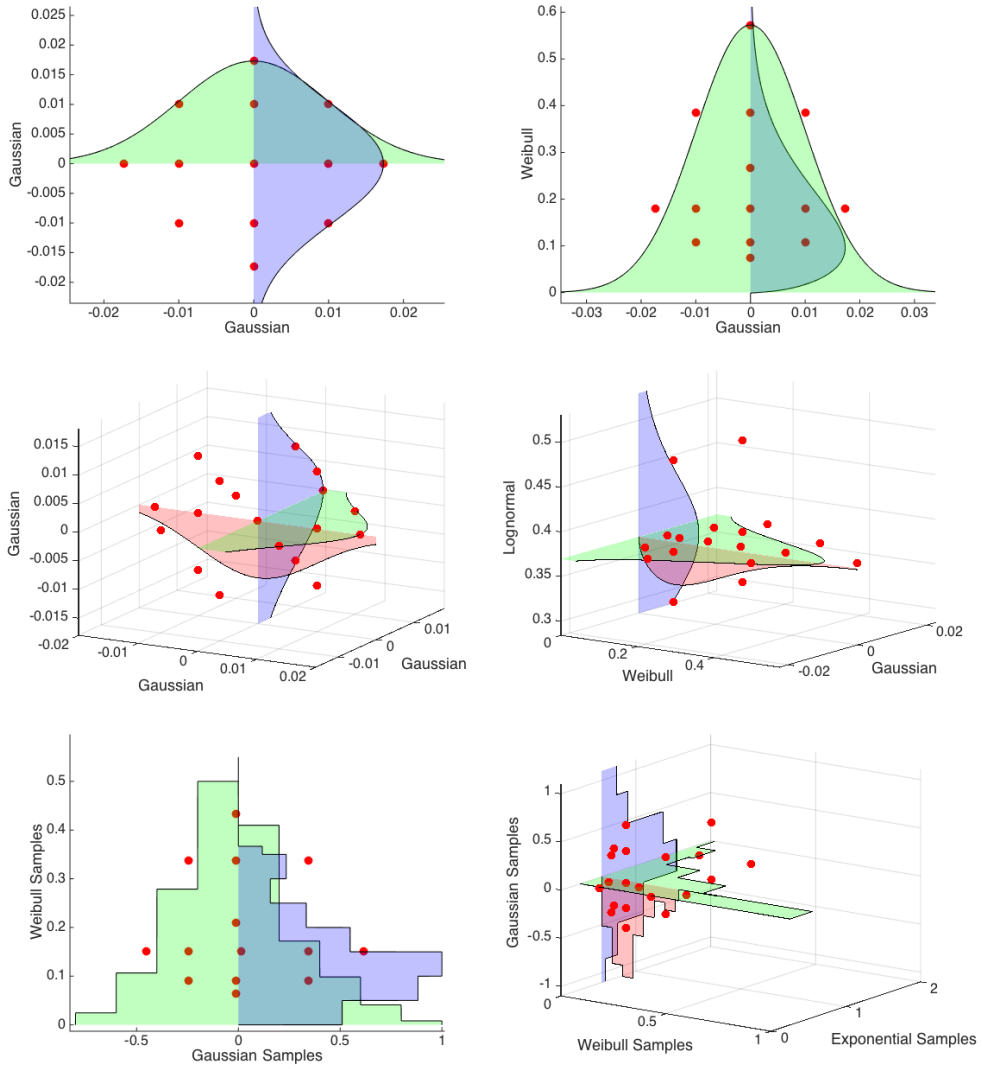


Figure 3.2: This figure illustrates the high level of generality and adaptability that is obtained by using SAMBA for quadrature grid generation. The graphs show various symmetric and asymmetric data-driven sparse grids for parametric distributions and random data histograms.

Using Gaussian integration the Sobol Indices S_k can be found by

$$S_k = \frac{Var\left[f_{\xi_k}\left(\xi\right)\right]}{Var\left[f\left(\xi\right)\right]} \approx \frac{\sum\limits_{k \in I_k} \alpha_k^2}{\sum\limits_{i=1}^{N_{Sp}} \alpha_k^2} \tag{3.11}$$

where the counting variable k sums up all rows of the multi-index matrix for which an index i_j is not zero. It is a collection of all results for which the input random variable ξ_i was varied.

3.1.2 Simple Example and Implementation Advice

The following two dimensional example illustrates how SAMBA can be implemented in matrix format. To obtain simple numbers two continuous distributions, namely the standard normal and the uniform distribution, are chosen: $N\left(0,1\right)$ and $U\left(0,1\right)$ and the Smolyak quadrature is chosen at level 1. The resulting sparse Smolyak index matrix for $l = 1$ is:

$$I_{l+1 \leq |i| \leq l+N_u} = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 2 & 1 \end{bmatrix}. \tag{3.12}$$

For more than 10 input variables, it is recommended to calculate the index matrix recursively because the memory consumption of saving the full combination matrix becomes very high. Using the SAMBA algorithm the collocation points corresponding to each input variable can be calculated

$$X^{(i_1)} = \begin{bmatrix} 0 & -1 & -1.73 \\ 0 & 1 & 0 \\ 0 & 0 & 1.73 \end{bmatrix} \quad X^{(i_2)} = \begin{bmatrix} 0.5 & 0.21 & 0.11 \\ 0 & 0.79 & 0.5 \\ 0 & 0 & 0.88 \end{bmatrix} \tag{3.13}$$

and then the corresponding weights

$$W^{(i_1)} = \begin{bmatrix} 1 & 0.5 & 0.16 \\ 0 & 0.5 & 0.66 \\ 0 & 0 & 0.16 \end{bmatrix} \qquad W^{(i_2)} = \begin{bmatrix} 1 & 0.5 & 0.28 \\ 0 & 0.5 & 0.44 \\ 0 & 0 & 0.28 \end{bmatrix}, \tag{3.14}$$

where the superscript index in brackets refers to the random variable. Recall that i_1 represents the first column of the index matrix in Equation (3.12) while i_2 goes through the second one. The sparse collocation points can then

obtained with

$$\xi_j = \bigcup_{2 \leq i_1 + i_2 \leq 3} \chi^{i_1} \times \chi^{i_2} \tag{3.15}$$

which is in detail:

$$[0] \times [0.5]\,, \quad [0] \times \begin{bmatrix} 0.21 \\ 0.79 \end{bmatrix}, \quad \begin{bmatrix} -1 \\ 1 \end{bmatrix} \times [0.5]\,, \tag{3.16}$$

so that the array of sparse collocation points becomes

$$\xi_j^{(i)} = \begin{bmatrix} 0 & 0.5 \\ 0 & 0.21 \\ 0 & 0.79 \\ -1 & 0.5 \\ 1 & 0.5 \end{bmatrix}. \tag{3.17}$$

For the conventional Smolyak algorithm, it can happen that in higher dimensions the same point is repeated several times. If that happens, the same points can be summarised to one position to decrease the number of necessary model runs, but the weights have to be added together. The corresponding sparse weights can be calculated through the tensor product

$$\omega_j = k_S \cdot \left(\otimes_{k=1}^{N_u} W^{i_k} \right) \quad j = \{1, 2, 3\} \tag{3.18}$$

where k_S is the Smolyak counting coefficient

$$k_S = (-1)^{l + N_u - \sum_{i=1}^{N_u} I_j^{(i)}} \begin{pmatrix} N_u - 1 \\ l + N_u - \sum_{i=1}^{N_u} I_j^{(i)} \end{pmatrix} \tag{3.19}$$

which is in detail:

$$-1 \cdot \prod_{i=1}^{N_u} [1] \otimes [1], \quad 1 \cdot \prod_{i=1}^{N_u} [1] \otimes \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \quad 1 \cdot \prod_{i=1}^{N_u} \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \otimes [1], \tag{3.20}$$

so that the array of sparse integration weights becomes

$$\omega_j^{(i)} = \begin{bmatrix} -1 \\ 0.5 \\ 0.5 \\ 0.5 \\ 0.5 \end{bmatrix}. \tag{3.21}$$

The mean and standard deviation of the posterior distribution of a model f can then be calculated through the scalar products

$$\mu = f\left(\xi^{(1)}, \xi^{(2)}\right) \circ \omega, \quad \sigma = \left(f\left(\xi^{(1)}, \xi^{(2)}\right) - \mu\right)^2 \circ \omega. \tag{3.22}$$

3.2 Validation of SAMBA

In this section, SAMBA is validated using a set of different input distributions and the results are validated against Monte Carlo Simulations with 10^7 samples. The used multivariate test functions are listed in Table 3.1. As shown by the 2D cross sections in Figure 3.3, the test functions were chosen to cover a wide range of possible response surfaces exhibiting, for instance, sinusoidal behaviour or exponential growth. To demonstrate the high potential of Smolyak's quadrature the two lowest possible levels were used. For accuracy comparison, also full tensor Gaussian quadrature of 3^{rd} order was performed. The stochastic input random variables used for 5D and 10D testing are listed

#	$f(\xi_1, \xi_2, ..., \xi_N)$
1	$\sum_{i=1}^{N} \xi_i$
2	$\exp\left(\sum_{i=1}^{N} \xi_i\right)$
3	$\sum_{i=1}^{N} \sin\left(\xi_i - 0.5\right)$
4	$1 \Big/ \sum_{i=1}^{N} \left(\xi_i - 1\right)^2$

Table 3.1: Multivariate test functions f used for 5 and 10 input random variables.

in Table 3.2. The ten distributions used for validation are the Normal (N), Uniform (U), Cauchy (t) or also called student-t with one degree of freedom, Log-normal (Nlog), Weibull (Wb), Exponential (exp), Gamma (Γ), Beta (β), generalised Pareto (gPD) and generalised Extreme Value distribution (gEVD). To include the effect of the fat tails of the Cauchy distribution, it was not only truncated but also squeezed to a shorter x range by calculating $t(1000x)$ instead of $t(x)$. In this way, a long part of the tails can be considered. The selected set of distributions does not only comprise Askey-scheme, but also more unusual distributions. From a conventional point of view, setting-up a sparse quadrature rule for this case would be difficult, but for SAMBA there is no difference to the use of a single univariate Gaussian rule. Similarly, also the truncation of the distributions can be performed without effort. For the

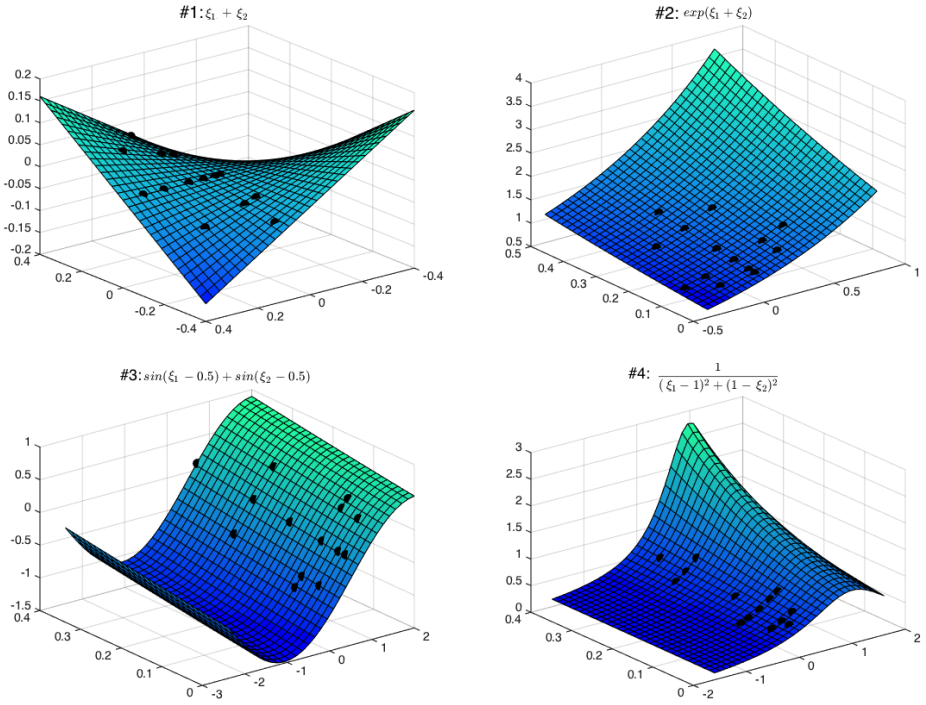


Figure 3.3: Response surfaces of the four multi-dimensional non-linear test functions from Table 3.1 including level 2 data-driven sparse grids for one Gaussian and one Weibull input distribution.

Input	Parametric PDF
ξ_1	$N(\mu = 0, \sigma = 0.1)_{[-1,1]}$
ξ_2	$U(a = 0, b = 0.1)_{[0,0.1]}$
ξ_3	$Nlog(\mu = 0, \sigma = 4)_{[0,1]}$
ξ_4	$Wb(a = 0.12, b = 1.5)_{[0,1]}$
ξ_5	$t(\nu = 1, x_t = 1000x)_{[-0.5,0.5]}$
ξ_6	$exp(\mu = 0.1)_{[0,1]}$
ξ_7	$\Gamma(a = 2, b = 0.1)_{[0,2]}$
ξ_8	$\beta(a = 1, b = 10)_{[0,1]}$
ξ_9	$gPD(\theta = 0, \sigma = 0.1, k = 3)_{[0,2]}$
ξ_{10}	$gEVD(\mu = 0, \sigma = 0.1)_{[0,1]}$

Table 3.2: Probability distributions used for the input random variables ξ_i during testing with their parameters (e.g. μ, σ, a, b) and the used truncation interval $[l, r]$.

continuous case, merely the integration boundaries need to be adapted. The

$f(\xi_1, .., \xi_5)$	MCS σ_{ref} $N = 5 \cdot 10^7$	Tensor ε_{rel} $N = 1024$ $P = 3$	SAMBA ε_{rel} $N = 11$ $\mathbf{l = 1}$	SAMBA ε_{rel} $N = 65$ $\mathbf{l = 2}$
1	0.281289	0.000182	0.000182	0.000182
2	0.491630	0.000142	0.040644	0.001254
3	0.267742	0.000139	0.008996	0.000065
4	0.021974	0.000015	0.006559	0.000021

Table 3.3: Relative error of standard deviation for five different input distributions compared to MC reference solution for the four test function in Table 3.1 and first five random variables of Table 3.2.

$f(\xi_1, .., \xi_{10})$	MCS σ_{ref} $N = 5 \cdot 10^7$	Tensor ε_{rel} $N \approx 1.05 \cdot 10^6$ $P = 3$	SAMBA ε_{rel} $N = 11$ $\mathbf{l = 1}$	SAMBA ε_{rel} $N = 231$ $\mathbf{l = 2}$
1	0.685013	0.000428	0.000428	0.000428
2	3.171731	0.000064	0.220369	0.035225
3	0.497968	0.000407	0.025444	0.000376
4	0.009219	0.000675	0.071664	0.005869

Table 3.4: Relative error of standard deviation for ten different input distributions compared to MC reference solution for the four test function in Table 3.1 and first ten random variables of Table 3.2.

relative errors compared to Monte Carlo for five input distributions are shown in Table 3.3 and for 10 distributions in Table 3.4. It can be seen that despite the asymmetry, the sparse asymmetric Smolyak quadrature results in maximum relative errors of order 10^{-1} and minimum relative errors of order 10^{-3} for the standard deviation. The full tensor quadrature is more accurate, but it needs significantly more points. The results in Table 3.3 and 3.4 are based on the moments calculated from continuous distributions.

3.2.1 Effect of Sample Size

The great advantage of SAMBA is that histograms can also be used. For this reason, sample sets were drawn from all input distributions $\xi_1, ..., \xi_{10}$. The obtained histograms were propagated directly through the test functions. To investigate the effect that the size of the input data set has on the accuracy of the solution, a Monte Carlo Simulation using SAMBA one thousand times for random draws with different seed numbers was performed. The accu-

racy was estimated with a comparison to a brute-force Monte Carlo solution with 10^7 samples. The propagated histogram bin sizes range from 10 to 10^6

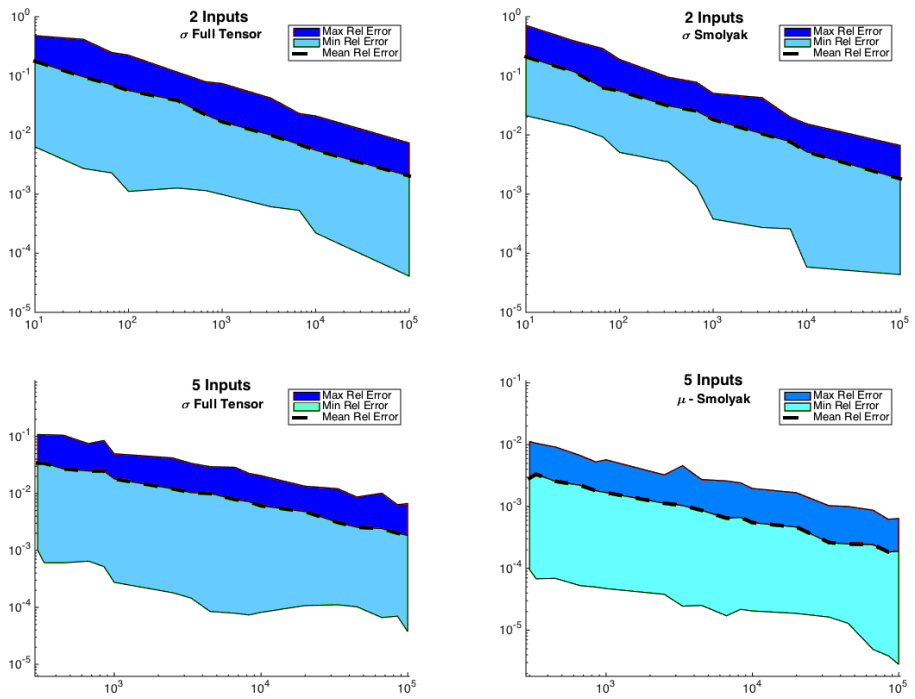


Figure 3.4: The four graphs show the relative error of the predicted mean including 100% confidence intervals for full tensor Gaussian and Smolyak's method compared to MCM for increasing sample size and for the first 2 and 5 input distributions from Table 3.2. It can be seen that the error grows with the number of inputs and decreases with the number of random samples available. The confidence intervals are created by repeating the uncertainty quantification process with different random samples.

random samples, which were drawn from the truncated input distributions characterised in Table 3.2. The obtained convergence plots are shown in Figure 3.4. The plots show the minimum and maximum error occurring in the Monte Carlo Simulation as upper and lower bound of the occurring relative error. Thus, the upper and lower lines in Figure 3.4 represent 100% confidence interval. Moreover, Figure 3.4 show that a 10% maximum relative error can be obtained from as few as 100 input samples for the standard deviation. Moreover, it was found that the convergence of the solution's accuracy depends on the number of input distributions. For more input distributions the general accuracy decreases. All the same, one digit accuracy can still be

obtained for 100 samples. However, the accuracy decreases to a mere 2 digits for 10^5 samples for 5 inputs, whereas it reaches 3 digits for 10^5 samples for 2 inputs. In conclusion, the results show that decent estimates of probabilistic output quantities can be obtained even if only small data sets are used as input distributions.

3.2.2 Error Compared to PDF Fitting

To actively compare the results of SAMBA and PDF fitting for different test samples, Gaussian distributions are fitted to each sample set of the previously performed 5 input Monte Carlo Simulation with test function 2. The fitted Gaussian distributions are propagated through the function using sparse PCE. The results shown in Figure 3.5 show three distinct regions:

1. For data sets with less than 100 samples, the use of the SAMBA and direct data is clearly preferable.

2. In the region between 100 and 1000 samples the accuracy of PDF fitting and SAMBA are comparable.

3. For data sets with more than 1000 samples, the use of the wrong distribution imposes a limit on the maximum possible accuracy.

Hence, SAMBA is more accurate for very small data sets and multiple inputs. Since the use of SAMBA is also easier than PDF fitting the figures demonstrate that the use of SAMBA is preferable in case little information is available. In general, they show that SAMBA is a safe alternative for engineers who are not experienced with performing statistical tests like the Anderson-Darling-Test to reduce the fitting error.

3.3 Rare Event Simulation using SAMBA

It is probably impossible to predict or completely prevent the occurrence of rare events with catastrophic consequences, the so-called Black Swans. However, it is possible to design systems which are more robust so that the occurrence of Black Swans becomes less likely. The next generation of aircraft engines should be designed with the intrinsic ability to deal with rare events in order to increase the engine life and reduce the probability of dramatic accidents. In order to achieve this, new UQ methods need to be developed which are capable of reducing the effect of the tails of the input PDF on the system performance. Efficiency is thereby one of the most crucial aspects.

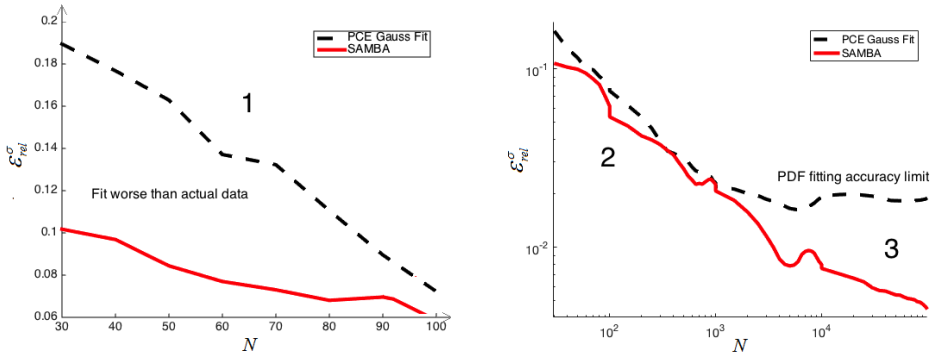


Figure 3.5: The two figures compare SAMBA with Polynomial Chaos based on Gaussian PDF fitting for the first five input distributions from Table 3.2 and for growing sample size. It can be seen that SAMBA is more accurate for scarce data and large data.

Realistic CFD models still have very long run-times, which make brute force approaches such as Monte Carlo Simulations difficult to use, even if using parallel high-performance computers.

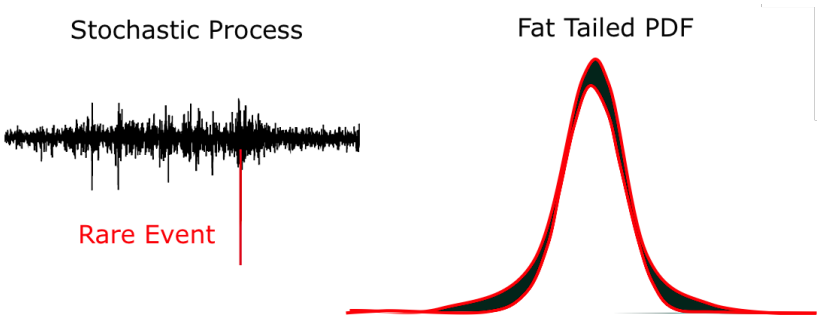


Figure 3.6: The left graph shows a time series with one extreme outlier. If this outlier leads to detrimental consequences it is called a Black Swan. The right graph shows a fat-tailed probability distribution that can be used to model rare, extreme outliers statistically.

In order to propagate rare events, fat-tailed distributions as shown in Figure 3.6 have to be used as input distributions. Because fat-tailed distributions have infinite moments, a meaningful propagation using Polynomial Chaos has so far been impossible. Non-intrusive polynomial chaos methods can only be applied to propagate a probability distribution or random data sets if the input distributions have finite moments. Some attempts have been made for the propagation of only the fat tails (truncated from the main part of the distribution) using Polynomial Chaos. This was done, for example, by Safta to develop basis functions for the fat tail of a log-normal distribution [125].

However, no general framework for complete fat-tailed distributions has been developed so far in the open literature. This is done for the first time in this chapter.

The crucial modification to use SAMBA for fat-tailed distributions is quite simple: truncation. Truncated heavy-tailed PDFs have finite moments, even if the tails are extremely long [82]. The simplest method to achieve this is to define a partially continuous function that is zero outside of a required fat-tailed domain. A mathematically more elegant solution is also explored for the first time in this chapter: smooth truncation using KoBoL distributions. KoBoL distributions are smoothly truncated so that they have finite moments of all orders while presenting fat tails up to a certain scale followed by a smooth cut-off. This approach is more elegant because it upholds the fundamental theorem of probability theory that the integral of a PDF should always be exactly one. For the truncated case this is only approximately true. It should be remarked at this point that in most engineering cases the truncation error is absolutely negligible. The smooth truncation approach was first suggested by Koponen [64] to truncate Lévy flights.

3.3.1 Fat-Tailed Probability Distributions

While there is no universal agreement about the terminology, the term heavy-tailed is prevalently considered as mathematically more general than fat-tailed and is used to refer to probability distributions which do not have exponentially bounded tails and therefore infinite moments [115]. Visually, heavy-tailed probability distributions can be recognised because they have higher peaks and kurtosis than for example the Gaussian distribution shown in Figure 3.7. A large family of heavy-tailed distributions is given by Alpha stable Lévy distributions with shape factor $0 \leq \nu \leq 2$. Other important distributions are the Student-t distribution and the Cauchy distribution. The most distinguishing feature of all heavy or fat-tailed distributions is that their tails decay as $f(x) = x^k$ and not as $f(x) = k^x$, for any variable x and constant k. While this difference may seem negligible at a first glance of their probability distribution in Figure 3.7, it has a strong effect on their moments. The raw moments of a random variable ξ with probability density function $f(\xi)$ are defined as

$$\mu_k = \int_{\xi \in \Omega} \xi^k f(\xi) d\xi \tag{3.23}$$

If $f(\xi)$ decays as $1/x^k$ and not exponentially, the moments μ_k can become infinite. The application of polynomial chaos methods to propagate uncertainty

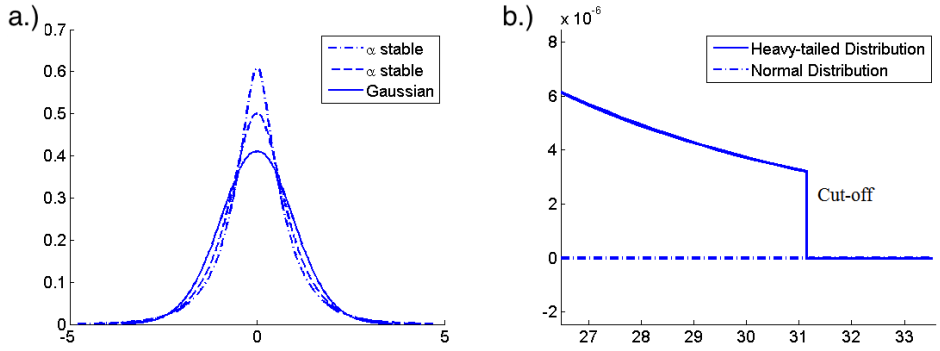


Figure 3.7: The left graph shows several alpha-stable distributions for different shape factors. The right graph shows a discontinuously truncated tail of a heavy-tailed distribution.

caused by heavy-tailed distributions has so far been kept back in engineering by the fact that heavy-tailed distributions have infinite moments. This can, however, be circumvented by acknowledging that a cut-off can be imposed at the boundaries of physical systems or at the boundaries of their feasible operating domain. If this cut-off l_{cut} is also imposed on the distribution of the heavy-tailed random variable, it ensures the finiteness of the moments. However, a cut-off extremely far away from the mean can create a huge magnitude difference between the moments. This may cause the numerical method of the data-driven Polynomial Chaos method to become ill-conditioned. The imposed cut-off can either be abrupt and discontinuous, as suggested by Mantegna and Stanley for alpha-stable distributions [82] or smooth and continuous as described by Koponen [64] or Dubrulle [36]. While a continuous cut-off is more satisfactory from a mathematical point of view, because the integral of the PDF remains exactly one, it does not make more sense from a physical perspective [36]. For example, in Lattice Boltzmann simulations the jumps of the particles are inevitably limited by the size of the box [57]. For this reason, a continuous but abrupt cut-off was suggested by Dubrulle [36]. However, Dubrulle remarks that this does not qualitatively affect the results. Since Dubrulle's results are merely valid for truncated Lévy distributions and not heavy-tailed distributions in general, an abrupt truncation related to the physically realisable values of a random parameter is more advisable.

3.3.2 Tempered Alpha Stable Distributions

For the special case that the physical cut-off is so far off the mean that the moments become too high, it is possible to fit a tempered Lévy stable distri-

bution to the original distribution and thereby impose a smooth exponential cut-off without imposing an unphysical sharp cut-off. The truncated distribution imitates the statistical properties of the non-truncated alpha stable Lévy distribution for a smaller (finite) sample numbers and the sums of tempered stable random variables converge to normal random variables for a very huge (infinite) number of samples. Therefore, the usual Central Limit Theorem is not violated. For random variables z_i with identical distribution $f_L(z)$ the truncated Lévy distribution $L_T(x)$ is defined as the limit of the sum of the random variables z_i in the Poissonian stochastic process (as explained in Feller [44]), whose characteristic function $\varphi(k)$ obeys:

$$\ln(\varphi(k)) = \int_{-\infty}^{\infty} \left(e^{-ikz} - 1\right) f(z)\, dz \tag{3.24}$$

where f_L is defined as

$$f_L(z) = \begin{cases} C|z|^{-1-\nu} & z < 0 \\ Cz^{-1-\nu} & z > 0 \end{cases} \tag{3.25}$$

Using Lévy's Inversion Formula the probability distribution $L_T(x)$ corresponding to the characteristic function can be calculated with

$$L_T(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \varphi(k) e^{-ikx} dx \tag{3.26}$$

To achieve a smooth exponential cut-off, the characteristic function $\varphi_X(t) = e^{\phi(X)}$ can be developed analytically by integrating an exponential cut-off into the probability distribution f_L from Eq. (3.25):

$$L_T(z) = L(z)\, e^{-\lambda z} = \begin{cases} C|z|^{-1-\nu} \cdot pe^{-\lambda|z|} & z < 0 \\ Cz^{-1-\nu} \cdot qe^{-\lambda z} & z > 0 \end{cases} \tag{3.27}$$

where λ is the variable used to control the cut-off. The characteristic function $\varphi_X(t) = e^{\phi(X)}$ can be derived as described by Nakao [100]

$$\phi_X(t) = C\Gamma(-\nu)\left[\left(\lambda^2 + t^2\right)^{\nu/2} \cos(\nu\theta) + i(q-p)\sin(\nu\theta) - \lambda^\nu\right] \tag{3.28}$$

for $0 < \nu < 1$ and with $\theta = \arctan\left(\frac{t}{\lambda}\right)$ and

$$\phi_X(t) = C\Gamma(-\nu)\left[q(\lambda + it)^\nu + p(\lambda - it)^\nu - \lambda^\nu - i\nu\lambda^{\nu-1}(q-p)\, t\right] \tag{3.29}$$

for $1 < \nu < 2$.

By using the expansion: $q(\lambda + it)^\nu + p(\lambda - it)^\nu = \left(\lambda^2 + t^2\right)^{\nu/2} \cos(\nu\theta) + i\beta \sin(\nu\theta)$ (where $\beta = q - p$ describes the symmetry) it can be seen that in the limit $\lambda \to 0$, the tempered stable distribution converges to the stable distribution. Unfortunately, an analytical expression for the probability distribution is not attainable for more than a few exceptional cases [100]. The probability density $L_T(x)$ from Eq. (3.26) can, however, be obtained numerically. By applying adaptive numerical integration suitable for oscillating integrals (like, for example, Gauss-Kronrod Quadrature) to Eq. (3.28) or (3.29) accurate numerical values for $f_X(x)$ can be calculated [64]. This was done in Figure

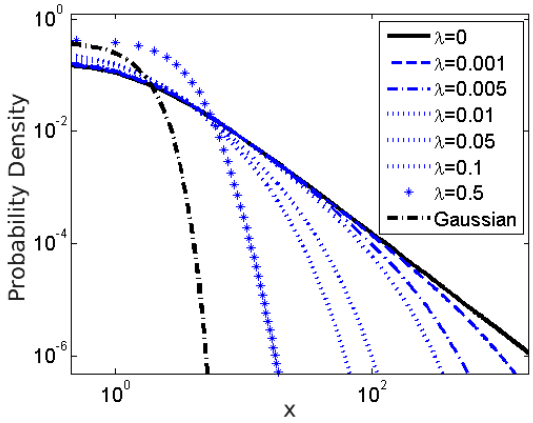


Figure 3.8: For tempered alpha stable distributions the length of the fat tail can be controlled through a cut-off parameter. This graph shows tails of varying truncation.

3.8 for various values of the cut-off parameter λ from purely fat-tailed (Lévy distribution) to fully Gaussian tails.

3.3.3 Fitting Alpha-Stable Distributions

The parameters of the characteristic function of a tempered Alpha-Stable Lévy distribution as characterised by Eq. (3.28) or (3.29) can be determined by using a Maximum-Likelihood approach. It is important to pre-set the cut-off parameter λ to the desired value before performing the distribution fitting because its value influences the overall behaviour of the distribution. A fitted distribution is shown in Figure 3.9. If the physical constraints on the stochastic parameter are known beforehand λ should be chosen such that the tails of the tempered distribution follow the tails of the not-truncated Lévy distribution in the relevant physical domain. For example: by choosing the cut-off parameter as $\lambda = 0.05$ and fitting a tempered alpha stable distribution

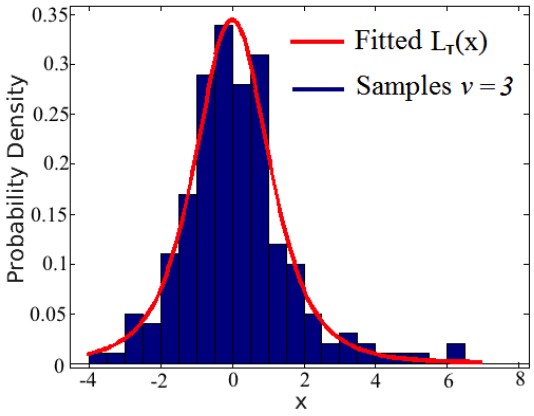


Figure 3.9: Example of a tempered alpha stable distribution fitted to fat-tailed Student-t samples.

to the histogram of 10000 samples obtained from a student-t distribution with 3 degrees of freedom as shown in Figure 3.8, the following parameters can be obtained: $C = 0.411569$, $\nu = 1.5268$ and $\beta = 0$, thus $p = 1$ $q = 1$. The optimal orthogonal polynomials for this exponentially tempered heavy-tailed distribution are approximately:

$$\begin{aligned} &1 \\ &x \\ &x^2 - 3.1525 \\ &x^3 - 262.52x \\ &x^4 - 2393.91x^2 + 6719.26 \\ &x^5 - 4898.38x^3 + 664205x \\ &x^6 - 771.769x^4 + 7413360x^2 - 18943630 \end{aligned} \tag{3.30}$$

3.3.4 Validation of SAMBA for Rare Events

In order to validate the proposed polynomial method it is applied to a simple one dimensional case, for which the model is given by the analytical formula

$$g\left(\tilde{t}\right) = 0.05\left(\tilde{t} - 10\right)^3 + 3\left(\tilde{t} - 10\right) + 120 \tag{3.31}$$

where $\tilde{t}$ consists of a deterministic and stochastic part: $\tilde{t} = t + \xi$, $t \in R$ and ξ is a random variable with three degrees of freedom Student-t probability distribution ($\nu = 3$). Eq. (3.31) is chosen as a simplification of the transient gas turbine temperature profile used in the real engineering example in application chapter, in Section 6.4. Figure 3.10 shows the optimal Gaussian and

fat-tailed collocation points for the halved distribution that occur because the time is assumed to be non-negative.

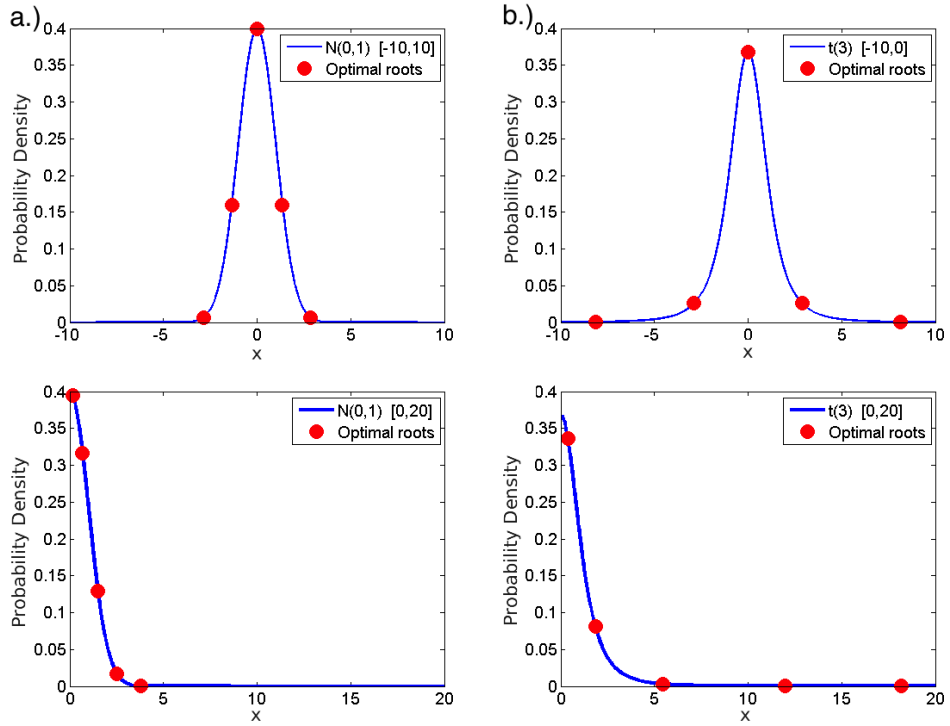


Figure 3.10: Optimal collocation points for the Gaussian distribution (left) and Student-t distribution (right). The bottom graphs illustrate that optimal points can also be calculated for parts of distributions.

For a heavy-tailed t-distribution with only 3 degrees of freedom, only the mean and variance are finite. Skewness, kurtosis and all higher moments are undefined. Since the finiteness of the moments is a necessary condition to apply arbitrary polynomial chaos, the distributions have to be truncated. In this case, it will be assumed that the stochastic space of ξ is limited between 0 and 20. The truncated heavy-tailed distributions have finite moments so that SAMBA can be applied. Particularly noteworthy is that SAMBA can even propagate the half probability distribution created at $t = 0$ and that the optimal collocation points for this case are entirely positive, as shown in Figure 3.11. The red arrows in both graphs have the same lengths at the same position in time. The results are validated with a Monte Carlo Method with 10^6 samples. The samples for the MCM are drawn from a heavy-tailed t-distribution with three degrees of freedom, $\nu = 3$, which has also been truncated to the interval $[0, 20]$.

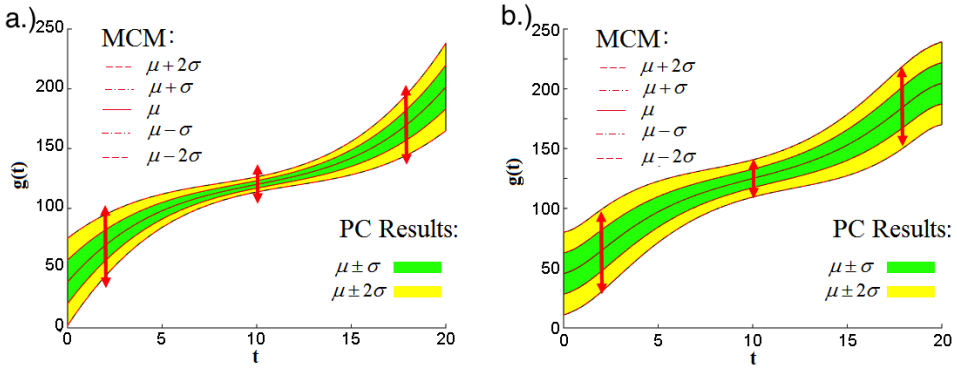


Figure 3.11: Output variation for a Gaussian input (left) and a heavy-tailed input distribution (right).

3.4 Applicability and Limitations

The mathematical simplifications offered by SAMBA come with certain limitations, which should be taken into account if the method is to be used accurately. First of all, it has to be noted that the use of the moments becomes less reliable for higher order polynomials. The results in the following section indicate that for a given number of inputs level 10 should not be exceeded. The origin of this effect can easily be seen, if the moments are obtained using numerical integration. As the calculated moments become higher, the integration error will cause an increasing deviation from the correct values. If the moments are obtained from data, the smaller data size of the input data will also result in lower overall accuracy. Moreover, it should be taken into account that the response surface should be smooth in order to use Smolyak quadrature. In addition, the distance between the collocation points should take fluctuations in the response function into account. Otherwise, aliasing errors will occur.

3.4.1 Integration Errors

If the moments are obtained from continuous distributions using numerical integration, integration errors have to be considered. In the author's experience, the use of state-of-the-art adaptive numerical integration techniques will allow an accurate calculation of the first 50 moments without exceeding error thresholds of 10^{-6}. For example, for two continuous Gaussian input distributions the polynomials up to order 28, that means at least the first 56 moments can be integrated accurately. Therefore, up to 25 different continuous input distributions at level 2 can be considered with SAMBA without large integration errors. Moreover, the accuracy of the numerical integration depends on

the chosen probability distribution and the length of the integration interval. For heavier tailed distributions, the moments grow significantly faster. In case the integration fails to converge to the required accuracy, it is possible to truncate the probability distributions. For example, by considering a heavy-tailed distribution on the interval $[-100\sigma, 100\sigma]$ instead of on the interval $[-\infty, \infty]$ integration errors can be significantly reduced, whereas the major part of the PDF remains unchanged.

3.4.2 Scarce Data Errors

No integration errors occur if the moments are obtained from histograms. However, the calculated moments will deviate more from what could be assumed their true value, the higher the moment becomes and the smaller the given data set is. Nevertheless, the calculated moments are at least based on the actually available information. In case the wrong distribution is chosen for a set of samples, the higher moments up to infinity will deviate even more from what can be considered their accurate value. As was shown in the method validation section using a Monte Carlo Simulation on SAMBA for different sample sizes, even small sample sizes result in good estimates of the moments of the posterior distribution. More evidence of the superior convergence of the data-driven Polynomial Chaos and several examples demonstrating it can be found in Oladyshkin and Nowak [107].

3.4.3 Convergence

As mentioned before, PCEs based on moments converge if the input random variables are independent, have finite moments and are determinate in the Hamburger sense. The determination of the Hamburger sense is especially important if the data-driven expansion is applied for fat-tailed PDFs, because if very long truncated fat tails are considered, the moment matrix can become indeterminate in the Hamburger sense. This may stop the method from converging to the correct solution as demonstrated in Ernst et al. [42]. The rate of convergence improves with increasing smoothness of the input-output mapping. In the cases reported by Oladyshkin and Nowak the convergence of aPC was superior to expansions based on incorrect parametric distributions. The crucial factor that decides whether the aPC will be more accurate than a fitted parametric PDF is the suitability of the fitted distribution. In case a given sample is small and very unrepresentative of the distribution, a fitted distribution will create a larger error than the direct use of the moments.

3.4.4 Higher Order Coefficient Errors

SAMBA is a pseudospectral projection method using an adaptive and anisotropic Smolyak sparse grid rule. This means that the quantities of interest are formulated as a spectral expansion in multivariate orthonormal polynomials and the Fourier coefficients α are calculated with a sparse numerical quadrature rule. This formulation is known to lead to large errors in the higher order Fourier coefficients, due to internal aliasing [25], if the expansion level is increased for the same number of input parameters.

To find the maximum order at which SAMBA as proposed in this work can be used for increasing expansion orders, the numerical integration examples as proposed in [26] were reproduced and tested for rising expansion level l. The two-dimensional test functions used in [26] are repeated in Table 3.5 . Figure

#	$f(\xi_1, \xi_2)$
1	$\xi_1^{10}\xi_2^{10}$
2	$e^{\xi_1+\xi_2}$
3	$\sin(5(\xi_1 - 0.5)) + \cos(3(\xi_2 - 1))$
4	$1 \Big/ \left(2 + 16(\xi_1 - 0.1)^2 + 25(\xi_2 + 0.1)^2\right)$

Table 3.5: Test cases used in Constantine [26] to demonstrate errors in higher order Fourier coefficients.

3.12 shows the resulting relative error in the standard deviation for increasing level compared to 10^8 Monte Carlo simulation samples. To emphasize the fact that SAMBA can be used for continuous distributions and discrete sample sets, the left side of Figure 3.12 shows the convergence assuming 2 independent continuous uniform distribution $U[-1, 1]$ as inputs. The right side shows the use of two different and independently created random sets of 50 samples drawn from $U[-1, 1]$ using a random number generator. Since the samples are not identical and the sample size is very small, the resulting expansion is not symmetric. In general, the following behaviour can be observed:

- The ill-conditioning of the mapping from moments to recurrence coefficients prevents expansion orders of higher levels than 21 for the continuous case and level 19 if only 50 samples are given. Consequently, the effect of the data error on the ill-posedness of the problem is not too large. The main reason that the method fails for higher levels than 20 is that the Cholesky decomposition is no longer positive definite. At order 21 the first eigenvalue of the Hankel matrix of moments becomes negative, namely $-4e^{-17}$.

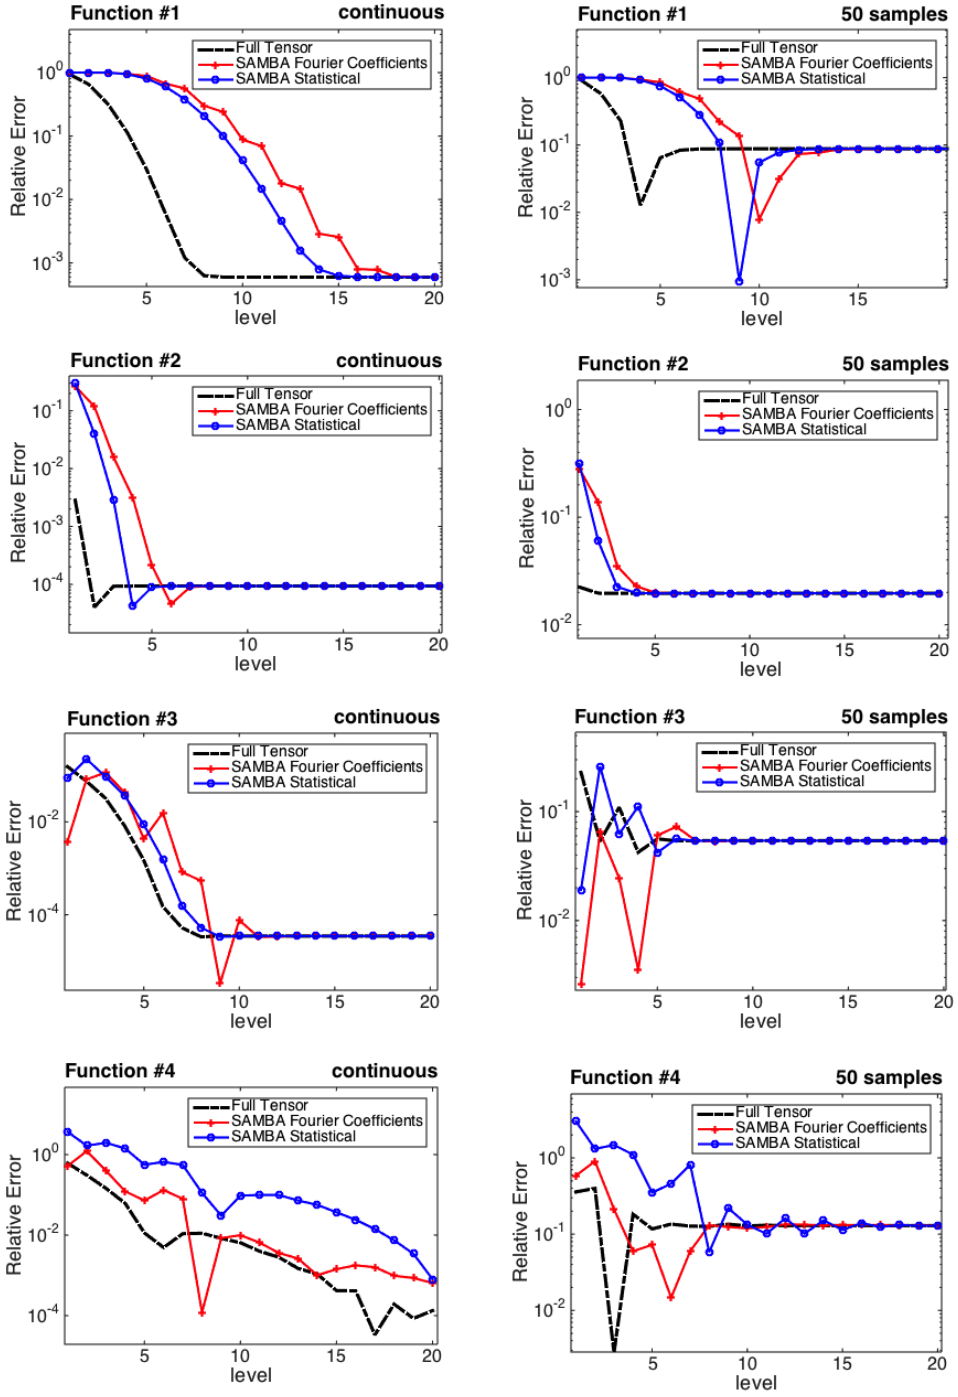


Figure 3.12: The figures show the relative error for the standard deviation calculated with SAMBA compared to 10^8 Monte Carlo samples. The four test functions are taken from Table 3.1. The graphs show that there is an accuracy limit for increasing order and that this accuracy limit is dependent on the number of samples available.

- The accuracy of the solution is reduced by the sample size. Since the continuous expansion uses the same distributions as the Monte Carlo solution, the accuracy achieved using only 50 samples is smaller. The use of the moments provides better results if the accurate underlying parametric distribution cannot be determined.

- The method converges to the correct solution for increasing level, even if only 50 data points are used.

- The maximum accuracy is reached at lower levels than 10 in case only limited information is available about the input distribution.

- The use of sparse methods can be preferable if data is very limited. The reason is that the additional effort spent on the tensor product is no use if the overall possible accuracy is bounded by the limitation in input information. If this is the case, the maximum possible accuracy can be reached significantly faster and with lower computational effort using SAMBA.

An illustration of how the error for the individual coefficients grows for increasing order is given by Figure 3.13. The figure shows the magnitude of

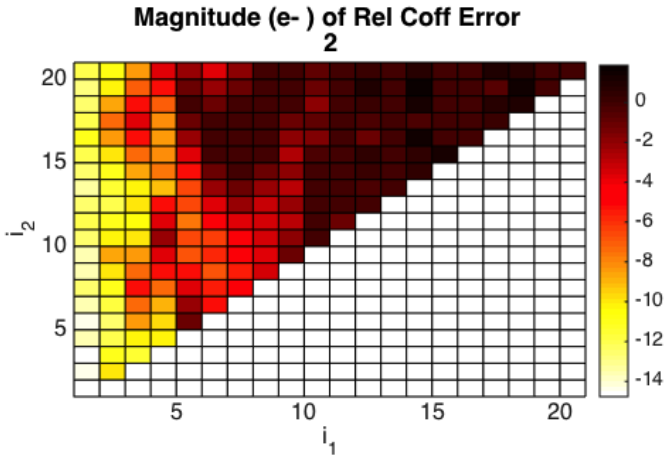


Figure 3.13: Magnitude of the relative error in the sparse Fourier coefficients based on data compared to analytic tensor product coefficients. It can be seen that the error is negligible for orders lower than five. Higher levels than 10, however, contain significant errors.

the relative error in the Fourier coefficients for function 3 from Table 3.5. It can be seen that the Fourier coefficients are highly accurate for lower levels.

The maximum error is only of magnitude e^{-4} for polynomial orders lower than 11, i.e. level 10 in Smolyak. For higher levels the error in the coefficient increases to around 100%. Comparing these results with the convergence behaviour from Figure 3.12 leads to the conclusion that the increasing error in the coefficients does not result in an additional error because no increase in relative error can be seen for higher orders. The error remains constant. For all cases, the error in the standard deviation stagnates for higher levels. This means that the additional computational effort for higher levels will most likely not result in improved accuracy.

In summary, the use of Smolyak's quadrature in SAMBA is justified and preferable to full tensor product methods, if the method is applied to a case where only scarce information is available for 10-20 input random variables. The coefficients up to level 10 can be determined accurately. Higher levels are unlikely to result in improved accuracy. In case accurate parametric distributions can be determined, the use of the SPAM [26] is recommended.

4 A Global Discontinuous Model Based on Arbitrary Points

Continuous Polynomial Chaos methods like SAMBA do not perform well if there are discontinuities in the modelling domain [15]. They require very high orders for convergence and create oscillations in the vicinity of the discontinuity, referred to as Gibbs Phenomenon. The prevalent solution to this problem is to partition the stochastic domain into continuous sub-domains, and then to employ continuous Polynomial Chaos separately to the individual sections. Such approaches are called local methods. Examples include the Multi-Element generalised Polynomial Chaos (ME-gPC) as shown in Wan and Karniadakis [143] and the Multi-Element Probabilistic Collocation method as shown in Lin and Tartakovsky [74]. The main issue with these methods is that the position of the discontinuity has to be known a priori. Sargsyan [128] solved this problem by using Bayesian inference to locate the discontinuity. While this approach works, it still needs a fairly high amount of model evaluations.

To reduce the number of necessary model evaluations to an absolute minimum, a global approach based on the computationally least expensive method is developed in this chapter: the Padé- approximation. Padé approximations are rational functions of two Polynomial Chaos expansions. When both are computed using SAMBA, arbitrary basis functions can be used. While the possibility of using arbitrary basis functions for Padé approximations is new, the main novelty of this chapter is that, so far, numerical quadrature was a necessary component to obtain a Padé approximation. This is computationally expensive and inconvenient for industrial applications. It is therefore replaced with a least squares approximation for the first time in this chapter. This means that the Padé approximation introduced here can be determined using an arbitrary set of collocation points and does no longer require fixed Gaussian quadrature points. The resulting method is also quite flexible with regard to switching between continuous and discontinuous problems. As a global least squares method based on SAMBA, it can fit models to both discontinuous and continuous non-linear response surfaces with a lower number of data points

than other methods, see [139]. Its only disadvantage is that previous authors have encountered stability issues related to the ill-conditioning of the least squares problem from randomly chosen collocation points.

To overcome the stability issue, the Padé approximation is combined with a recently developed method called Optimal Quadrature Sub-sampling proposed by Seshadri [134]. The method is proven to significantly improve the numerical condition of the least squares problem by using an optimal choice of points from a tensor quadrature instead of random samples. The points are chosen to ensure a good condition number of the design matrix. In this way, one can obtain more accurate results than through random sampling while using a significantly lower number of model evaluations.

To further guarantee the reliability of the Padé response surface, it is automated using statistical learning approaches based on re-projection error, smoothness, and over-fitting indicators. With regard to automating discontinuous UQ methods, Lucor [80] performed a study about adaptive methods for uncertainty quantification. Among them were the generalised Polynomial Chaos (gPC) method as proposed by Lucor and Su [81], the Padé-Legendre approximation detailed in Chantrasmi [19], Wiener-Haar Wavelets approaches by Le Maıtre [69] and the Simplex Stochastic Collocation (SSC) method developed in Witteveen [153] and Edeling [39]. Most recent works in the area move towards the strategy suggested in this work and adapt techniques from machine learning like Gorodetsky [53] who created an automatic labelling procedure of different model regions using polynomial annihilation.

To begin with, the terminology and notation are adjusted to be more suitable for least squares pseudospectral approximations of Padé approximations [134].

4.1 Matrix Notation

Let $u(\xi)$ describe the response surface of a physical model whose input uncertainty is described through the vector of random variables ξ. The Padé Legendre approximation conventionally relies on Legendre polynomials. This restricts the probability space to random variables with uniform Probability Density Function (PDF) $\rho(\xi)$. Polynomial basis functions that are optimal for other PDFs are trivial to create with SAMBA, but less well-documented in the context of Padé approximations. The function $u(\xi)$ is in $L^2[-1,1]$ and

can be approximated through the pseudospectral approximation

$$u(\xi) \approx \sum_{i=1}^{c(M)} \hat{p}_i \psi_i(\xi) \qquad \hat{p}_i = \int u(\xi)\, \psi_i(\xi)\, \rho(\xi)\, d\xi. \tag{4.1}$$

The number of linear combination terms is $c(M) = \frac{(d+M)!}{d!N!}$ for a total order polynomial expansion and $c(M) = (M+1)^d$ for a tensor product expansions [12]. The coefficients $\hat{p}_i$ can be calculated through the (possibly weighted) discrete Fourier transform

$$\hat{p} = \Psi u \ \ (unweighted) \quad or \quad \hat{p} = \Psi W^2 u \ \ (weighted). \tag{4.2}$$

The diagonal matrix W contains the weights belonging to the tensor grid collocation points γ_i. The matrix Ψ contains the $c(M)$ orthogonal Legendre polynomials evaluated at the collocation points. Usually, the points are chosen as the tensor product of univariate quadrature rules. Here, however, a more general set-up is aimed for: N is the number of collocation points and can be chosen arbitrarily. The matrices Ψ and W can be set up as

$$\Psi = \begin{bmatrix} \psi_1(\gamma_1) & \cdots & \psi_{c(M)}(\gamma_1) \\ \vdots & \ddots & \vdots \\ \psi_1(\gamma_N) & & \psi_{c(M)}(\gamma_N) \end{bmatrix}, \; W = \begin{bmatrix} \sqrt{w_1} & \cdots & \\ \vdots & \ddots & \vdots \\ & & \sqrt{w_{c(M)}} \end{bmatrix} \tag{4.3}$$

The vectors $\hat{p}$ and u are defined as $\hat{p} = \begin{bmatrix} \hat{p}_1 & \cdots & \hat{p}_m \end{bmatrix}$ and $u = \begin{bmatrix} u(\gamma_1) & \cdots & u(\gamma_N) \end{bmatrix}^T$. Note that the polynomial coefficient vector can be truncated to $m \leq N$. In this case, fewer basis functions are fitted to the same points. Another possibility is the use of Least Angle Regression (LAR) to determine the most active set of basis functions [8]. In general, the number of model evaluation points should be equal to or larger than the number of basis functions. The authors even had some success solving under-determined systems using least-norm solutions. This could potentially allow the use of even fewer points [32]. However, approximation errors and uniqueness need to be explored further. The multivariate polynomial basis functions ψ_i are calculated as the product of univariate Legendre polynomials $\prod_{i=1}^{d} \psi_{I_{i,j}}(\xi_i) \ \forall j \in 1, ..., c(M)$. The multi-index $I_{i,j}$ is used to reference the polynomial order of the individual univariate components of the tensor product space. For example, for two random variables it can be defined as $I_{i,j} = \begin{bmatrix} [0,0]^T, & [0,1]^T & [1,0]^T & [0,2]^T & [1,1]^T & \cdots & [M_1, M_2]^T \end{bmatrix}$ for

$j = 0, 1, 2...$ $I_{i,j}$ and the corresponding number of basis functions $c(M)$ can be defined in various ways. Here, total index sets are used. They are constrained by the number of sub-samples used, and it is recommended to define $c(M)$ as the number of rows in the index matrix. For simplicity with two parameters, Chantrasmi [20] defined $c\,(M) = (M+1)\,(M+2)\,/2$. By introducing the matrices

$$A = W^T \Psi^T \qquad b = W^T u \tag{4.4}$$

one can solve Eq. (4.2) as a weighted linear least squares problem in L^2

$$\hat{p} = \arg\min \|Ap - b\|_2^2 \,. \tag{4.5}$$

Several authors have previously shown that computing the Fourier coefficients using least squares approximation is computationally significantly cheaper than employing quadrature techniques [111]. The drawback of this computational advantage is that some random sets of collocation points - for instance, an unlucky set of Latin-Hypercube or pseudo-random samples - can lead to a severely ill-conditioned least squares problem and consequently large approximation errors. Seshadri et al. [134] therefore proposed a new deterministic method called Optimal Quadrature Sub-sampling that aims at obtaining the best least squares estimate by reducing the condition number of the linear system for a given set of basis functions.

4.2 A Stabilised and Automated Least Squares Padé Approximation

This section details the numerical method to obtain accurate representations of complex non-linear and discontinuous response surfaces. All Polynomial Chaos components are computed using SAMBA. Two innovative methods that complement each other are added. The first component is the multi-variate Padé approximation. As a global non-intrusive approximation method it provides a very efficient way to determine discontinuous response surfaces using a very low number of physical model evaluations. However, it has known stability issues. It is therefore combined with a stability regularisation method called Optimal Quadrature Sub-sampling. In order to do so, various adaptations are necessary. These will be described in the sections below. Finally, statistical learning methods are used to automatically ensure reliability and convergence of the combined method.

4.2.1 New Least Squares Algorithm

A Padé approximation is a rational function of two polynomials. In a deterministic setting, it has been used as a generalisation of the Taylor series expansion to calculate an approximation of a function given its power series. In the context of Uncertainty Quantification, a Padé approximation of a function u(ξ) of a uniform random variable vector ξ is the ratio of two Polynomial Chaos expansions $P(\xi)$ and $Q(\xi)$

$$u\left(\xi\right) \approx \hat{M}_{Pade}(\xi) = \frac{P\left(\xi\right)}{Q\left(\xi\right)} = \frac{\sum_{j=1}^{c(M)} \hat{p}_j P_j\left(\xi\right)}{\sum_{j=1}^{c(L)} \hat{q}_j Q_j\left(\xi\right)}, \tag{4.6}$$

with numerator order $c(M)$ and denominator order $c(L)$. The coefficient vectors $\hat{p}$ and $\hat{q}$ need to be determined in order to determine a rational pseudospectral approximation of u. The terms P_j and Q_j refer to the multivariate products of the univariate Legendre series for ξ ordered by the multi-index matrix $I_{i,j}$. Our multivariate Padé approximation is based on Chantrasmi [19]. It differs from Chantrasmi's in the sense that it allows the use of an arbitrary set of collocation points to combine the method with Optimal Quadrature Subsampling or an industrial database. Let $\psi_j(\xi_1)$ be a Legendre function with M_1 terms. Then, $\psi_j(\xi_1) = 1, \xi_1, \frac{1}{2}\left(3\xi_1^2 - 1\right), \frac{1}{2}\left(5\xi^3 - 3\xi_1\right)$, for $j = 1, 2, .., M_1$. The discrete scalar product corresponding to $\rho(\xi)$ is defined through numerical Gauss quadrature as

$$\langle P, Q\rangle_M = \sum_{i=1}^{M+1} P\left(\gamma_i\right) Q\left(\gamma_i\right) w_i. \tag{4.7}$$

with optimal quadrature points matrix Γ, whose rows are denoted as γ_i, and optimal weights vector w. With a quadrature rule of order M, the integral in Eq. (4.7) can be accurately solved for polynomials P_j, whose univariate components $\psi_{I_{i,j}}$ both have a smaller maximum order than $2M - 1$. For Legendre polynomials, the orthogonality condition of the multivariate polynomials P_j can be obtained through the simple formula $\langle P_i, P_j\rangle = 1/((i + 1/2)\left(j + 1/2\right))\delta_{i,j}$ for $i, j \in N$.
An overdetermined system of linear equations to determine $\hat{q}$ can be derived by projecting Qu onto the basis functions $c(M)$ to $c(M + k)$:

$$\langle Qu, P_j\rangle_{c(M)} = \langle P, P_j\rangle_{c(M)} = 0 \quad \forall j = c(M) + 1, ..., c(M + k) \tag{4.8}$$

which can be rearranged it into

$$\begin{bmatrix} \left\langle uP_1, P_{c(M)+1} \right\rangle_{c(M)} & \cdots & \left\langle uP_{c(L)}, P_{c(M)+1} \right\rangle_{c(M)} \\ \vdots & \ddots & \vdots \\ \left\langle uP_1, P_{c(M+K)} \right\rangle_{c(M)} & \cdots & \left\langle uP_{c(L)}, P_{c(M+K)} \right\rangle_{c(M)} \end{bmatrix} \begin{bmatrix} \hat{q}_1 \\ \vdots \\ \hat{q}_{c(L)} \end{bmatrix} = 0. \quad (4.9)$$

The index i in Eq (4.9) runs from $c\,(M)$ to $c\,(M+k)$ because the projection error of the multi-dimensional Padé approximation onto the Legendre basis is only zero up to order $c(M)$. From $c(M)$ to $c(M+k)$ the projected error is merely minimized in the least-squares sense and the condition $c\,(M+k) - c\,(M) > c\,(L)$ must be uphold to ensure that Eq (4.9) is always overdetermined. The variable $k \in \mathbb{N}$ should be greater than 1 and it is beneficial for stability to keep $k \geq 2$.

As suggested by Chantrasmi [19], Eq (4.9) can be rewritten in an easier to implement way by introducing the three matrices

$$\begin{aligned} B_{i,j} &= P_i\,(\gamma_i)\,w_j \quad\;\; c\,(M)+1 \leq i \leq c\,(M+k)\,, \quad 0 \leq j \leq c(M) \\ C_{i,j} &= u_i\,(\gamma_i)\;\delta_{ij} \quad 0 \leq i \leq N, \quad 0 \leq j \leq N \\ D_{i,j} &= P_j\,(\gamma_i) \quad\;\; 0 \leq i \leq c(M), \quad 0 \leq j \leq l \end{aligned} \quad (4.10)$$

The denominator design matrix A_q is then $A_q = BCD$. One important detail here is that the numerical quadrature of the inner products requires the solution of u in all n numerical quadrature point locations. This is a temporary simplification and it will be removed in future versions of the code. However, since Q is merely used to precondition the actual polynomial approximation, the caused error is not large and decreases by adding points. Based on A_q, one can find q through solving the minimisation problem

$$\min_{\|\hat{q}\|=1} \|A_q \hat{q}\|^2 \quad (4.11)$$

It can be uniquely solved using a Single Value Decomposition of $A_q = USV$, where V contains the eigenvalues of $A_q^T A_q$. The sought solution is then given by the last column of V $q = Ve_1$, because the first eigenvalue is the smallest. The matrix A_q is of size $(c\,(M+k) - c\,(M)) \times c\,(L)$. This procedure is known as generalised least-squares Padé approximation [55]. In it, however, only the denominator coefficients are found using weighted regression. The polynomial coefficients $\hat{p}$ of the numerator expansion $P(\xi)$ are calculated by employing a quadrature formula. To combine the method with Optimal Quadrature Sub-sampling, the coefficients need to be fitted to arbitrary points using a

least squares regression. This is straightforward if the problem is regarded as standard non-intrusive Polynomial Chaos problem with pre-conditioner matrix Q:

$$\hat{p} = \arg\min \|\hat{p}\Psi - Qu\|_2^2 . \tag{4.12}$$

For higher generality, it is useful to replace the denominator polynomial index $c(L)$ with a manual integer L. Using $c(L)$ can limit the model to an unnecessary high order for the denominator during the automatic fitting process. Moreover, $c(L)$ excludes the possibility that the best fit is continuous. L is hence used here in order to keep the method applicable to both continuous or discontinuous models.

4.2.2 Stabilisation via Optimal Quadrature Sub-sampling

Optimal Quadrature Sub-sampling (OQS) is a design of experiment (DoE) approach developed by Seshadri [134] that helps to find well-conditioned least squares estimates. Instead of solving the least squares problem in Eq. (4.5) in a set of arbitrary random points [61, 111] or in all available tensor grid collocation points [75], the authors suggest a sampling strategy that reduces the condition number of the design matrix $A = W^T\Psi^T$. To that end, the authors perform a low rank approximation via rank-revealing QR factorization of the design matrix A. Since it is assumed that N physical model evaluations are available, the design matrix can be approximated with

$$A = Q_N R_N P_N \underbrace{+\tilde{A}_N P_N}_{error} . \tag{4.13}$$

The matrix Q_N holds the chosen columns N of A and the matrix $\tilde{A}_N$ holds the remaining columns, which have not yet been chosen. Its Frobenius norm can be used to describe the low rank approximation error. The permutation matrix P_k reveals which choice of collocation points will provide the best conditioned solution to the least squares approximation problem. The vector containing P_k's first k diagonal elements is denoted as π_{OQS}. This is the optimal sub-sampling index set. Subsequently, the pruned and better conditioned minimisation problem

$$\hat{p}_* = \arg\min \|A(\pi_{OQS})\hat{p} - b(\pi_{OQS})\|_2^2 \tag{4.14}$$

can be solved. An elegant Matlab implementation of OQS can be found in the authors' original paper [134].

4.2.3 Automation Using Statistical Learning

With the Least Squares Padé Approximation described above, a transonic CFD code can be evaluated in an arbitrary set of parameter collocation points. This is most useful for the industrial application of the method, where an already available database of test cases can be recycled. In addition, CFD codes do occasionally not converge in optimal quadrature locations. This issue can be easily avoided using arbitrary points. To ensure a reliable fit leave-one-out cross-validation is used and the least squares problem is solved using QR decomposition. When starting from a manually chosen set, it found in this work that further samples should be added using the following three principles. First, the domain should be equally covered. Second, regions with high gradients are to be preferred. Third, the position of additional points should be chosen so that the condition number is reduced. If CFD convergence issues occur, it can be beneficial to interpolate to the numerically more efficient points, bearing in mind that proximity is key for accuracy. A concise summary of the most relevant lines of Matlab codes is shown below.

```
% numerator design matrix
A_p = sqrt(w)*Psi_matrix;
% QR pivoting of design matrix
[¬,¬,Pi] = qr(A_p', 'vector');
% select first l entries
Pi = Pi(1:l);

% denominator design matrix
A_q = B_q*C_q*D_q;
% determine denominator coefficients
[U,S,V] = svd(A_q); q_coef = V(:,end);

% use Q as preconditioner
best = sqrt(w(Pi))' * (Q(Pi).*u_sol)';
% determine numerator coefficients
p_coef = A_p(Pi,:) \ best;
```

1. The approximation should have a small re-projection error (goodness of fit): $u(\gamma_1, \gamma_2) - \frac{P(\gamma_1,\gamma_2)}{Q(\gamma_1,\gamma_2)} \approx 0$
2. The approximation should have no discontinuities or overshoots apart from the physical discontinuities actually present in the data (smoothness condition)

The main advantage of the least squares Padé approximation is that many numerator and denominator polynomials can be fitted to a given set of points. Therefore, criteria are needed to determine the most suitable parameters for M and L. The goodness of fit of the surrogate model is usually evaluated using an L_2-norm error. Minimising only the re-projection error, however, can lead to the well-known problem of over-fitting. This means that the fitted model only describes the solution in the given data points, but not in between. The effect is best known for causing high amplitude oscillations between points when fitting higher order polynomials under the name of 'Runge Phenomenon'.

Over-fitting errors can be reduced by minimising the prediction error instead of the re-projection error. For a given dataset, the points are divided into learning and training set, so that the model's predictive behaviour can be evaluated from training data that were not involved in the determination process. In this work, Leave-one-out cross-validation is used. K-fold cross-validation is usually more reliable. However, since the number of collocation points in the used examples is very low omitting many points for model determination is impractical [111, 139].

For the special case of linear regression, a very elegant theorem is available that allows determining the Leave-one-out error ε_{LOO} in a single model fit instead of having to perform N fits for N points. The result can be quickly derived by extending Eq (4.2) using the normal equation

$$\hat{u} = \Psi\hat{p} = \Psi\left(\Psi^T\Psi\right)^{-1}\Psi^T u = Hu \tag{4.15}$$

If the vector h is defined to contain the diagonal elements of H: $h = diag(H)$, the Leave-one-out error can be computed as

$$\varepsilon_{Loo} = \frac{1}{N}\sum_{i=1}^{N}\frac{u_i - \hat{u}_i}{(1-h_i)^2} \tag{4.16}$$

To quantify unwanted discontinuous outbreaks in the response surface, Chantrasmi [20] suggested the use of the total variance as an indicator for smoothness

$$\tau_{TV} \approx \sum_{i=1}^{N-1}\sum_{j=1}^{N-1}|\hat{u}_{i+1} - \hat{u}_i|. \tag{4.17}$$

By combining these indicators, the optimal trade-off between prediction error and smoothness can be chosen for a given number of points by solving a multi-

objective optimisation problem with cost function

$$\arg\min_{M,L} \left(\omega_1 \log\left(\varepsilon_{LOO}\right) + \omega_2 \frac{\tau_{TV}}{\max\left(\tau_{TV}\right)} mean\left(\log\left(\varepsilon_{LOO}\right)\right) \right). \tag{4.18}$$

where ω_1 and ω_2 are the weights of the different objectives. The normalisation and multiplication by the error mean is added to ensure equal weighting of both quantities. Moreover, the logarithm of the errors makes it easier for an optimisation method to identify differences. Both weights, ω_1 and ω_2, are set to one in the following examples because an equal importance of the objectives is assumed a priori. It is possible to choose the weights differently if one of the two objectives is preferable for an application. For some applications it may be possible that the linear cost function does not yield a satisfactory result. For such cases, a different approach than the linear combination method may be more suitable: both the global criterion method or the ϵ-constraint method, tend to yield better results for multi-objective optimisation problems. For more information see [22].

To ensure the reliability of this method in the presence of discontinuities, Eq (4.18) has to be complemented with an Inf and NaN filter. All values larger than 10 are consequently filtered from the error and smoothness arrays and replaced with one to recognise them. Figure 4.2 shows a variety of model evaluations for different parameters of m and l and the final choice of the method.

4.3 Validation

The OQS Padé Matlab Algorithm is demonstrated by applying it to a simple analytical example: two fairly smooth domains created by two-dimensional sinus curves divided by a sharp rising edge created by a hyperbolic tangent function:

$$z = \tanh\left(10\xi_1\right) + 0.2\sin\left(10\xi_1\right) + 0.3\xi_2 + 0.1\sin\left(5\xi_2\right) \tag{4.19}$$

The resulting 3D surface is taken after the response surfaces obtained for the thermal efficiency in the later described film cooling shock-interaction problem. It can be seen in Figure 4.1 how the choice of optimal sub-samples from the tensor rule allows reconstructing the previously used *tanh* response surface with high accuracy using only 18 collocation points. In comparison, the full tensor quadrature rule with a full 16 points misses most of the details and

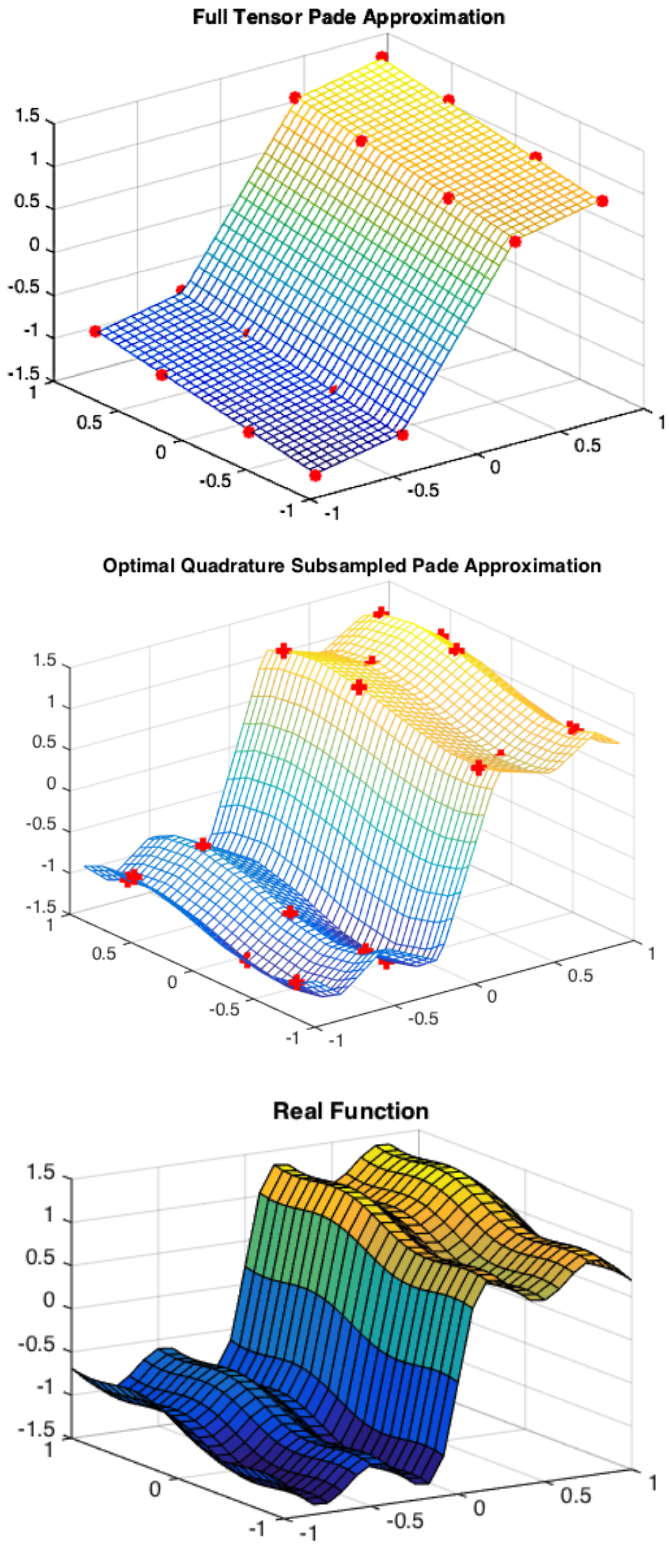


Figure 4.1: Response surfaces of a hyperbolic tangent function recreated from a low number of data points. The surface at the top was determined using the classic Padé approximation based on 16 Gaussian quadrature points. The surface in the middle is based on the new Padé approximation developed in this work based on 18 optimally sub-sampled points.

looks more like a linear approximation.

The graph in Figure 4.2 shows the Pareto front for leave-one-out error and smoothness parameters. It can be seen that the smoothness is more important for the Padé approximation than for common approximations. The error alone can be minimised even if there are singularities between the given collocation points. A reliable surface can be chosen only if the variation of the area between the points is taken into account and convergence is reached in the values. The found response surface is obtained for $[M, L, N] = [6, 1, 18]$. The mean and variance are estimated as $\mu = 1.97e-04$ and $\sigma = 0.9476$ compared to Monte Carlo Simulation with 10^7 samples of $\mu = -3.45e-04$ and $\sigma = 1.009$. The method converged because $\epsilon_{L2} = 4.43e-15$ and $Err_{Loo} = 1.0778e-59$ and the smoothness error is 1.32.

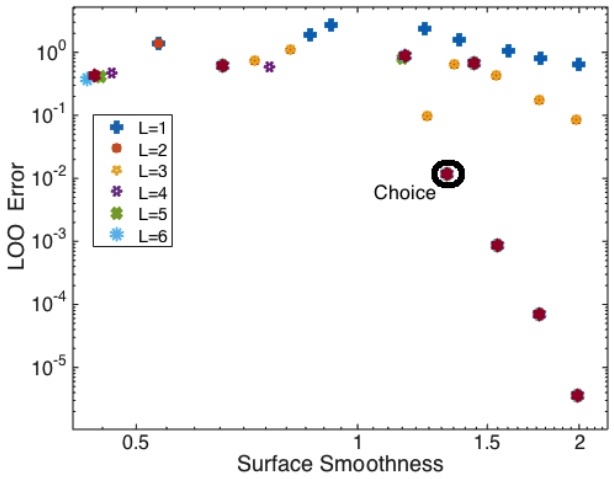


Figure 4.2: Demonstration of how an automatic trade-off between the leave-one-out error and the surface smoothness parameter is determined.

5 Multi-Fidelity to Reduce Model-Form Uncertainty

One of the biggest problems that uncertainty causes in industrial simulations is rarely discussed in the academic literature: epistemic model-form uncertainty. The reason is that it is very difficult to find purely statistical or mathematical solutions for it. Probabilistic methods cannot make up for the fact that simplifying assumptions in physical models misrepresent the reality. Engineers in the industry are more often interested in understanding this bias than the stochastic variability of their predictions. However, there is one completely general tool in the arsenal of uncertainty quantification that can reduce model-form uncertainty: multi-fidelity methods.

In this chapter, two state-of-the-art non-intrusive multi-fidelity frameworks based on Polynomial Chaos are tested and modified to combine two different kinds of CFD simulations: Reynolds Averaged Navier-Stokes (RANS) and Direct Numerical Simulations (DNS). It corrects the wrong behaviour of the RANS solution using data from an immersed boundary DNS code [67, 66]. DNS is the most accurate but also most expensive method currently available in CFD whereas RANS is the most frequently used approach in industry but less accurate. More details about the methods can be found below.

The original aspect of the multi-fidelity methods suggested in this chapter is that they are intended for the sole purpose of uncertainty quantification and robust design. Modelling errors in the lower fidelity model are not corrected by calibrating the model but within a statistical post-processing step. Thus, the RANS model still contains errors, but a design exploration study including RANS and DNS data, performed during post-processing, will be RANS model error-free. The advantage of this approach is that reliability analysis or optimisation can be performed at higher physical accuracy without requiring any interference with the physical models. This concept of non-intrusiveness makes the method quite widely applicable. It can be trivially applied to any combination of high and low fidelity models within industry.

The core concept of the approach is that the necessary high number of samples needed for statistical reliability is achieved using the lower fidelity

model (here RANS) and the physical accuracy is ensured using the higher fidelity model (here DNS). The main objective is to reduce the computational cost because purely DNS based uncertainty quantification studies are still far too expensive for widespread use. This is another novelty of this section. For the first time, enough DNS simulations of a CFD case were performed to conduct a meaningful non-intrusive uncertainty quantification study.

The eventual goal of this multi-fidelity research is to find an uncertainty quantification framework that can accurately and efficiently quantify the effect of model uncertainty in high safety applications such as aircraft engineering. In particular, geometric profile manufacturing errors are relevant in this field. However, manufacturing process tolerances [141] and destructive effects taking place during operation (like fouling [119, 17]) inevitably create components that vary individually and randomly from the design intent. Although these variations are usually small, their impact can become disproportionately large under turbulent flow conditions [13, 38]. Many authors have therefore attempted to quantify the effects of manufacturing-induced geometric variability by combining CFD codes with probabilistic methods. In the following, two multi-fidelity Polynomial Chaos methods, one based on SAMBA and one based on a least-squares method, will be used to account for these variations and reduce model-form uncertainty.

5.1 Multi-Fidelity Polynomial Chaos

The concept of exploiting multiple levels of fidelity has only recently entered the field of UQ. Giles [49] proposed to speed-up the solution of stochastic partial differential equations using multi-level Monte Carlo sampling. Following their work, Ng [103] implemented a multi-level Monte Carlo method for engineering models. For CFD, Doostan [33] suggested a bi-fidelity model combining denser and coarser grid within a RANS simulation framework. Within the field of Polynomial Chaos, two authors recently developed new methods: Ng [102] developed an approach using different levels of sparse grids and Palar [111] developed an approach using a multi-level point collocation. This work builds upon these methods developed by Ng and Palar. They are only modified slightly. Unlike Ng [103] sparse grids based on optimal Gaussian quadrature points are used, which are generally not nested. However, SAMBA as described in Chapter 3 be set up in a nested way. Unlike Palar [111], the optimal Gaussian collocation points are reused for the second approach of weighted linear regression. More importantly, a weighted regression method

is introduced that uses weights that are derived from the model correlation between RANS and DNS.

5.1.1 Nested optimal SAMBA grids

First, the formulas from SAMBA that are needed for the multi-fidelity setup method are quickly repeated. For d arbitrary independent random variables ξ_j, the univariate Gaussian quadrature rules based on SAMBA is defined as $\{U^j\}_{j=1,...,d}$ as

$$U^j = \sum_{k=1}^{m_j+1} f(\gamma_{j,k})\omega_{j,k}, \tag{5.1}$$

where m_j is the order of the individual polynomial expansion for random variable ξ_j and the vectors $\gamma_{j,k}$ and $w_{j,k}$ with $k = 1, ..., m_j + 1$ are the optimal quadrature collocation points and weights for probability distribution j. An adaptive and anisotropic sparse Smolyak quadrature operator $S(q, l)$ for d variables and of level $l \in \mathbb{N}$, where the convention is followed by introducing the abbreviation $q = d + l$, can be defined as a recursion formula of the tensor product $\otimes$ as

$$S\left(q,d\right) = \sum_{\substack{l+1 \\ \leq |i| \leq \\ l+d}} (l-1)^{q-|i|} \begin{pmatrix} d-1 \\ q-|i| \end{pmatrix} \otimes_{j=1}^{d} U^j \tag{5.2}$$

where $|i| = I_{i,1} + ... + I_{i,d}$ is the degree of the i^{th} row of the multi-index matrix $I_{N_p,d}$. It is used to truncate the Polynomial Chaos expansion from the tensor product set to the sparse Smolyak set. Equation (5.2) needs only a reduced number of quadrature collocation points and weights collected in the arrays Γ and W. For higher levels, sparse grids with more points can be obtained that integrate a function $f(\xi)$ with increased accuracy.

To set up a multi-fidelity framework, the RANS model response needs to be evaluated f_{RANS} in more points than the DNS response f_{DNS}. A difference in levels of $r \in \mathbb{N}$ is assumed between RANS and DNS and defined as

$$f_{RANS}\left(\xi\right) \approx S_{q,d}\left[f_{RANS}\left(\xi\right)\right] \qquad f_{DNS}\left(\xi\right) \approx S_{q-r,d}\left[f_{DNS}\left(\xi\right)\right]. \tag{5.3}$$

The resulting rule is nested because the lower-order DNS grids are a subset of the highest-level RANS grid, that is, $S_{q-r,d} \subset S_{q,d}$ for $r \in 1, .., l-1$. For example, a level three grid includes the sparse grids of level one and level two. Figure 5.1a illustrates this concept. The shown DNS points are obtained from a level one grid and the RANS points from a level three grid. Although this

is not clearly visible, the DNS points are also contained in the RANS grid.

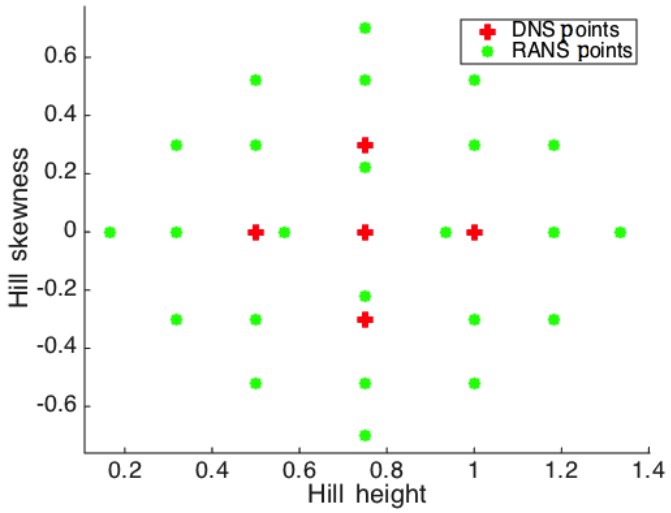


Figure 5.1: A multi-fidelity nested Gaussian sparse grid comprised of five DNS and 29 RANS solutions.

The method to combine multiple grids to obtain Polynomial Chaos coefficients is described Ng [102]. With such a nested framework, the quality of a DNS response surface can be approximated by correcting the RANS response through DNS data on the lower level sparse grids

$$\tilde{f}_{DNS}(\xi) \approx S_{q,d}\left[f_{RANS}(\xi)\right] + S_{q-r,d}\left[C_{DNS}(\xi)\right]. \tag{5.4}$$

The correction function can be either additive or multiplicative

$$C(\xi) = f_{DNS}(\xi) - f_{RANS}(\xi) \qquad C(\xi) = \frac{_{DNS}(\xi)}{f_{RANS}(\xi)} \tag{5.5}$$

or a combination of the two [102]. An elegant side effect of this approach is that the multi-fidelity notation can be easily transferred from the quadrature rule to the Polynomial Chaos coefficients. This is practical because Polynomial Chaos coefficients can be used to calculate output moments, parameter sensitivities and skewness metrics at the fraction of the cost of other sampling methods [134]. Applying the sparse quadrature operators $S_{q,d}$ to the uncorrected RANS higher level expansion gives

$$S_{q,d}\left[f(\xi)\right] = \sum_{i \in I_{q,d}} \alpha_i \Psi_i(\xi). \tag{5.6}$$

Similarly, the lower level expansion DNS correction function can be expressed as

$$S_{q-r,d}\left[C\left(\xi\right)\right] = \sum_{i \in I_{q-r,d}} \alpha_{C,i} \Psi_i\left(\xi\right). \tag{5.7}$$

In this way, Equation (5.4) can be reformulated using Polynomial Chaos expansions as

$$\tilde{f}_{DNS}\left(\xi\right) \approx \sum_{i \in I_{q,d} \setminus I_{q-r,d}} \alpha_{RANS,i} \Psi_i\left(\xi\right) + \sum_{i \in I_{q-r,d}} \left(\alpha_{RANS,i} + \alpha_{C,i}\right) \Psi_i\left(\xi\right). \tag{5.8}$$

With this methodology only the lower order polynomial coefficients are corrected by DNS. The higher order coefficients are calculated from the RANS model alone. This is advantageous for parameters, for which RANS and DNS are well correlated, but the RANS is biased.

5.1.2 Weighted Regression

Instead of using numerical quadrature, the Polynomial Chaos coefficients can also be retrieved by solving a linear system of equations. This approach is called point-collocation non-intrusive Polynomial Chaos. The system that needs to be solved is

$$\begin{bmatrix} \Psi_1\left(\eta_1\right) & \cdots & \Psi_{N_p}\left(\eta_1\right) \\ \vdots & \ddots & \vdots \\ \Psi_1\left(\eta_N\right) & & \Psi_{N_P}\left(\eta_N\right) \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \vdots \\ \alpha_N \end{bmatrix} = \begin{bmatrix} f\left(\eta_1\right) \\ \vdots \\ f\left(\eta_N\right) \end{bmatrix}. \tag{5.9}$$

The sample array η in Equation (5.9) can be filled with an arbitrary choice of model evaluation points of $N \geq N_p$: random samples, quasi-random samples, tensor product or sparse collocation points. Reasonable accuracy can be achieved if this system is oversampled [60] and least squares linear regression is used to solve it

$$\hat{\alpha} = \arg\min \left[\alpha^T \Psi\left(\xi\right) - f\left(\xi\right)\right]^2 \tag{5.10}$$

For a few random variables this can be done using the normal equation: $\hat{\alpha} = \left(\Psi^T \Psi\right)^{-1} \Psi^T Y$, whereas for higher numbers of variables the gradient descent method becomes increasingly preferable. As CFD data generated for the nested grid approach described in the previous section are reused for the regression approach, no quasi-random samples used. Instead, the DNS solutions obtained in the sparse grid collocation points are used. This approach is often referred to as probabilistic collocation and more details can be found in [78].

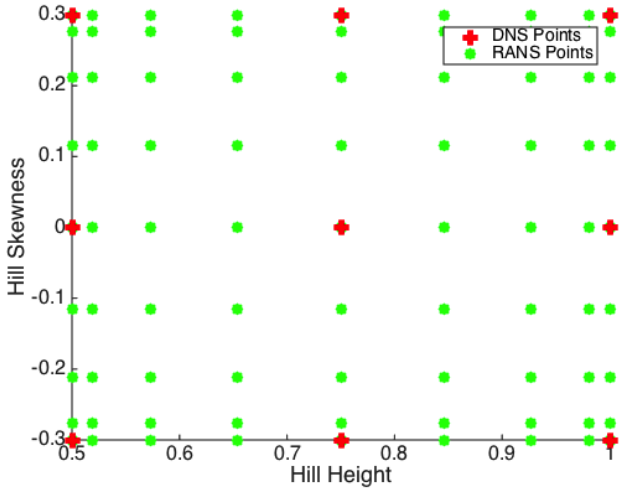


Figure 5.2: The multi-fidelity grid used for the weighted regression comprised of nine DNS and 81 RANS solutions.

Instead of solving the linear regression problem directly it is transformed it into a weighted least squares problem

$$\hat{\alpha} = \left(\Psi^T \lambda \Psi\right)^{-1} \Psi^T \lambda Y. \tag{5.11}$$

As weight vector λ is used. It contains the model correlation coefficient R^2 between RANS and DNS model. The R^2 correlation between RANS and DNS model is characterised as

$$R^2 = \left(\frac{\sum_{i=1}^{N} \left(y_i^{[DNS]} - \mu^{[DNS]}\right)\left(y_i^{[RANS]} - \mu^{[RANS]}\right)}{\left(\sum_{i=1}^{N} \left(y_i^{[DNS]} - \mu^{[DNS]}\right)^2\right)\left(\sum_{i=1}^{N} \left(y_i^{[RANS]} - \mu^{[RANS]}\right)^2\right)} \right)^2 \tag{5.12}$$

where $\mu^{[DNS]}$ and $\mu^{[RANS]}$ are the means of all y_i for DNS and RANS, respectively. The reason for this modification of the method [111] is that it provides an automatic indicator for the multi-fidelity scheme to decide if modelling errors affect a quantity of interest and how the two different models should be weighted. When RANS is more accurate for a quantity of interest, the correlation is better and a higher weight is assigned. If the correlation is low, because the RANS model is too inaccurate, the DNS solution is weighted

more. In addition, the mean absolute relative error $|M|_{rel}$ defined as

$$|M|_{rel} = \frac{1}{N} \frac{\sum_{i=1}^{N} \left| y_i^{[DNS]} - y_i^{[RANS]} \right|}{y^{[DNS]}} \tag{5.13}$$

is used to quantify the offset between the two models. More details about multi-fidelity point collocation can be found in [111].

5.2 Demonstration for Turbulent Flow over Periodic Hills

As test case, turbulent flow over periodically arranged hills is considered at a Reynolds number of $Re = 700$ (Figure 5.3a) as described for example by Marquillie [83] or Breuer [11]. Despite its simple geometry, this test case is nevertheless challenging for RANS models. It is a widely known problem that RANS models struggle with the accurate prediction of the size and length of the separation bubble at the downstream side of the hill. This uncertainty quantification study, therefore, uses quantities of interest (QoI) that are connected to the separation bubble. A significant geometric variability is introduced by parametrising the hill shape stochastically and try to find the correct behaviour of the resulting multitude of flows as accurately and efficiently as possible. To keep the computational cost low and create a method of potential interest for industry, the RANS equations are solved with a commercial CFD code and with commonly used k-ω -SST eddy viscosity model. Due to the non-intrusiveness of the approach, the method can trivially be used with other turbulence models. For illustration purposes, mainly one model is used in this work.

In the following example, the focus is on eddy viscosity models based on the Boussinesq hypothesis, like k-ϵ [68], k-ω [151] and k-ω SST models [86]. They approximate the effect of the Reynolds stresses on the mean flow field by prescribing a turbulent viscosity model [116]. While this works well for fully attached turbulent flows, there are usually significant discrepancies for flows close to separation or where the streamline curvature is large [120]. Various innovative methods have been suggested in recent years to approach this problem. Dow [34] suggested solving an inverse problem using an adjoint method to correct the turbulent viscosity field of the RANS model using DNS data. Other authors proposed a physics-based Bayesian framework [155] or the use of random matrix theory [157]. Most recently, Wang proposed a data-driven,

physics-informed machine learning approach for predicting discrepancies in RANS modelled Reynolds stresses [147, 156].

All flow conditions are defined identically to the DNS reference case described by Breuer in [11]. The details for the DNS simulation are described in subsection 5.2.1 and for the RANS simulations in subsection 5.2.2. Both the DNS and RANS simulations are performed on a normalised domain with height $h = 1$, density $\rho = 1$ and bulk velocity $U_b = 1$ [11]. In physical dimensions these parameters become $h = 0.028[m]$, $\rho = 1.225[kg/m^3]$ and $U_b = 0.3437\ [m/s]$ with a resulting mass flow rate of $\dot{m} = 0.024[kg/s]$, see Hellsten [58]. The boundary conditions are of a periodic type in the streamwise and span-wise directions and zero-velocity type in the vertical direction. The Reynolds number of $Re = 700$ is based on the bulk velocity above the hill U_b and the hill height h: $Re = hU_b/\nu$ in the reference test case. This definition is not applicable to the range of simulations here because the hill height varies. Therefore, the mass flow rate is calculated based on the reference case and kept constant in both DNS and RANS for all simulations.

The reference test case by Breuer [11] is chosen because a highly reliable experimental and computational database is available. The flow separates from the curved surface of the hill, recirculates in its downstream side and reattaches at the flat channel bottom. Although this is a simple scenario, CFD methods based on solving the RANS equations still struggle with predicting the size and length of the separation bubble. It is important for this work that the shape of the curved surface strongly affects the reattachment position. Consequently, considerable model-form uncertainty can be expected from RANS simulations if geometrical uncertainty is present in the hill shape. Geometric uncertainty is introduced by parametrising the hill shape with two variables: a hill height parameter h and a hill skewness parameter s. The mean height is chosen at 75% of the reference case and the hill shape is skewed by asymmetrically stretching and compressing the upstream and downstream sides of the hill. For example, $s = 0.3$ means that the upstream part of the hill is stretched by a factor of 1.3, whereas the downstream side is compressed by a factor of 0.7. The two normal distributions $N(\mu, \sigma^2)$ for the height and skewness are defined as

$$h \sim N\left(0.75, 0.25^2\right) \qquad s \sim N\left(0, 0.3^2\right). \tag{5.14}$$

To avoid drawing random samples from outside of a physically meaningful domain, both distributions are truncated so that $h \in [0, 1.1]$ and $s \in [-0.3, 0.3]$. For the Polynomial Chaos analysis Hermite polynomials based on the Gaussian distribution are used. The computational domain is shown in Figure 5.3a

and the geometric uncertainty created by varying hill size and skewness is shown in Figure 5.3b.

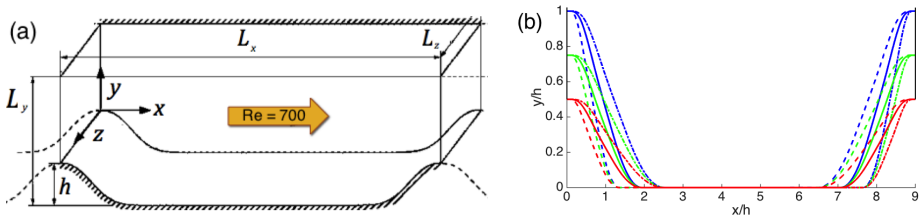


Figure 5.3: (a) Computational domain used for the CFD simulation.
(b) Illustration of how the hill shape is changed parametrically to create the assumed geometric uncertainty.

5.2.1 Direct Numerical Simulation

The high-fidelity 3D simulations presented in this study and shown in Figure 5.3 are performed for a non-dimensionalised computational domain of size $9h \times 3.036h \times 4.5h$, where h is the height of the hill. The domain is discretised on a regular Cartesian mesh with $512 \times 257 \times 128$ mesh nodes. The mesh is slightly refined in the y direction near the walls following the recommendation of Breuer [11]. The spatial resolution is as good as the ones found in previous numerical studies for the same Reynolds number of $Re = 700$.

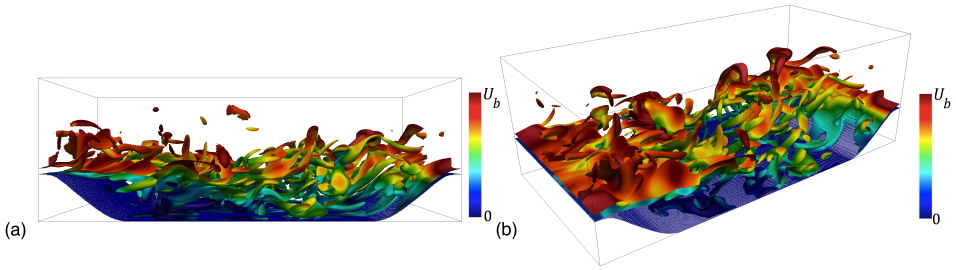


Figure 5.4: Instantaneous DNS visualisation of the vorticity modulus (for $||W|| = 5Uc/h$) coloured by the stream-wise velocity.

The coordinate system is orthonormal, with coordinate x in the stream-wise direction, y in the vertical direction and z in the span-wise direction. The boundary conditions are of periodic type in the stream-wise and span-wise directions and zero-velocity type in the vertical direction. The initial condition for the velocity field is a Poiseuille profile $(1 - (y/d)^2)$, where d is half the height of the channel $d = 3.036h/2$.

First, the simulations are launched without any hill in order to get a fully developed flow. In order to drive the flow, a pressure gradient is adjusted dynamically at each time step to maintain a constant mass flow rate through the channel. In the set-up here, the reference mass flow rate is computed from the Poiseuille profile and is equal to

$$\dot{m} = \rho \frac{1}{2h} \int_{-d}^{d} 1 - \left(\frac{y}{d}\right)^2 dy. \tag{5.15}$$

In the normalised set-up one obtains $\dot{m} = 2/3$, because $\rho = 1$ and $h = 1$. This mass flow rate is kept constant for all the simulations. Then the hill is introduced into the computational domain and the data are collected only after a transitional period when the flow is fully developed again. The time step is equal to $0.0005h/U_b$, where U_b is the bulk velocity. For all calculations, turbulent statistical data have been collected on a time period $T = 150h/U_b$, with a spatial average in the span-wise direction in order to reach a statistical convergence. The modelling of the hill is performed using a customised

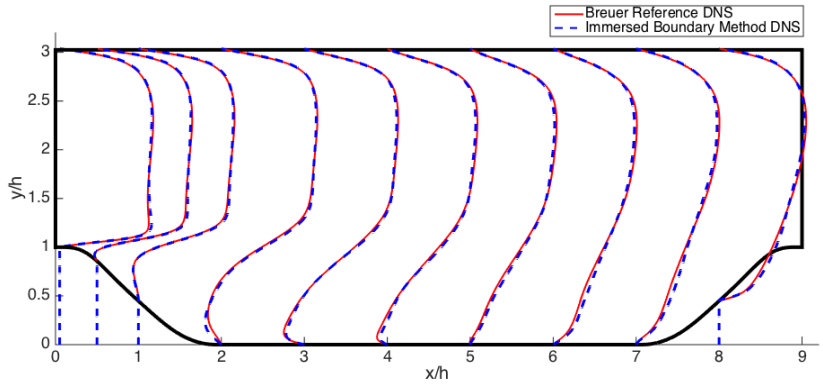


Figure 5.5: DNS validation for $Re = 700$ using horizontal velocity profiles.

immersed boundary method based on a direct forcing approach whose forcing field $\mathbf{f}(\mathbf{x}, t)$ to ensure a zero-velocity boundary condition at the wall of the hill. To solve the incompressible Navier-Stokes equations the high-order flow solver Incompact3D is used, which is based on sixth-order compact schemes for spatial discretisation on a Cartesian mesh and a third-order Adams-Bashforth scheme for time advancement. Note that the convective terms are written in skew-symmetric form. This form reduces aliasing errors while remaining energy conserving for the type of spatial discretisation considered here. To treat the incompressibility condition, a fractional step method is used to solve a Poisson equation. For efficiency reasons, this equation is solved in spectral

space using appropriate 3D Fast Fourier Transforms (FFT). In order to have a strict equivalence between finite-difference operators in physical space and spectral operators, the concept of modified wave number introduced by Lele [73] is used, which reduces the accuracy of the spectral operators to sixth-order accuracy. Note that the divergence-free condition is ensured up to machine accuracy. The pressure mesh is staggered from the velocity mesh by half a mesh to avoid spurious pressure oscillations.

A validation of the velocity field and turbulent kinetic energy comparing the present code to reference data obtained by Breuer [11] is shown in Figure 5.5. It can be seen in Figure 5.5 that the DNS solver used here agrees well with the results in the literature.

More details about the code, especially about the original treatment of the pressure in spectral space, can be found in Laizet [66]. Because of the size of the simulations, the parallel version of Incompact3D has been used for this numerical work. Based on a highly scalable 2D decomposition library and a distributed FFT interface, it is possible to use it on thousands of computational cores. More details about this efficient parallel strategy can be found in Laizet [67].

5.2.2 RANS Simulations

ANSYS Fluent is used for the 2D RANS simulations [62]. Unlike the DNS, the RANS simulations are performed in 2D. While this creates a larger discrepancy between the models, it makes a multi-fidelity method more efficient. The idea of any multi-fidelity method is to reduce the computational cost by combining an accurate expensive model, with a much faster cheap model. The cheaper this fast model is the higher is the gain from a multi-fidelity framework. Hence, even though the use of 2D RANS is less accurate, in the context of a multi-fidelity framework its use will pay off as long as the two methods are still sufficiently correlated. This will be investigated in this section. Various RANS models have been tested for their suitability in a multi-fidelity framework. Their velocity profiles are compared in Figure 5.6. Since all RANS models come with individual advantages and disadvantages, it is found that a comparative study of RANS models for multi-fidelity purposes becomes rather similar to a literature study of the suitability of different RANS models for the considered test case. Many such studies are already available in the literature [58, 118]. Since the focus of this work is on the feasibility of multi-fidelity modelling for uncertainty quantification, only one model is investigated: $k-\omega$ -SST due to its suitability for predicting separation. In Fluent, the same di-

mensionless domain is used as in the DNS simulations. The only information needed by the commercial software for periodic boundary conditions is the mass flow rate. The same mass flow rate $\dot{m} = 2/3$ is chosen as in the DNS simulations.

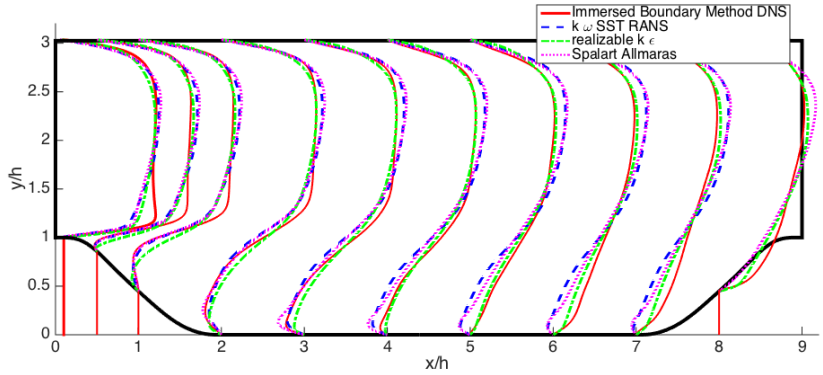


Figure 5.6: Comparison of horizontal velocity profiles for different turbulence closure models.

The mesh is constructed appropriately for RANS simulations from curvilinear blocks with 192 × 120 nodes as described in Hellsten [58] with refinement towards the lower and upper walls. A convergence study for the RANS model was carried out by Hellsten [58], who found that the effect the grid-spacing on the solution remained practically the same for 192 × 120 as it was for 384 × 256 control volumes.

The authors conducted an additional study to see how much lower the grid resolution could be chosen. Our motivation was that the benefits of a multi-fidelity study are the higher the cheaper the low-fidelity model is. It is found that the RANS solver still converges to a quite acceptable description of the flow at 64 × 32. A higher resolution was chosen anyway because the selected quantities of interests are all connected to the bubble size: a coarser grid makes it more difficult to accurately find the edges of the separation bubble.

5.2.3 RANS and DNS Compatibility

In total nine DNS simulations and 105 RANS kω SST simulations were performed at a grid size of 192×128. 81 of the RANS simulations were performed in a Clenshaw-Curtis grid and 29 in a higher level Gaussian sparse grid configuration. Figure 5.7 shows a comparison of the horizontal velocity streamlines of DNS and RANS for the geometric parameter combinations that form the corners of the simulated random domain: $h = 1, s = -0.3$,$h = 1, s = 0.3$,$h =$

$0.5, s = -0.3$,$h = 0.5, s = 0.3$. The main difference that can be observed is that the k-ω - SST RANS model overestimates the reattachment point and the length of the separation bubble. Moreover, it is visible that the hill height has a significantly larger effect on the size of the separation bubble than the hill steepness. From the large number of possible quantities of interest (QoI),

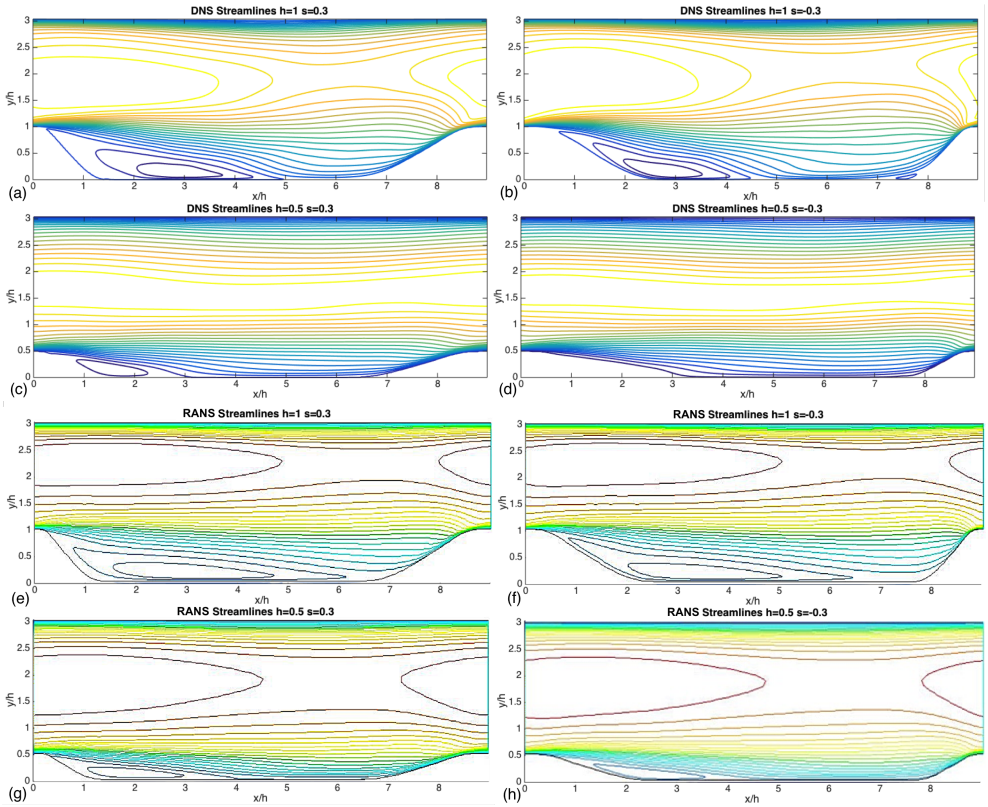


Figure 5.7: Horizontal velocity streamlines visualisation comparing DNS (a,b,c,d) and RANS simulations (e,f,g,h).

four quantities that are chosen that demonstrate and summarise the findings best. The minimum velocity inside the bubble is used as an indicator of the global accuracy of the flow; the height of the separation bubble at $x/h = 2$ and $x/h = 4$ to have one indicator in each zone of good and bad model agreement, and the reattachment location, because it is a widely known weak spot of RANS models. Since a visual comparison of all QoI is not possible for many parameters, the correlation function is used as described by Equation (5.12) and suggested by Palar [111] to determine compatibility. The results in Table 5.1 show that the correlation between the two model responses for the maximum horizontal back-flow velocity and reattachment are fairly low compared to what was used in Palar [111]. The main reason is the large conceptual dif-

QoI	$R2$	$\lvert M\rvert_{rel}$
Reattachment Location	0.8504	0.6414
Height of the separation bubble at $x/h = 2$	0.9570	0.1049
Height of the separation bubble at $x/h = 4$	0.7840	2.0988
Maximum horizontal back-flow velocity at $x/h = 2$	0.6088	0.6197

Table 5.1: Correlation test of DNS and RANS for separation bubble parameters.

ference between the two models. A minor reason can be, however, that only nine points are available for comparison. Nevertheless, the data confirm that the height of the separation bubble at $x/h = 2$ is well described by RANS, whereas the height at $x/h = 4$ is not.

For a more detailed analysis, the velocity profiles are plotted at $x/2 = 2$ and the friction coefficient along the bottom wall in Figure 5.8 using three different hill heights. Graph (a) in Figure 5.8 shows that the RANS model follows the trend of the DNS velocity profiles reasonably well for varying geometries. Figure 5.8b, however, explains the bad correlation at $x/h = 4$ and for the reattachment location. The RANS model's over-prediction of the reattachment location creates an area of bad correlation in the region after the bubble. To automatically calculate the reattachment location for the 9 DNS and 81 RANS simulations, the intersection of the skin friction coefficient with zero in the interval $x \in [2, 6.5]$ is calculated numerically, see Figure 5.8b. To ensure the correctness of the found locations, all values were checked manually.

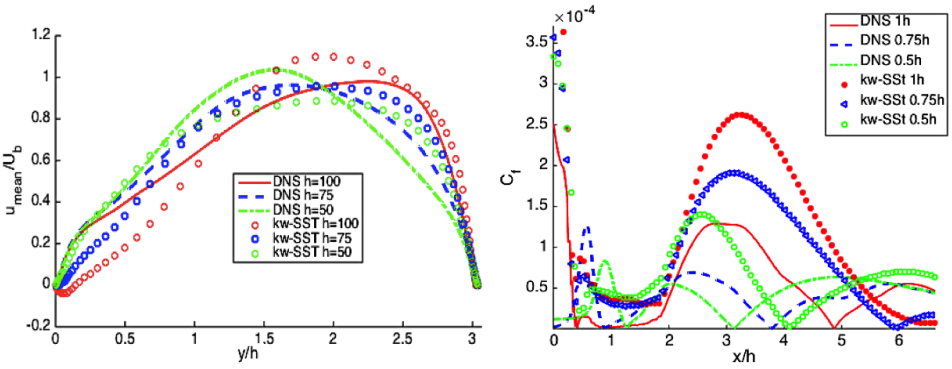


Figure 5.8: (a) Velocity profile trend validation of k-ω SST with DNS at the location $x/h = 6$. (b) Skin friction trend validation of k-ω SST with DNS along the x-axis.

To investigate how the conceptual difference between RANS and DNS manifests itself in the PDFs of an uncertainty quantification study, comparisons

between pure RANS and DNS PDFs are shown in Figure 5.9 for all investigated quantities. The DNS PDFs are created by sampling a linear response surface between the available DNS points. This creates unusually sharp edges in the distributions. It is done because the straight edges illustrate that no information about the true curvature between these DNS points is actually available. The reader should therefore interpret straight lines in the DNS PDFs as areas where the curvature of the PDF is unknown due to data scarcity. However, the general trends and regions of the DNS PDFs are quite reliable and the DNS response surfaces are physically more accurate than RANS. The RANS response surfaces, on the other hand, are obtained using a non-intrusive Polynomial Chaos method based on 81 collocation points and can thus be considered stochastically more accurate than the DNS distributions (more accurate in the sense that its shape is better converged because more data points were available).

Figures 5.9e - 5.9h show that the modelling errors in RANS can have a significant effect on the shape of the resulting probability distributions. Most strikingly, the PDF of the reattachment location is wrongly skewed to the right and favours the region around $x/h = 7$. Surprisingly, the RANS and DNS probability distributions for minimum velocity in Figure 5.9a and 5.9b look quite alike despite the calculated low correlation. This shows that high correlation and low bias between the data is not a sufficient indicator for multi-fidelity model compatibility and further research is needed. To better understand this discrepancy the RANS and DNS response surfaces are plotted for reattachment location and minimum velocity in Figure 5.10. The graph in Figure 5.10a shows a fairly good agreement between RANS and DNS minimum velocities. Only a valley of lower velocity in the DNS is wrongly left out by the RANS model. This trend partly explains the lower correlation.

The reattachment response surfaces in Figure 5.10b show that the trend of increasing bubble size for higher hill height is resolved well by both models. However, there is a sharp bend in the response surface. This is caused by the fact that the k-ω -SST model over-predicts the position of the reattachment location. As the investigated geometric uncertainty is large with respect to the domain size, the bubble growth is impeded by the next hill, see Figure 5.7. This probably creates the sharp bend in the RANS response surface, which in turn creates the large discrepancy between the two model PDFs in Figure 5.9.

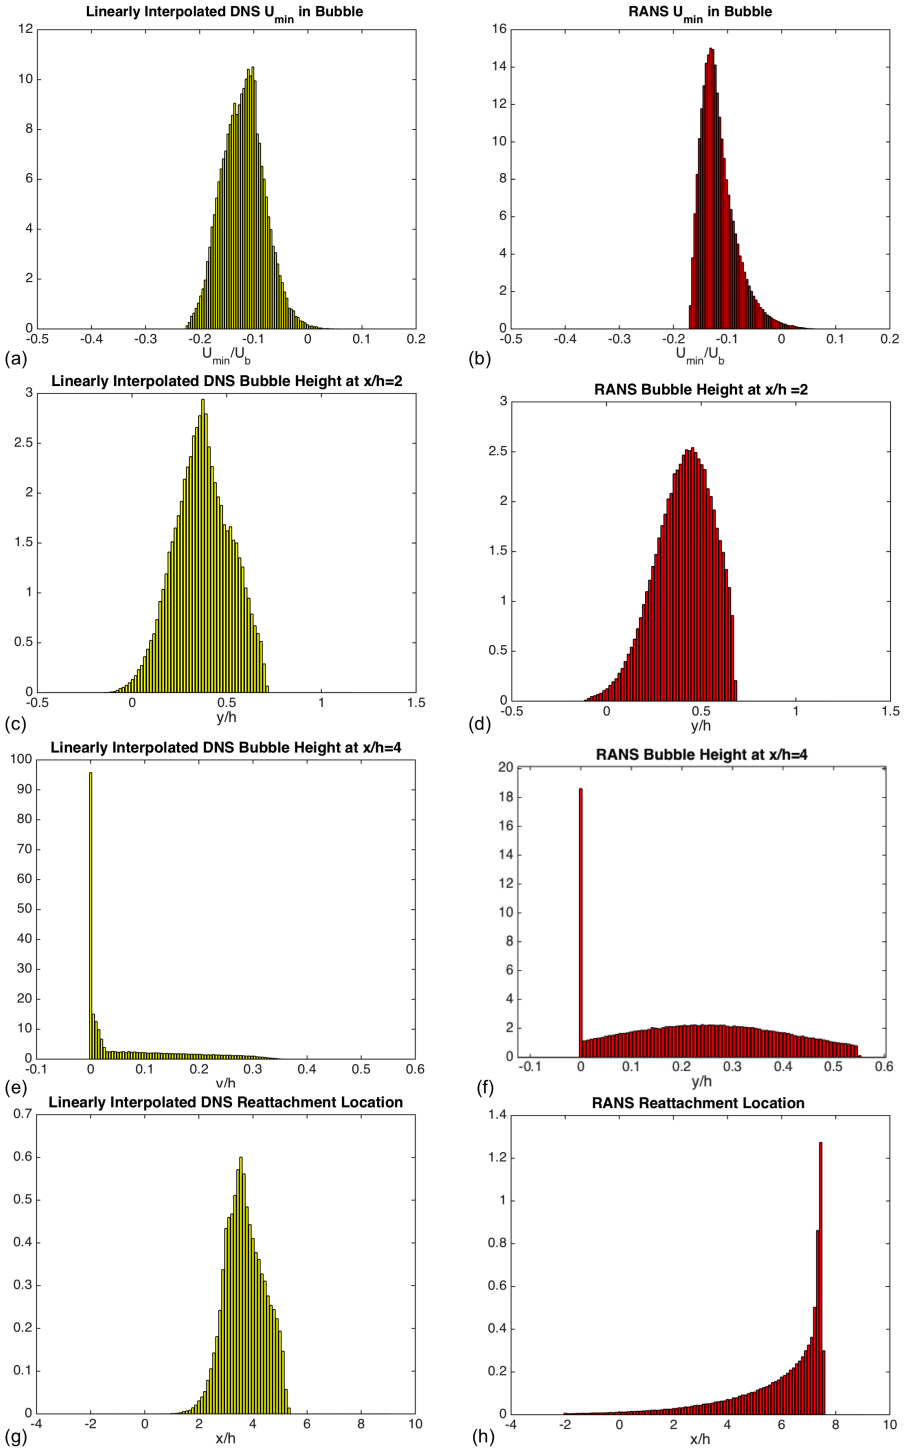


Figure 5.9: Probability distributions obtained from pure DNS and RANS data: (a,b) minimum velocity in bubble, (c,d) bubble height at $x/h = 2$, (e,f) bubble height at $x/h = 4$ and (g,h) reattachment location.

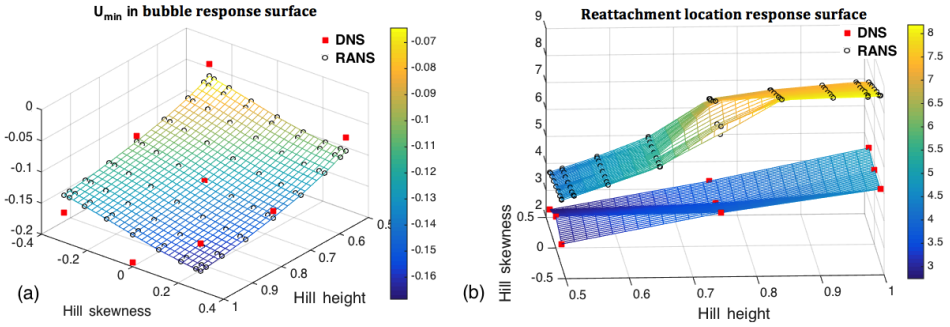


Figure 5.10: Response surfaces for the maximum back-flow velocity in bubble (left) and the reattachment point location (right).

5.2.4 Multi-fidelity Results of RANS and DNS

In this last section, the mathematical methods described in Section 5.1 are applied to create a multi-fidelity model combining RANS and DNS. Ideally, the PDFs found from this multi-fidelity model should be compared to a Monte Carlo Simulation of pure DNS results. However, this is computationally infeasible and must consequently rely on the linearly interpolated PDFs in Figure 5.9 for guidance. Due to the physicality of the underlying flow equations (no large variation in the design space are expected because the flow changes gradually) and the fact that nine simulations are available that span all corners of the domain, the ranges should be quite trustworthy. Only the shape of the distribution should change using a more refined sampling method. Probability distributions obtained using the sparse nested grid and weighted regression method are displayed in Figure 5.11. Compared to the pure RANS PDFs both multi-fidelity PDFs show a significantly higher resemblance to the pure DNS PDFs in Figure 5.9. Due to the higher number of samples added by the RANS simulations they look smoother and more converged. However, judging by the reference PDFs in Figure 5.9 the weighted regression PDFs seem like a better compromise between the range of the DNS distribution and the shape of the better converged RANS distribution. This is the case because the sparse grid Polynomial Chaos coefficients are split into corrected and uncorrected coefficients. If there is a large discrepancy between the models this results in a very inconsistent model. The weighted regression method, on the other hand, smoothes this discrepancy because it finds the Polynomial Chaos coefficients by minimising the fitting error. Moreover, it is found that the weighted regression method can be easier calibrated than the sparse grids approach. Since the weights can be chosen manually, the operator can choose a trade-off.

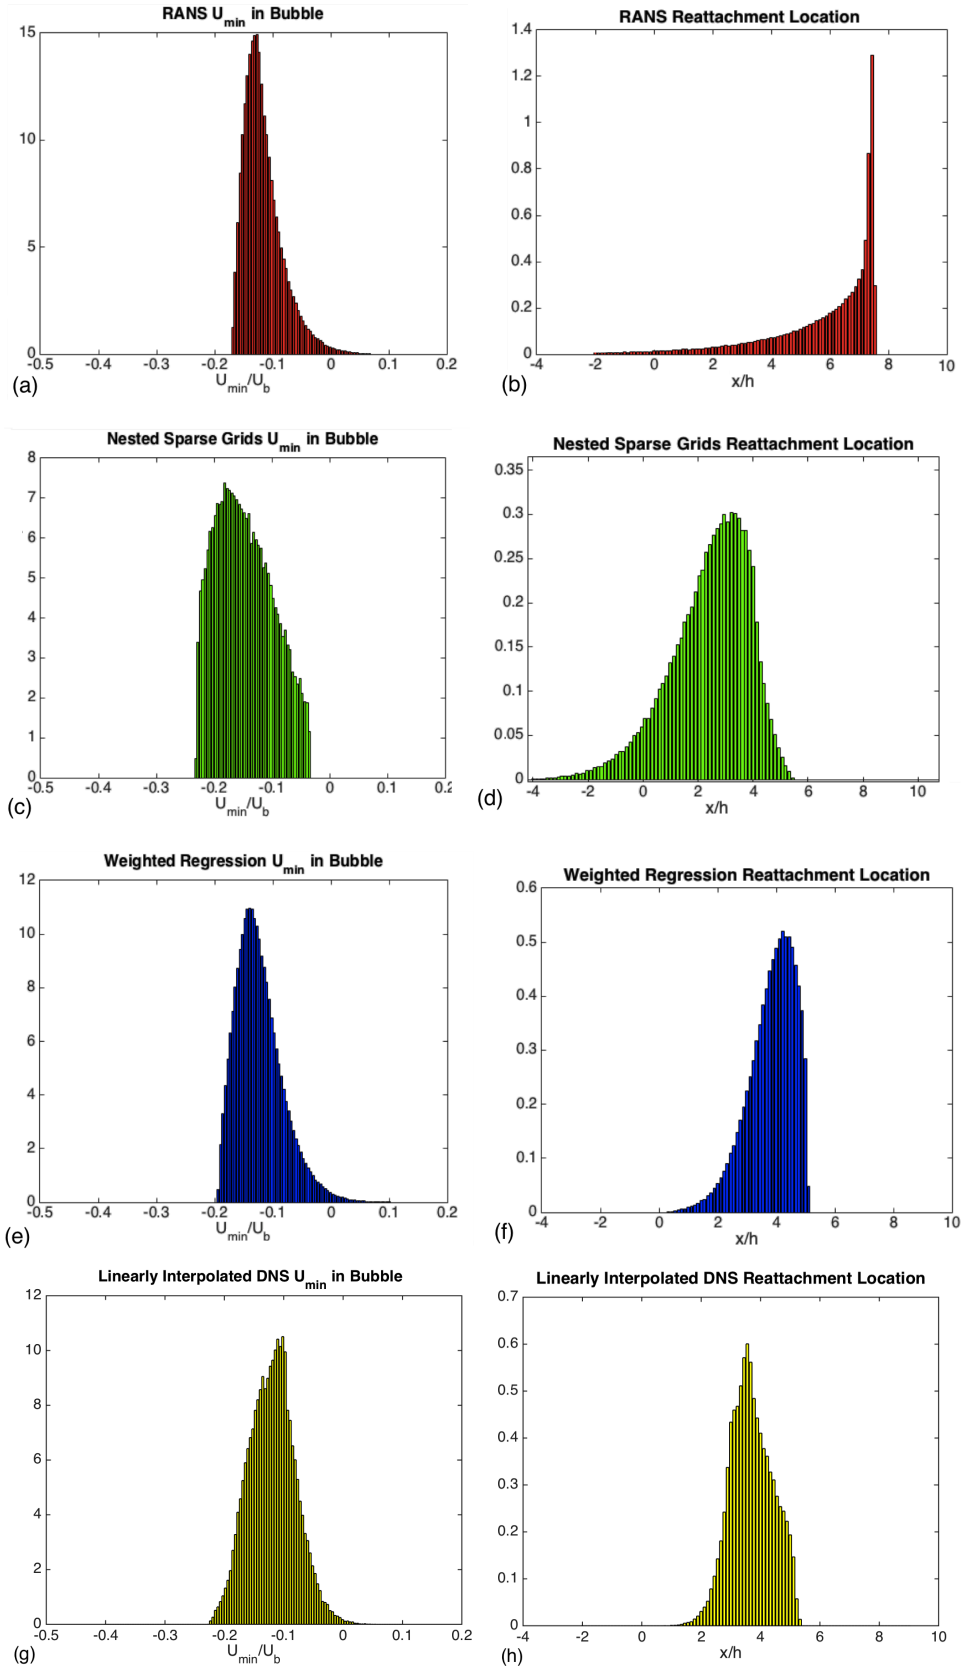


Figure 5.11: Multi-fidelity PDFs (c-f) obtained for minimum velocity inside the bubble (left) and reattachment location (right) compared with pure RANS (a,b) and pure DNS (g,h) results.

6 Engineering Applications

Data-driven Polynomial Chaos is particularly useful for industrial engineering applications. It offers a simple method that better fits to realistic test conditions. Due to the high cost of experiments in engineering, real test data are generally scarce. Moreover, data can also be highly skewed or multi-modal [107]. While non-parametric data is a statistical challenge for Askey scheme based Polynomial Chaos expansions, it makes no difference within the framework of data-driven Polynomial Chaos. Thus, the method presented in this work simplifies the access to advanced statistics for engineers with little statistical experience.

Moreover, data-driven PC offers a new answer to the question of how to deal with scarce information about input distributions. This newly discovered feature was already demonstrated with a simple diffuser example in Section 2.5. The example showed how the construction of a parametric PDF for a case with limited data always involves a lot of subjective assumptions. The data-driven Polynomial Chaos, however, offered a possibility to propagate only the given information without making assumptions.

If a Polynomial Chaos expansion is based on a wrong probability distribution, it severely diminishes the convergence behaviour [107]. As data-driven PC is based on available data and not assumed distributions, it converges always. Convergence problems are very rare. Ernst demonstrated that only a few distributions, like the infinite log-normal distribution, are exceptions. It was shown in Section 3.2 through validation with Monte Carlo Simulation that the truncated log-normal distribution can also be accurately propagated using data-driven Polynomial Chaos. The first users of arbitrary Polynomial Chaos, Oladyshkin and Nowak, showed that only moments are propagated in all Polynomial Chaos approaches. This is important because even a well-determined input PDF will at best have the same or similar first $2N$ moments as the data-driven PC expansion of order N [106]. Therefore, data-driven Polynomial Chaos offers by default the most reliable evaluation of the effect that limited input data has on model outputs among all Polynomial Chaos expansions.

6.1 Race Car Aerodynamics

Race cars are produced in low quantities and for maximum performance. Thus, their designing process is characterised by manufacturing data shortage and complex Computational Fluid Dynamics (CFD) simulations with long run times. A CFD simulation is performed for a unique set of physical model parameters, boundary conditions, and an idealised domain geometry. However, a single CFD simulation cannot take into account the stochastic variations that affect the flow system in the real world, show sensitivities in the design space or indicate parameter configurations that would result in a better performance of the design with regard to a set of engineering objectives. A design that fulfils all these conditions is called a robust design.

Obtaining a robust design for an aerodynamic system is no trivial task. Race cars are subject to many random variations that introduce uncertainty into the down-force and drag performance, such as ride height, front wing direction and pitch angle variations. To accurately predict the car performance during a race it is important to study the effect of these random variations. Such an assessment cannot be performed with the standard deterministic CFD approaches used in F1. Even with regard to stochastic approaches, no efficient method has so far been suggested within the car design literature that addresses the scarce data problem. The reason is that most efficient uncertainty quantification methods fit probability distributions to the scarce data that make large assumption errors as described above. In the following two examples, SAMBA is thus applied to two realistic car 3D CFD simulations: a NACA 0012 tip wing and the front part of an F1 car. The probabilistic variations of the lift and drag of these two configurations are calculated.

6.1.1 Wing Tip Vortex

The first test case is the simulation of the wing tip vortex around a NACA0012 airfoil. The wing and its computational domain are shown in Figure 6.1. A steady-state simulation was performed by solving RANS equations in OpenFOAM with a k-ω turbulence closure model. The mesh has 2.3Mn nodes and a grid convergence study was conducted. Moreover, a grid-independence study was carried out to make sure that the various stochastic solutions are independent of grid effects. Such a study is quite important in uncertainty quantification investigations to ensure that observed variability is not created by mesh effects. From a CFD modelling point of view, this test case is challenging for two reasons. First, there are no sharp edges. This makes it difficult

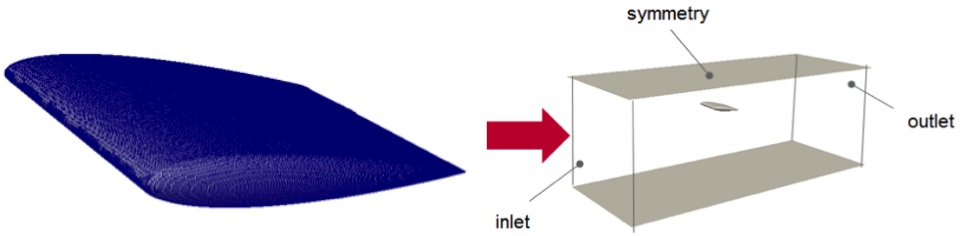


Figure 6.1: Meshed 3D NACA0012 airfoil (left) and the computational domain used for the CFD simulation (right).

to have a scheme that does not dissipate the vortex. The second reason is that vortices in the near wake at large Reynolds numbers are hard to resolve. A detailed study of numerical and experimental methods can be found in [23]. In our simulation, the Reynolds number is $Re_c = 7.410^5$ at a nominal angle of attack is 5 degrees. The Reynolds number is computed based on the chord length c as $Re_c = cU_\infty/\nu$, where U_∞ is the free stream velocity and ν is the kinematic viscosity. Due to the physical complexity, most research has focussed on conducting higher fidelity simulations to improve the physical understanding. See for example [79] for an implicit LES simulation. Due to limited computational resources, a RANS solver is used. A streamline illustration of the vortex is shown in Figure 6.2. Since the main focus in this area

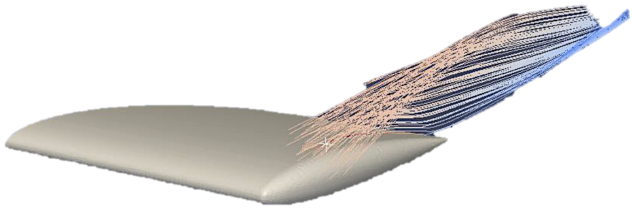


Figure 6.2: RANS streamlines illustrating the wing tip vortex.

of research has been on improving the CFD modelling, the effect of the variability in wing position has been neglected. Such a study is newly carried out in this work. Our results show the large effect of this variability. We investigate the effect of small variations in the wing orientation with a variability of about 3 degrees. Three parameters for the NACA0012 case are assumed to be affected by uncertainty: the three angles describing a solid rotation along the x, y and z axis. The random data and probability distributions for these variable parameters are shown in Figure 6.3. The effect of the input variability on lift and drag of the airfoil are summarised in Table 6.1. The results are non-dimensionalised and shown in percent compared to the reference values.

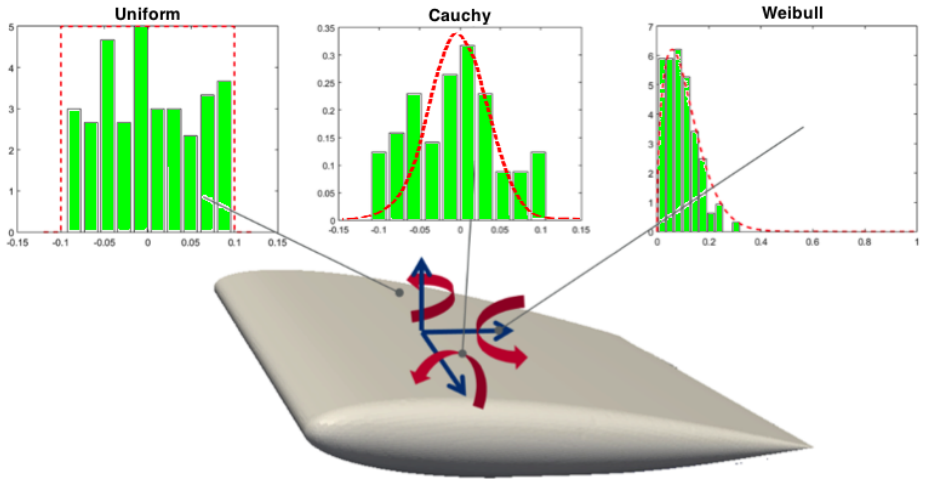


Figure 6.3: Random samples used to create a stochastic positioning of the NACA0012 airfoil.

The importance of performing uncertainty quantification studies can be seen in the fact that a difference of almost 20% from the deterministic CFD simulation is observed. The likely regions are shown by the entire output histograms

QoI	Expected Value [%]	Standard Variation [%]
Lift	81.711	17.50
Drag	84.53	23.36

Table 6.1: Mean and standard deviation of lift and drag under the influence of the imposed random variables (in percent and normalised by the deterministic reference values).

in Figure 6.4 for the lift and drag of the NACA0012 profile. Both the drag and the lift performance are skewed distributions with fat tails towards the left. This is an important discovery. It means that in rare cases the performance of the configuration can be significantly lower than anticipated. For a robust design, there should be no tails and the uncertainty should be clearly contained within a bounded interval.

6.1.2 F1 Race Car Front Wing

The next test case is the front of an F1 race car shown in Figure 6.5. As can be seen, the domain of the simulation was sufficiently large to ensure that there are no significant interactions between the car and the walls of the simulation. The dimensions for the domain are shown in Figure 6.5. Again the open source CFD package OpenFOAM is used to perform incompressible steady-

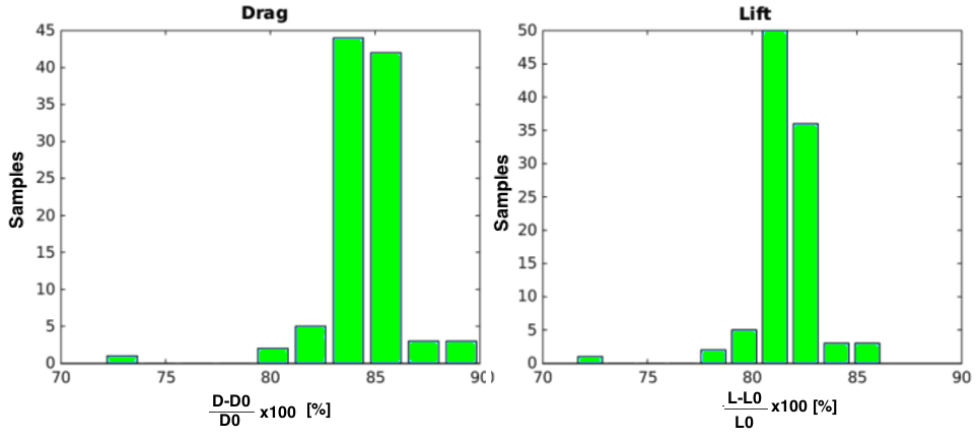


Figure 6.4: Variability of lift and drag given the uncertainty in the positioning of the NACA0012 airfoil.

state RANS simulation with k-ϵ turbulence closure model. The mesh is an unstructured grid, with prismatic layers near the walls and it was generated in SnappyMesh. A grid sensitivity study was carried out monitoring integral parameters, such as lift and drag, and the residuals of the numerical simulations. The geometry is automatically meshed and modified by a shell script. The floor was set as a moving boundary at the same velocity as the free stream inlet velocity. The impact of the distance between the inlet and outlet plane was quantified and found not to interfere with the car simulation. To ensure that the tires do not contact the floor tangentially and do not reproduce the wheel deformation with contact to the ground, a small amount of cut-off was made to ensure a larger contact area, as common practice in car CFD.

To understand the impact that small variations in the orientation of the components have on the car performance, four parameters for the F1 front wing are assumed to be affected by random variations: the height and relative distance from the ground of the wing, and similarly the height and relative distance from the ground of the car body. The rotating wheels and the moving floor are modelled in the simulation as shown in Figure 6.6. The vectors on the right-hand side in Figure 6.6 illustrate the stochastic movement of the parts. The used surrogate measurement data are shown in the histogram in Figure 6.7. The numbers shown are non-dimensional. SAMBA is used to propagate the random data and to evaluate the mean value of the car down-force under the influence of the prescribed random variations. The mean and variation of the main performance parameters (down-force, the centre of pressure position) are generally more representative of the truth than deterministic values based on subjective assumptions. Table 6.2 shows the mean parameters resulting

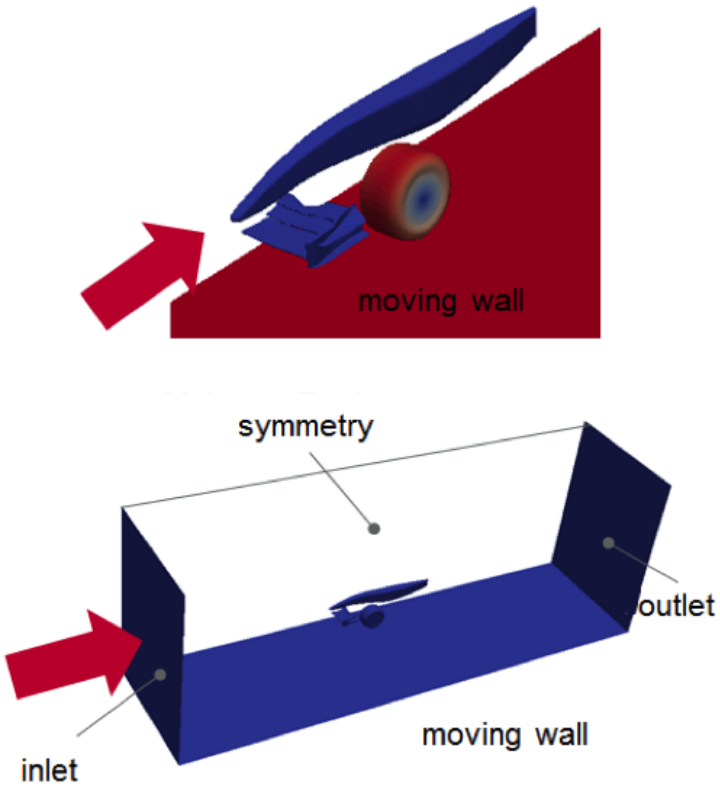


Figure 6.5: F1 front wing CAD model and CFD simulation domain.

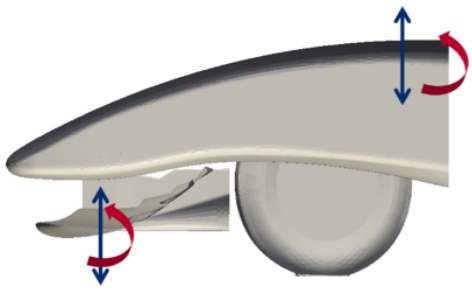


Figure 6.6: Illustration of the imposed rotational and translation uncertainty in the car body and front wing.

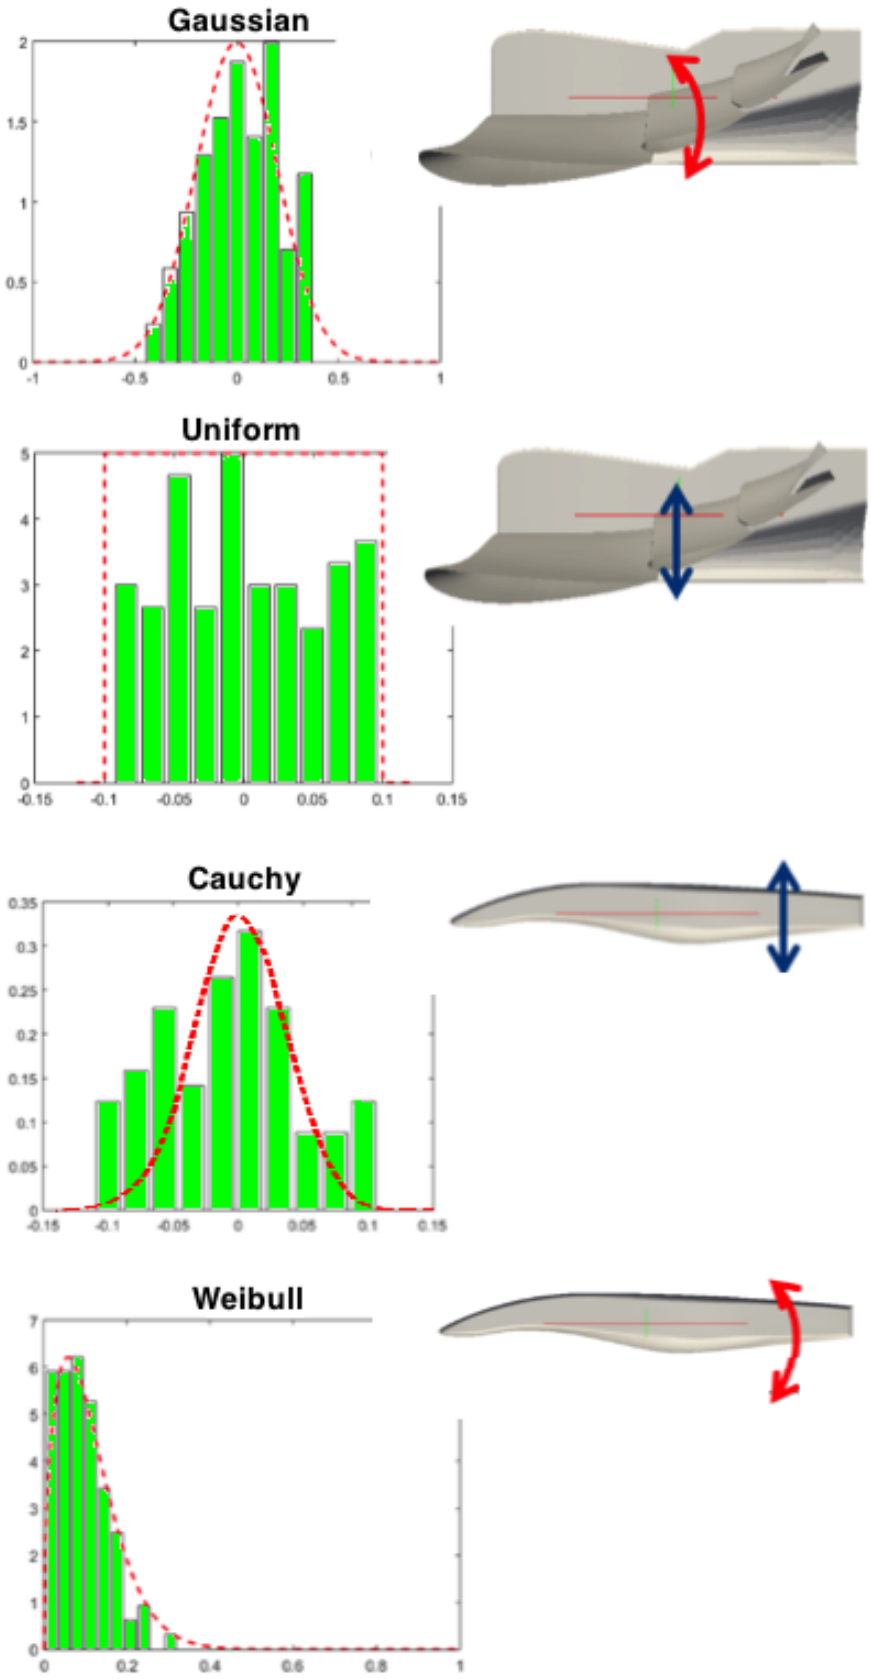


Figure 6.7: Random sample histograms and the directions of how they are assumed to affect the positioning of car model components.

from the UQ study for lift, drag and down-force. As before, the data has been normalised by the deterministic baseline configuration. The variability of lift and drag here is smaller than for the wing tip scenario. However, the lateral down-force deviates significantly from the deterministic prediction. Detailed

QoI	Expected Value [%]	Standard Variation [%]
Lift	108.12	3.74
Drag	99.15	1.57
Half Lateral Force	130.63	11.066

Table 6.2: Mean and standard deviation of the lift, drag and lateral force under the influence of the imposed random variables (in percent and normalised by the deterministic reference values).

output probability distributions are displayed in Figure 6.8. All distributions are skewed towards the left this time. For drag and lift, this can be favourable because the worst-case behaviour is better limited.

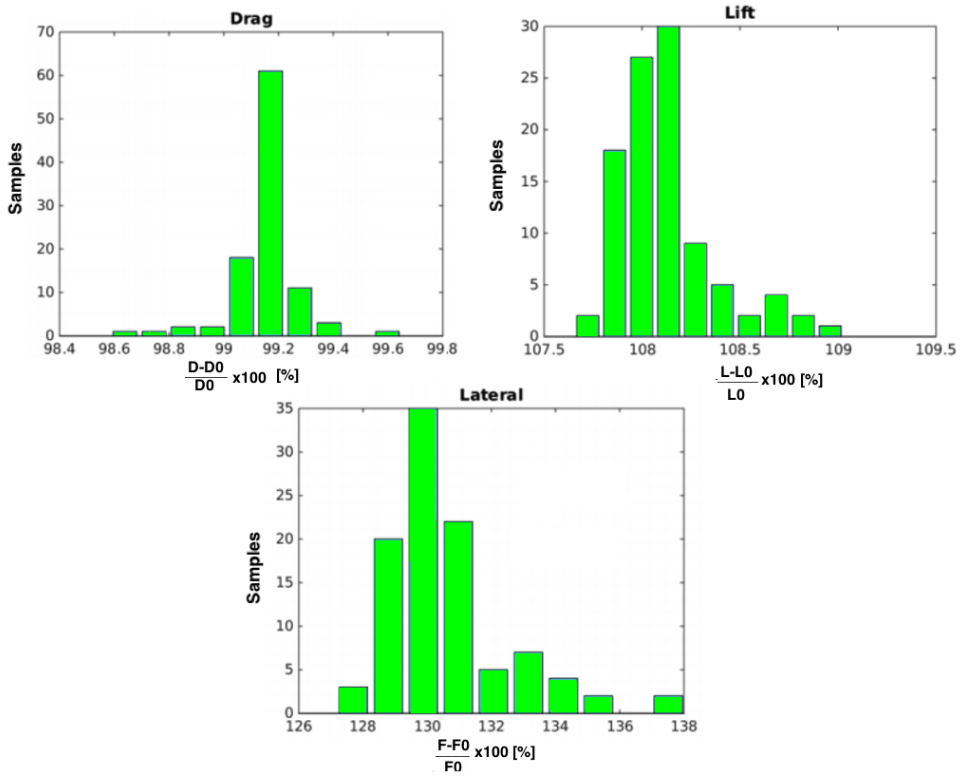


Figure 6.8: Variability of lift, drag and lateral force acting on the car for the assumed position uncertainties.

To summarise, both race car test cases demonstrate that large performance variations result from position uncertainty. It was also demonstrated how

useful the inclusion of random data using data-driven Polynomial Chaos is to analyse the behaviour of a car under variable conditions. As a next step, these probability distributions can be used to obtain more robust car designs in order to improve the probability to win.

6.2 Turbomachinery

There is a growing need for reliability in gas turbines used for power generation. The reduction in the size of the core engine has increased the impact of manufacturing deviations or assembly errors. Therefore, Uncertainty Quantification (UQ) techniques are becoming an increasingly valuable tool to improve the performance of turbomachinery. They can take the stochastic variations due to manufacturing errors and in-service degradation into account [96]. Stochastic variations are usually small, but they can still have a big impact on the engine life. For example, the variation of the film cooling diameter due to manufacturing imperfection can result in an engine life prediction that is off by 30 % [13]. For this reason, UQ is going to play a major role in increasing the reliability of gas turbine designs in the future.

In recent years, Non-Intrusive Polynomial Chaos has become the most common UQ technique in turbomachinery design [92]. The growing popularity of non-intrusive methods in the field is partly due to the fact that they treat industrial CFD solvers as a black box. However, there are still three problems:

1. Cases in turbomachinery are characterised by a variety of different input types such as histograms (from experimental data) or normal PDFs (design rules). Thus, the application of UQ requires the adaptation of ad hoc methods for each individual case.

2. With current methods, parametric PDFs have to be determined for all inputs. This is difficult if only a few samples are available. In gas turbines, however, the collection of statistical information is difficult, expensive and having scarce information is the norm.

3. For Non-Intrusive Polynomial Chaos methods, the number of required simulations grows exponentially with increasing numbers of input uncertainties - the so-called 'curse of dimensionality'.

The most frequently used distribution in turbomachinery is the Gaussian distribution. It is used to represent manufacturing or performance variability [10, 92, 15, 131, 76]. The assumption of normality can be justified by referring to the Central Limit Theorem and correct for many cases. For example,

previous studies on compressor blades have confirmed that geometric manufacturing errors of surface deviations are approximately normal distributed [47, 38]. Normality is not always the case though. Panizza showed recently for manufacturing process variability that all of the encountered distributions were distinctly non-Gaussian [112]. In particular, for small sample sizes, the assumption of a Gaussian distribution can make a big difference from the actual distribution of the data. This is a major problem in UQ for turbomachinery since for most cases there is only little experimental data available. The reason is the high cost associated with performing measurements and the fact that some quantities, like the stochastic variation of the inflow velocity or pressure [113], are difficult to measure.

SAMBA is introduced in this section as a single method to deal with the above problems for UQ in turbomachinery. It can deal with real-sized industrial problems in a non-intrusive and highly efficient way, it can combine various types of data and it is the only method propagating any available random information without making wrong probabilistic assumptions. Most importantly, it is more accurate than other methods for scarce data as shown in Chapter 2. In the following, SAMBA's capability to deal with an uncertainty quantification problem in gas turbine problems is demonstrated by considering geometric manufacturing uncertainties affecting the blade surface in a Low-Pressure Turbine (LPT) and tip gap leakages in a multistage turbine.

6.2.1 Manufacturing Errors in Low-Pressure Gas Turbine Blades

Previous studies have shown that gas turbine compressors are especially sensitive to these uncertainties [52, 77]. In some cases, it may be possible that almost no prior knowledge about the probability distribution is available.

The following simulation is a combination of experimental data from two works [141, 94]. For the given case all input uncertainties are assumed to be uncorrelated. The absolute values of the manufacturing errors shown in Figure 6.9 are not based on measurements taken on the H2 profile. They are, however, based on actual manufacturing data of an IPT. Thus, the shape of the points, their location and the relative importance correspond to real measurement values: the data are not perfectly Gaussian-like pseudo-random numbers and are skewed. This highlights the importance of a moment based formulation as SAMBA for turbomachinery problems. The suction side standard deviation has a value that is 20% lower than the values on the pressure side. This is consistent also with in-service degradation effects, taking into account that

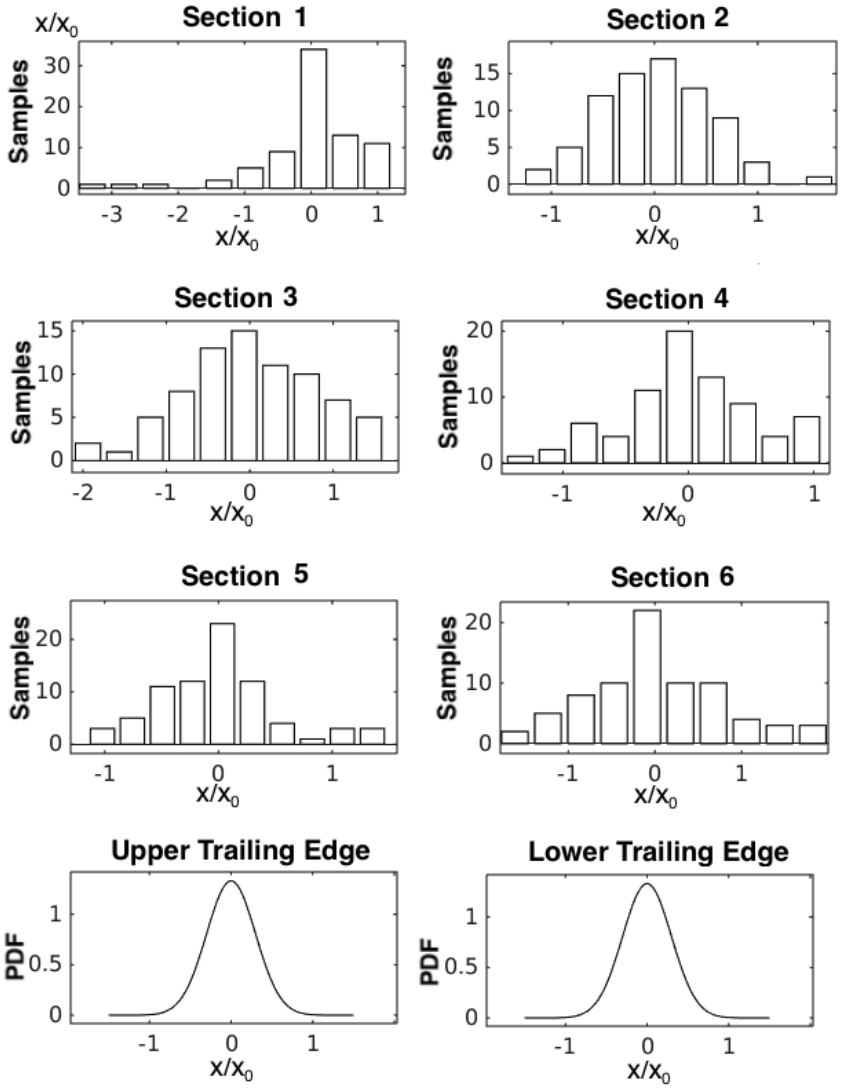


Figure 6.9: Normalised measurements of geometric manufacturing deviations for various blade regions.

the pressure side is more exposed towards local depositions.

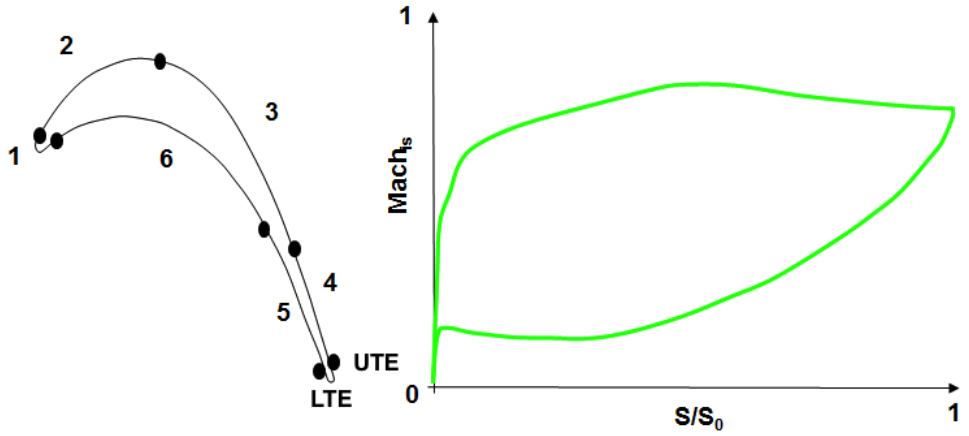


Figure 6.10: Schematic of the H2 airfoil, not in scale, with 6 sections and 2 points varied according to 6 different measurement sets and two PDFs given in Figure 6.9 and isentropic Mach number.

The geometry was deformed smoothly in each section upholding continuity between the sections in such a way that the point in the middle of each section exactly follows the assumed input histograms. This is, of course, a major simplification regarding the actual geometric deformation. However, it is sufficient to demonstrate the application of SAMBA. The suction side was divided into 3 regions, from leading edge to maximum peak (2), from maximum peak to halfway (3), and the rear part of the suction side (4). The pressure side, that is going to have a minor impact on losses, in two regions, high curvature side (6) and the rear part (5). The leading edge was treated separately (1) and the trailing edge is not measured. For this reason for the trailing edge, two arbitrary PDFs have been chosen, showing also the ability of SAMBA to mix continuous PDFs and histograms. These regions are shown in Figure 6.10. The chosen airfoil H2 was derived from the work published in [94]. The geometrical data are summarised in Table 6.3 and in this case, the lower pitch value has been used. To estimate the profile pressure losses p_L due to the given stochastic manufacturing uncertainty, a 2D in-house solver with a k-ω Langtry-Menter transition model is used. Using SAMBA for uncertainty quantification the model was run 17 times using 8 optimal input values for the 6 given measurement histograms and the 2 distributions.

The pressure losses p_L in Figure 6.11 are shown with respect to the normalized KSI pressure losses shown by Montomoli et al. [94]. The strong influence of the uncertainty is seen in the fact that the mean value of the losses is 20% higher than in the reference solution.[94].

Parameter	Unit	Value
Chord	[mm]	27
Design inlet flow angle	[deg]	36.7
Design exit flow angle	[deg]	-59.04
Aspect ratio	[-]	3.78
Trailing edge thickness	[mm]	0.41
Pitch/chord	[-]	0.784, 0.896
Number of blades	[-]	9,8
Surface roughness	[-]	0.185, 0.037
Axial distance bar to cascade	[mm]	20

Table 6.3: Cascade geometrical parameters as defined in Montomoli [94]

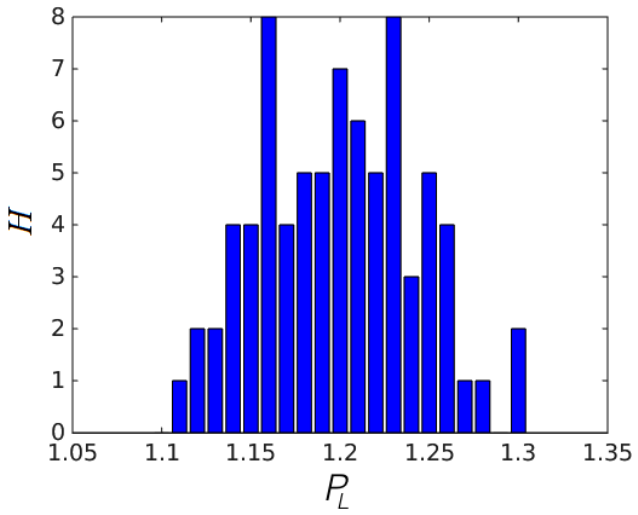


Figure 6.11: Histogram of pressure losses p_L of H2 profile for the 6 manufacturing uncertainty histograms and 2 assumed PDFs.

In order to estimate the relative importance of each zone, to improve the manufacturing process of a specific area, a sensitivity analysis using Sobol Indices was performed. The resulting indices are summarised in Table 6.4. They show that the geometric variations on the suction side have a larger effect on the pressures losses than on the pressure side of the airfoil despite the fact that their input variation was smaller. In particular, section 3, contributes to almost 85% of the occurring total variance. This is consistent with what was found experimentally and it is driven by the adverse pressure gradient on the suction side at very low Reynolds number. Overall the variation in the leading

Section	Sobol index
1	0.0205
2	0.330204
3	0.8587
4	0.0365
5	0.0020
6	0.0615
UTE	0.0003
LTE	0.0001

Table 6.4: Sobol sensitivity indices showing the effect that manufacturing uncertainty has on the pressure losses dependent on region. It can be seen that the middle to end section of the suction side is most important.

edge shape, as well as the beginning of the suction side and the upper side of the trailing edge have a smaller importance and their relative impact is more or less the same. The impact of uncertainty of the trailing edge shape is minimal and can be neglected. The results are consistent with the open literature and the overall goal of this application was to prove the flexibility of SAMBA with several different input distributions for turbomachinery applications.

6.2.2 Tip Leakage in a Multistage Gas Turbine

Most uncertainty quantification studies for turbomachinery design are performed for highly simplified geometries without experimental data for comparison. The following UQ study is thus rather exceptional. It is supported by complex three-dimensional CFD simulations performed by Mazzoni [84] and experimental results performed by Rosic [121]. This example shows how UQ can be used as a tool to identify possible errors in a CFD simulation by

comparing UQ results with experiments. The experiments were performed on

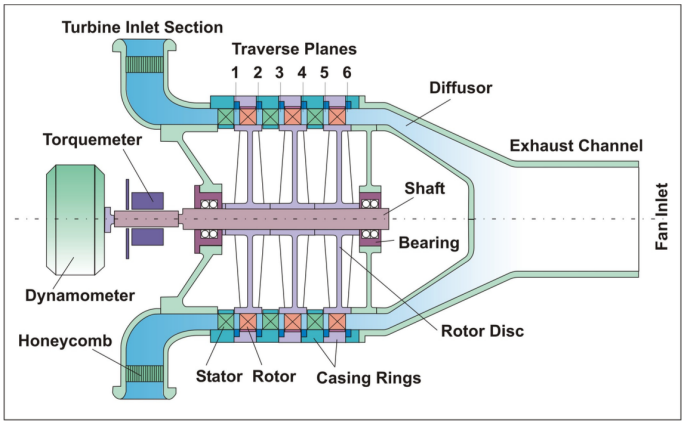


Figure 6.12: Schematic illustration of the experimental turbine, from Mazzoni [84].

a low-speed multistage air turbine that was designed to represent the first few stages of a high-pressure steam turbine shown in Figure 6.12. Atmospheric air enters the inlet section radially and passes through a flow straightener before entering the turbine. Air is drawn through three turbine stages by a fan downstream of the turbine. The blading was designed to represent typical high-pressure steam turbine conditions, and all parameters satisfy this condition. Only the Reynolds number, based on exit velocity and true chord, is significantly lower than that in a real turbine.

Both stator and rotor blades are shrouded. The rotor shroud cavities geometry and sealing arrangement investigated here is described in Figure 6.13. The configuration is representative of a commercial turbine design as it allows for the axial movement of the rotor shaft relative to the casing. The shroud forms a radial clearance of 3mm with the casing. Using two approximately 0.75mm clearance radial seals results in an over-shroud leakage flow of approximately 2%. The mechanisms of shroud leakage losses are highly dependent on the particular shroud and stage geometry. Therefore, it is difficult to develop universal design rules and to correlate different geometrical and operational turbine parameters to generic shroud design. A map of the possible turbine efficiency changes caused by different shroud modifications was obtained by Rosic [122, 123]. Another difficulty is that the blade tip clearance changes during operation due to rotor unbalances, dynamic deformations and blade vibrations. This can have a huge impact on the safety, lifetime and the energy efficiency of turbomachines and it is impossible to consider this without UQ methods.

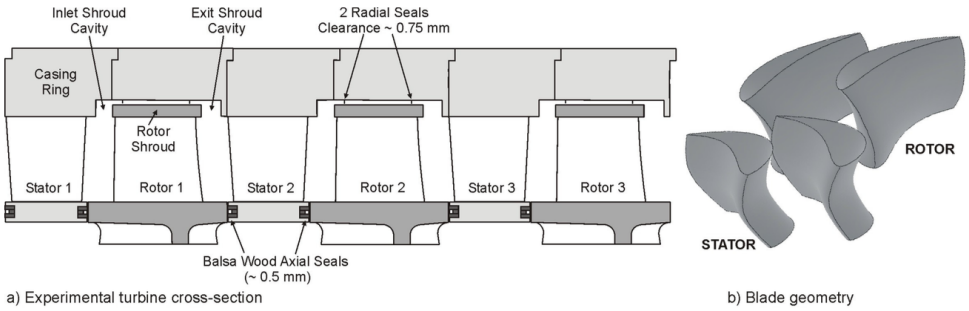


Figure 6.13: Representation of the experimental turbine showing the arrangement for shroud sealing, from Mazzoni [84].

For this reason, the geometric uncertainty in the radial seals clearance of a datum shroud geometry is stochastically varied in the following by adding SAMBA to the CFD simulations. It is well known that the rotor gap can change the gas turbine efficiency, but the impact of the random variation of the clearance height has not been investigated before.

The CFD simulation was set-up by Mazzoni as in [84] and is only summarised in the following. It was performed using TBLOCK, a multi-block structured grid solver developed by John Denton [31]. A steady or unsteady

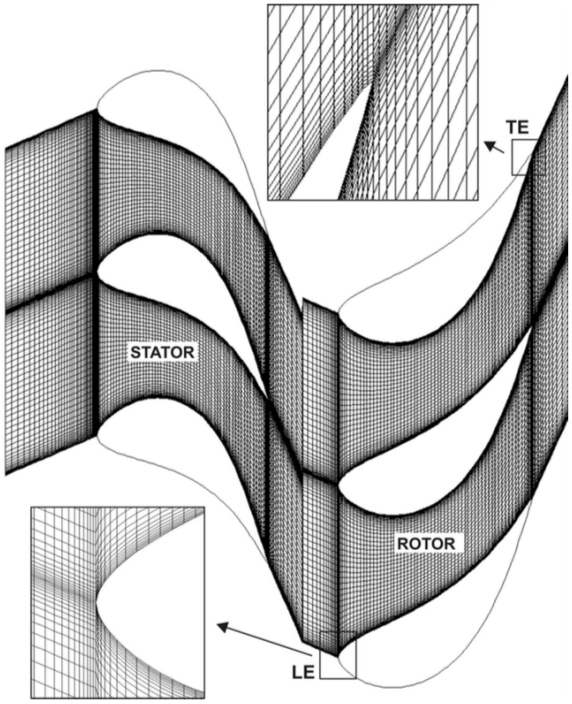


Figure 6.14: TBLOCK computational domain, from Mazzoni [84].

RANS solver using the finite volume method was employed, with the explicit SCREE scheme [30] for steady simulations and a dual-time stepping method for unsteady calculations. Turbulence is modelled using a mixing length approach. The laminar to turbulent boundary layer transition was modelled by assuming a laminar boundary layer up to a specified point on each blade and end-wall surface and a turbulent one downstream. The mixing plane approach allows reversed flow across the mixing plane, which may occur near the end-walls close to the leakage cavities. To accelerate convergence, multigrid and spatially varied time steps are used. The domain was divided into 72 blocks. In particular, the leakage flow paths were divided into the necessary number of rectangular blocks using structured H meshes. Two different levels of mesh refinement were analysed. For the finer mesh, the main blade passage was modelled using 124 grid points in the axial, and 114 grid points in both pitch-wise and radial directions as shown in Figure 6.14.

In Figure 6.15, CFD results of the total pressure coefficient downstream of the second and third stators and a comparison with experimental results are shown. Both pitch-wise averaged profiles (left) and 2-D contour map (right)

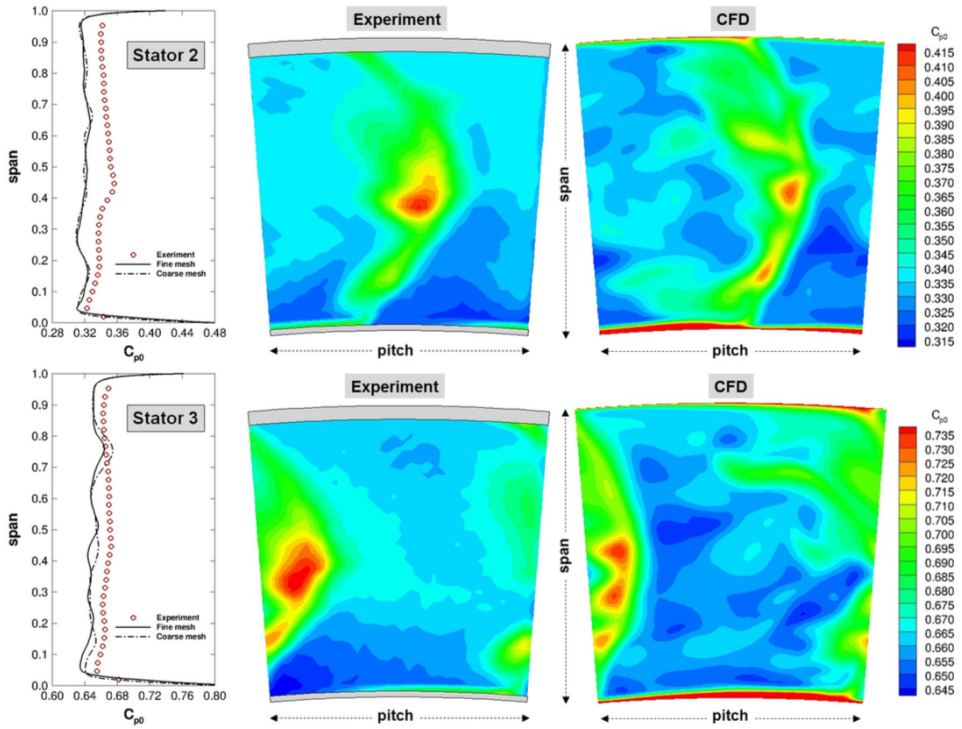


Figure 6.15: Measured and predicted total pressure coefficient downstream of stator 2 & 3, from Mazzoni [84].

are presented. The measurement plane covers one pitch from 3% to 95% span.

As it can be observed on the pitch-wise averaged profiles, numerical simulations on both coarse and fine mesh reproduce the pressure levels and the main trend along the span well. A comparison between 2-D contours shows that TBLOCK calculations capture the main flow features reasonably well, such as the locations and intensity of the main loss core, the hub end-wall secondary flow, and the migration towards the mid-span of the concentrated low momentum flow associated with casing end-wall boundary layers and leakage.

As no stochastic measurement data is available for the clearance gap, it is assumed to be a Gaussian random variable with distribution $N(0.75, 0.03356)$ as shown in Figure 6.16. This assumption is made based on the available literature on manufacturing uncertainties, which tend to display Gaussian behaviour [112]. The results of the UQ study are shown in Figure 6.17 for the

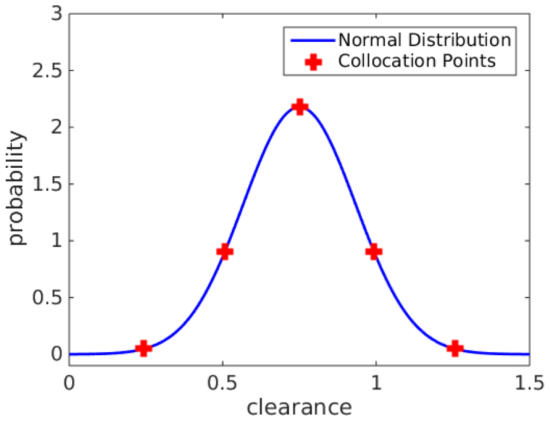


Figure 6.16: Probability distribution used as input PDF for the clearance gap including optimal Gaussian collocation points.

stagnation pressure uncertainty and in Figure for the yaw angle uncertainty 6.18. It can be observed in the first stage that the impact of uncertainty in the rotor tip is affecting only the higher part of the rotor. Only about 10% of the span (100-90%) is influenced by the rotor tip random variation. A similarly large random variation in this region can be observed for rotor 2 and rotor 3. For the second stator, the peak of uncertainty is surprisingly at 50% and the region between 40% and 60% of the span is most affected. For the second rotor, the entire span is affected. In the last stage, the standard deviation for the stator peaks at 78%, with a region of increased variation higher between 65% and 80%. The last rotor is most strongly affected by the uncertainty with a strong variation from 20 to 80%.

To make sense of these results, it has to be understood that there are different mechanisms involved: a change in rotor tip gaps modifies the mass flow

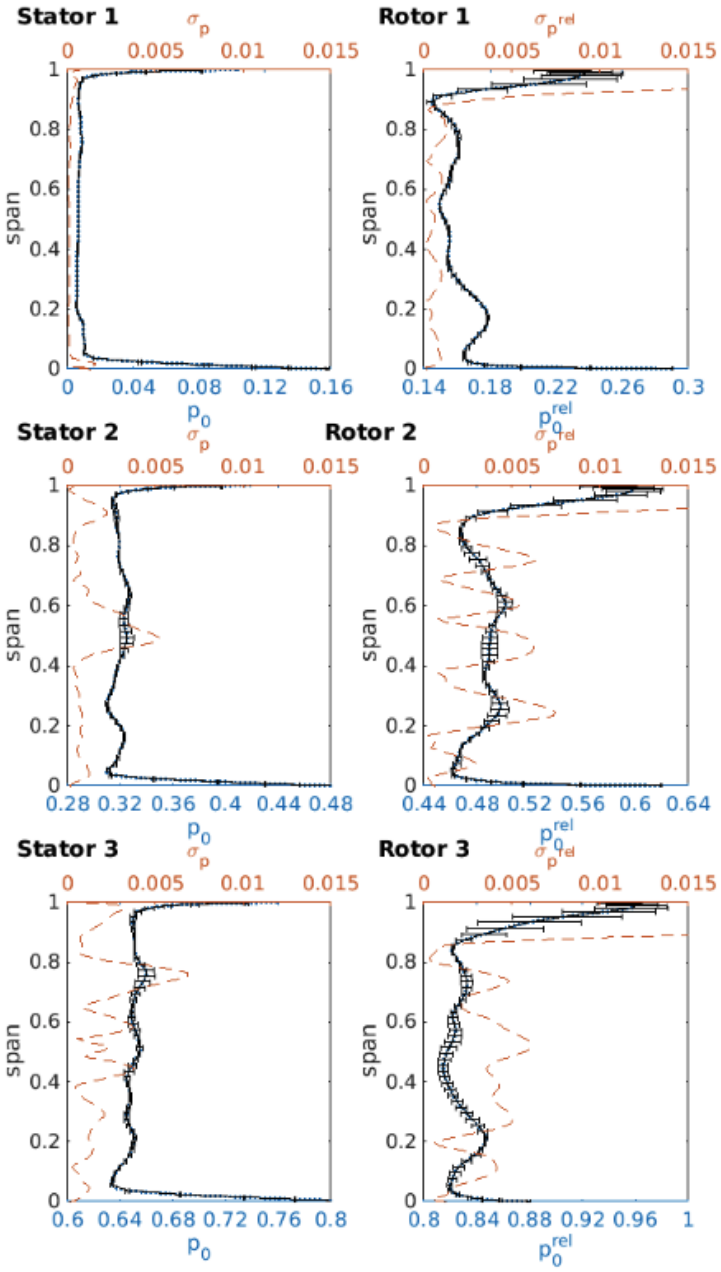


Figure 6.17: The graphs show the effect of the assumed clearance gap uncertainty on the stagnation pressure for the three rotor and stator stages.

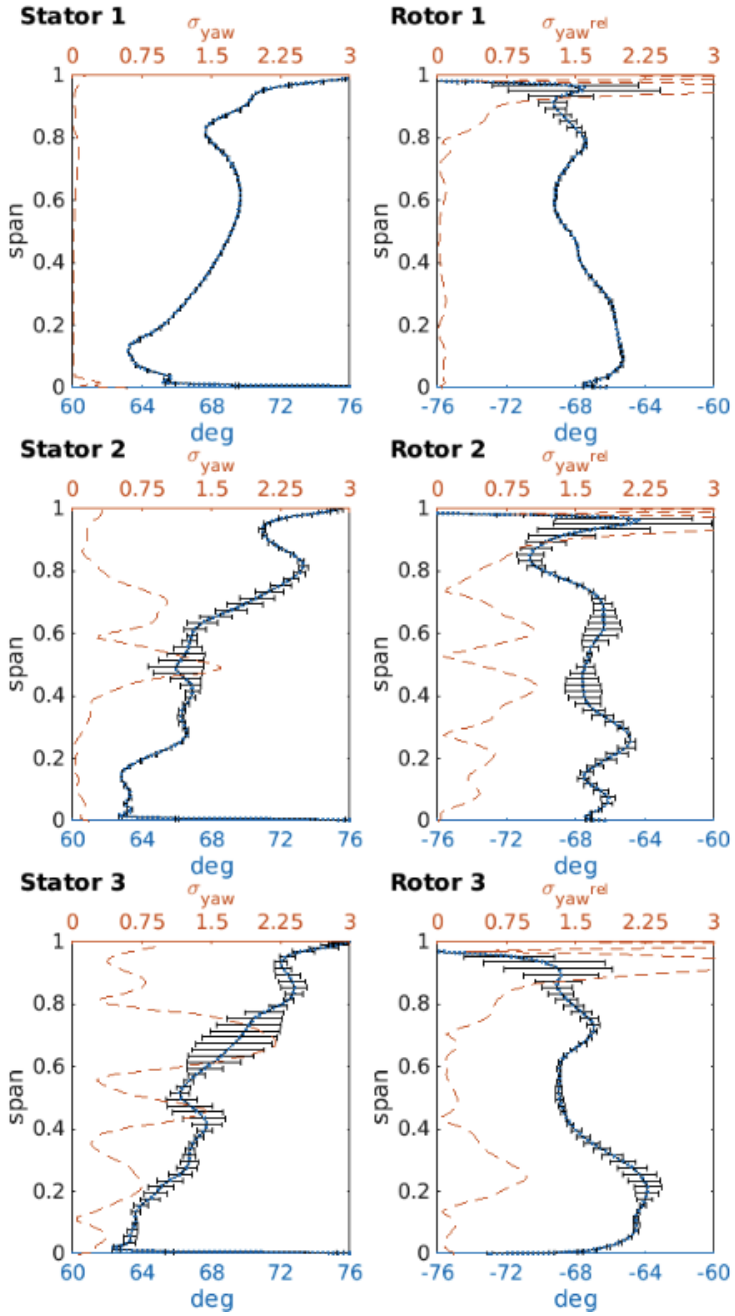


Figure 6.18: The graphs show the effect of the assumed clearance gap uncertainty on the yaw angle for the three rotor and stator stages.

distribution and this alters the radial uncertainty. At the same time, the impact of the leakage flow modifies the work carried out by the rotor due to the variation of the inlet and exit angle. This is clearly shown in Figure 6.18 where the exit angles indicate the variation in load associated with each row. The found regions with high standard deviation agree with the regions in which the CFD results differ largely from the experimental in the original work of Rosic et al. [121].

In summary, the performed uncertainty quantification study suggests a reason for the disagreement between CFD and experiments: it can be associated with a small variation of the tip gap in the machine that is different from the CFD model.

6.3 Structural Dynamics

The uncertainties affecting structural mechanics simulations are to some degree comparable to those in CFD simulations. However, the modelling techniques are quite different. The prevailing method is Finite Element Simulation (FEM). When modelling a structure with FEM the intricacies lie in finding the right number of degrees of freedom (DoF) in the model compared to reality, using appropriate shape functions, modelling its materials, damping, joint properties and boundary conditions. Examples for varying conditions are the structure's relation to the ground or in-flight thrust load fluctuations. The biggest uncertainty usually is the loading, which can vary quite stochastically during a structure's lifetime. In addition, there are also fabrication-induced errors that are difficult to consider and measurement data tends to be noisy. Data is usually not as scarce for structures as it is for Formula 1 simulations or as costly as it is for turbomachinery, but the uncertainties tend to be of greater magnitude. This has led to the consistent use of safety factors in structural engineering. However, safety factors lack scientific rigour, produce sub-optimal designs and make it difficult to predict the reliability of a structure. Uncertainty quantification methods provide a more objective and quantitative approach to uphold safety while they help to reduce costs and increase efficiency. They make it possible to search for suitable answers in a wider design space. In addition, they simultaneously provide expressions of confidence in a given estimation.

6.3.1 NASA's Space Launch System

The UQ methods presented here for structural safety analysis were designed for the new NASA Space Launch System (SLS). The FEM model or results of the actual SLS are confidential and cannot be presented in this work. Thanks to the courtesy of Lucas Horta, however, released results [59] from the Ares I-X in Figure 6.19 can be shown. The Ares I-X is a reasonably close relative of the new SLS. Like the SLS, the Ares I-X is a crew launch vehicle. However, the Ares I-X consisted of a functional four-segment Solid Rocket Booster (SRB) stage and the fifth segment was only a mass simulator. The SLS uses two five-segment solid rocket boosters and a centre booster segment. It is designed to carry humans farther into space than ever before, on missions to Mars and even beyond to possibly redirect Asteroids. Figure 6.20 shows a detailed modal analysis carried out for the entire FEM model of the Ares I-X. According to NASA, the new SLS will significantly increase crew safety. To that end,

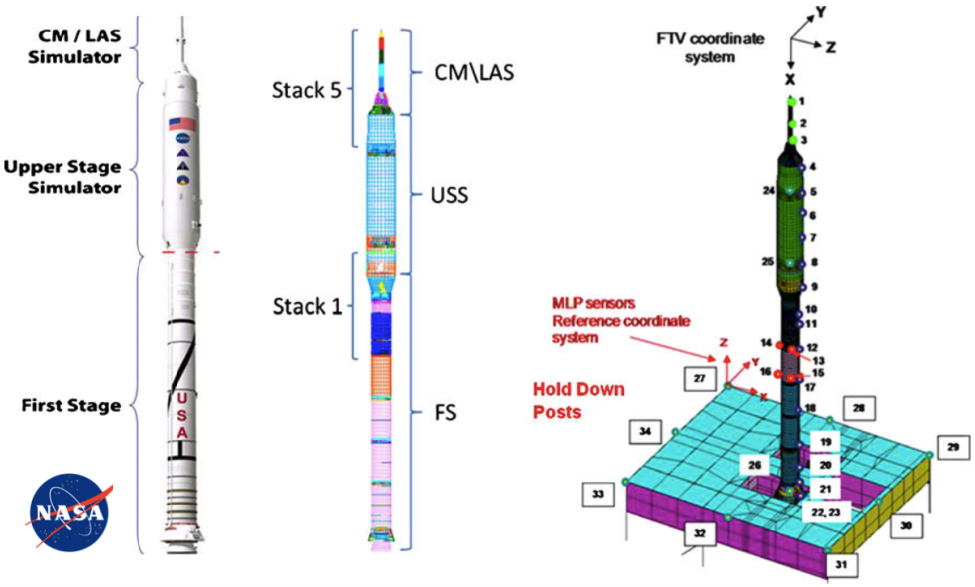


Figure 6.19: Ares I-X vehicle schematic, Finite Element Model and measurement instrumentation positions, taken from Horta [59].

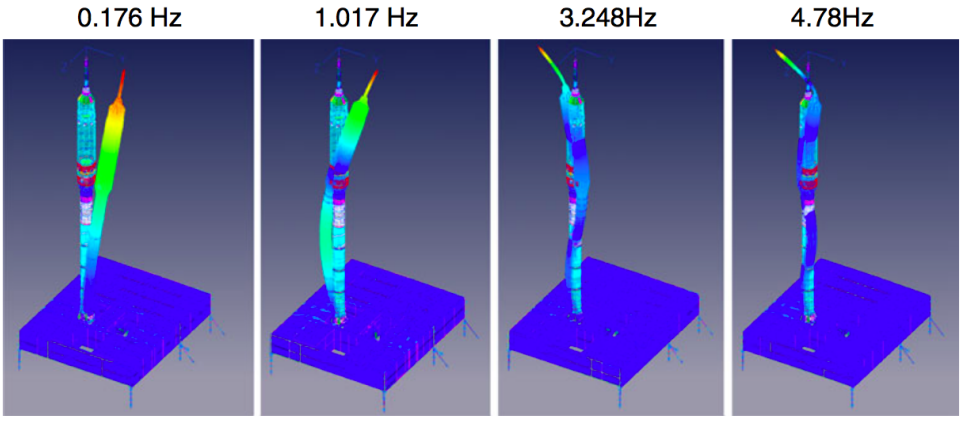


Figure 6.20: Modal analysis of the FEM model of the Ares I-X, taken from Horta [59].

uncertainty quantification already plays an important role during the designing and testing process. In this section, SAMBA will be used to perform a global sensitivity analysis to localise erroneous model components.

FEM Model Updating

Not only the stability but also the reliability of a structure can be assessed from a FEM model. Discrepancies between simulation and real structure (represented by experimental results that are also biased by measurement errors) can hence adversely affect the design quality, maintenance, and repair schedule of a structure [89]. NASA, therefore, performs a process called FEM model calibration or FEM model updating before the maiden flight of each new space vehicle, where the measurement results of the right graph in Figure 6.19 are used to improve the FEM model of Figure 6.20. Model calibration using experimental data is a useful tool to improve a FEM model's capability to accurately predict the integrity of the actual structure [89]. It is the process of updating design parameters in the FEM model using an optimisation approach in order to reconcile FEM model predictions with experimental test data. A schematic of this process is illustrated in Figure 6.21.

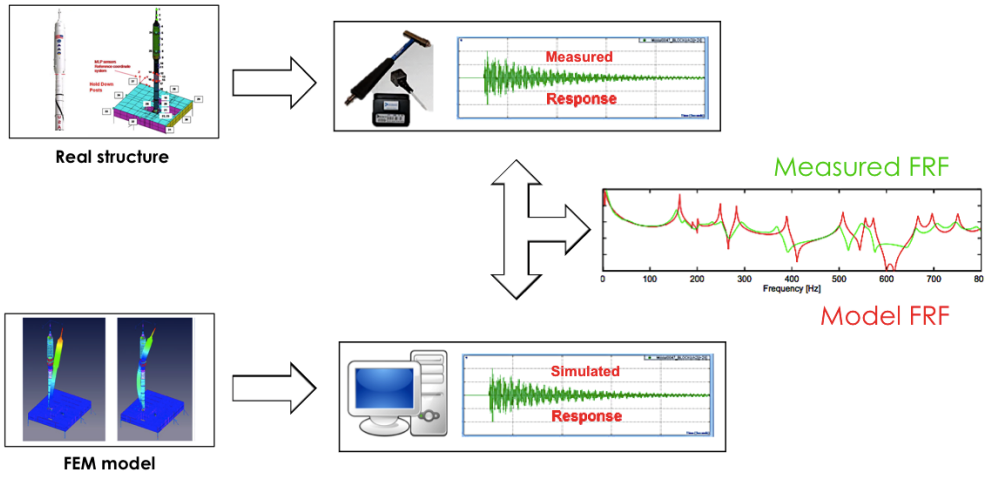


Figure 6.21: Illustration of the concept of FEM model updating: the modal response of the model is updated through parameter changes until it agrees with the experimental modal response.

The agreement between test and FEM analysis is quantified using a cost function that usually compares frequency response function (FRF) data measured from sensors placed in strategic locations on the rocket structure and a response prediction based on the FEM model. A schematic of the FEM

model (middle) and the locations of the measurement instruments (right) for the modal test are shown in Figure 6.19.

Even though FEM model updating has been researched extensively for decades, no technique has been universally accepted yet. This is partly the case because model updating is a mathematically very challenging problem. In order to update a FEM model from experimental data, a high-dimensional non-linear optimisation problem has to be solved that is characterised by many local minima and infinitely many solutions with parameters that are often strongly correlated. The solution of such an optimisation problem is computationally very expensive. In particular, if the runtime of the FEM model is long and no good initial solution is available. Providing a satisfactory initial solution, however, becomes increasingly more difficult the larger the FEM model is due to model-form uncertainty (modelling welds, joints, materials, etc.).

These problem characteristics, especially the periodicity in the solution space, can make it questionable if updated parameters are actually physical. In order to maintain physicality, it is consequently crucial to introduce tight, physically feasible bounds for all variables during the optimisation process.

Here, the Predictor-Corrector Method developed by Henning Grafe and implemented by Sean Kenny (NASA) is used [54] for model updating. It is an iterative frequency response technique that employs gradient information for convergence acceleration and two unique frequency based correlation coefficients as the cost function. However, as SAMBA is a non-intrusive uncertainty quantification technique any other updating method could be used as well.

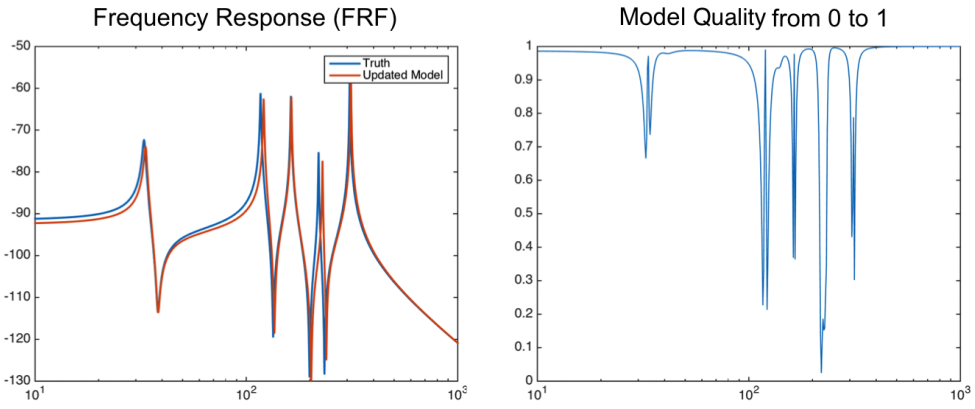


Figure 6.22: The left graph shows the model and experimental frequency response. The right graph shows the difference between them, calculated with the correlation coefficient used in the predictor corrector method.

The main issue with all updating methods is that in order to efficiently minimise the error between simulation results and experimental data, the erroneous model parameters should be identified before starting the calibration process. To find these parameters, a sensitivity study usually precedes the calibration to reduce the dimensionality. The computational cost decreases quite rapidly when reducing the numbers of parameters. Moreover, an improved agreement is less likely to be caused by falsifying correct parameters if these are removed beforehand.

For the Ares I-X, NASA used an analysis of variance (ANOVA) approach based on Monte Carlo Sampling [59]. In the following, SAMBA is used to improve the efficiency of the ANOVA analysis that precedes the model calibration To find sensitive elements, Polynomial Chaos Sobol indices are used as introduced in Section 3.1.1. The NASA project engineers usually identify around 50 parameters that are relevant for the model calibration step. These are the parameters that are entered into the ANOVA analysis. Uniform probability distributions are assumed to describe their uncertainty and the distribution intervals are chosen based on the expertise of the project engineers. Due to the confidentiality of the NASA FEM model, the UQ approach is demonstrated in the following for a 'toy problem' of comparable physicality. For the toy problem, a 5% uncertainty for 32 mass and 32 stiffness parameters were assumed resulting in 64 uncertain model parameters.

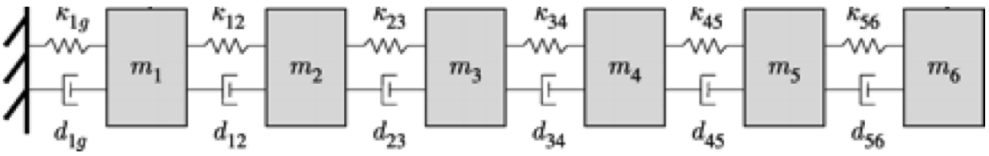


Figure 6.23: Illustration of the simplified problem (called 'toy problem') used to test and demonstrate the efficiency of new UQ methods.

Due to the volatile nature of the problem, a sparse Gaussian quadrature approach proved too inaccurate to deal with the volatile structural response surface. Instead, the principles of statistical learning have to be used to build a reliable Polynomial Chaos emulator. This was done in the following steps:

1. A random sample set of 100 Halton Sequence Samples was performed, where all 64 random parameters and the model were executed in all these locations.

2. The data set was split into learning and training data and a Polynomial Chaos model was fitted to the training data to estimate the prediction uncertainty

3. Further points were added to the Polynomial Chaos response surface utilising gradients in the response surface until a convergence criterion in the prediction error was fulfilled

The above procedure resulted in convergence (with leave-n-out point prediction errors lower than e^{-2}) after 472 full model runs. In comparison, 500 Latin Hypercube samples were performed for the Ares I-X with no mathematical guarantee of the reliability of the Polynomial Chaos surrogate model. To further improve the efficiency of the Polynomial Chaos method, it is recommendable to use Least Angle Regression [8] to solve the least squares problem. It efficiently identifies the polynomial chaos coefficients that are responsible for the largest errors and thus cleverly reduces fitting errors.

Figure 6.24 shows the results of the Sobol sensitivity study. The left graphs show the values of the masses and stiffnesses investigated and the right graphs their effect on the error cost function.

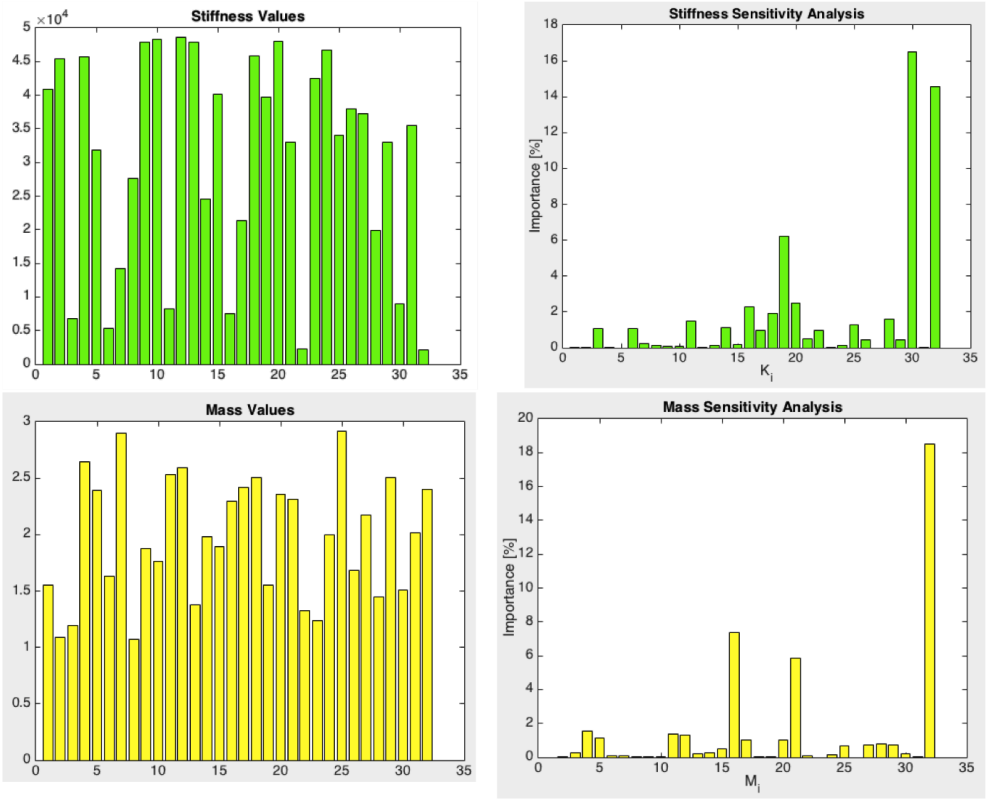


Figure 6.24: Result of the global sensitivity analysis for 64 parameters. It can be seen that the number of parameters that is important for model updating is smaller than the total number.

The above example shows that SAMBA can be used as an effective global

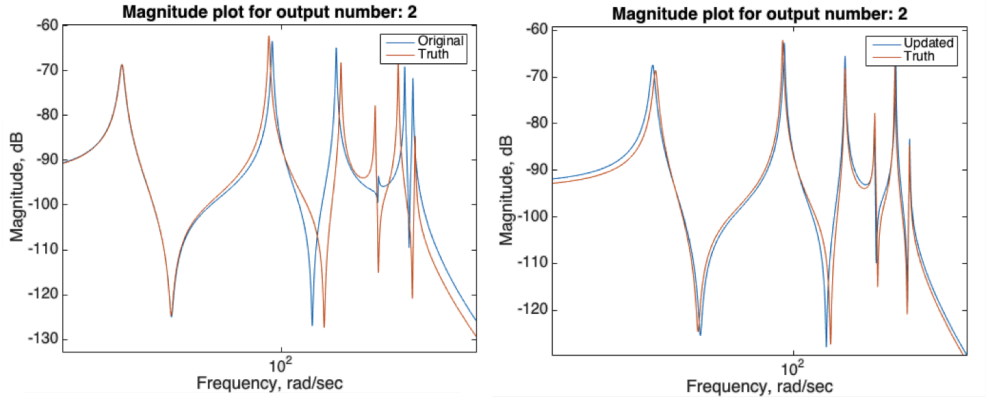


Figure 6.25: Improved calibration within the lower [0,200Hz] frequency domain using only active/sensitive parameters instead of all parameters.

sensitivity analysis tool within complex and computationally expensive simulations. In particular, if the parameter space is not sparse, as it is the case for finite model updating. SAMBA requires an additional computational effort of 5-6xN model runs, where N is the number of uncertain parameters. For the investigated test cases, on average, around 50% of parameters could be eliminated and deterministic calibration could be made 30% more accurate.

6.4 Rare Events in Aircraft Engine Simulations

Despite the decline in the number of aviation accidents since the 1990s and their low number generally, a single aircraft accident causes high losses in terms of human lives and property. From an insurer's point of view, it has one of the highest risks of potential catastrophes associated with a single accident, even though aviation is one of the smallest markets. Despite the high level of safety of modern jet engines, there is a continuing quest for the highest possible safety standard.

Even with today's designing techniques and manufacturing methods, it is still very difficult to achieve the safety margins required by the aviation authorities. One of the main problems is that the uncertainties introduced by manufacturing errors reduce the capability to predict the engine life accurately. Even state-of-the-art high-fidelity CFD methods assume idealised geometries, which do not exist in reality. They cannot be manufactured and their performance differs from real-world geometries.

Several authors have confirmed that the small stochastic variations caused by manufacturing deviations can result in considerable differences between the deterministically predicted and the actual system. For example, [9] found that the flow structure and heat transfer obtained for a digitalised real geometry was completely different from the results obtained for the idealised CFD geometry. [13] found that the typically occurring manufacturing variation of the cooling hole diameter in nozzles can reduce the life of a blade by 33%. Hence, disregarding the discrepancy between idealised model and reality can reduce the lifetime of components.

With UQ methods, manufacturing errors in the geometry, materials and flow parameters can be considered. This makes it possible to predict the overall variation due to manufacturing errors and to optimise the geometric shape, taking these errors into account. Previous research has shown that, by considering stochastic surface variations, geometries can be found which are much less affected by manufacturing errors [96]. This means that it is possible to design safer engines with uncertainty-based designing approaches without expensive changes in manufacturing methods, just by using today's technology.

Designing methods using UQ have been proven to work and yield CFD simulation results which are much closer to experimental observations [96]. However, they still have limitations. The most important limitation is widely known throughout the field but little discussed: in many cases, no information about the PDF of the stochastic parameters is known. Obtaining this knowl-

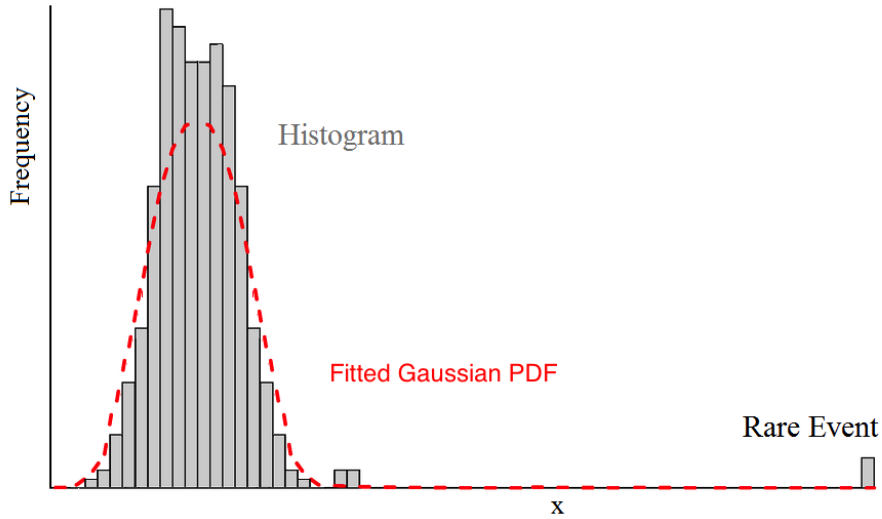


Figure 6.26: Example data set for which data-driven Polynomial Chaos is more reliable than PDF fitting for rare event simulation. The Gaussian PDF smoothes out the rare event in the tail region.

edge can be impossible, or, at the very least, requires a detailed knowledge of statistical data analysis. For this reason, engineers most frequently assume that their random parameters are normally distributed and defend this choice by referring to the Law of Large Numbers. This choice leads to an underestimation of the occurrence of rare events, such as aircraft accidents, as can be seen in Figure 6.27. It shows the number of fatal aircraft accidents from 1978 to 2007. One can observe that the empirical distribution deviates from a log-normal distribution. Instead, it exhibits a heavy-tailed (power-law) behaviour. UQ under the assumption of normality would have predicted no more than 100 fatal aircraft accidents during the last 40 years. In reality, however, there were 2166 accidents between 1978 and 2007 [142].

While it is, of course, far-fetched to attribute the total number of aircraft accidents to technical failure only, the graph still makes a point: in reality, distributions are not perfectly Gaussian and rare catastrophic events occur. In current UQ methods, however, this is not the case. If all unknown random parameters are assumed to have thin, exponentially decaying tails, the models can never predict a realistic probability for fatal system failure. As Weitzman argues for climate change models [149], it is much safer to assume PDFs with a relatively high probability of extreme events for cases where no or only a little information is available about the PDF. This holds for CFD as well: it is safer to use heavy-tailed distributions for unknown stochastic parameters. Unfortunately, their fatter tails result in infinite variance and higher stochastic

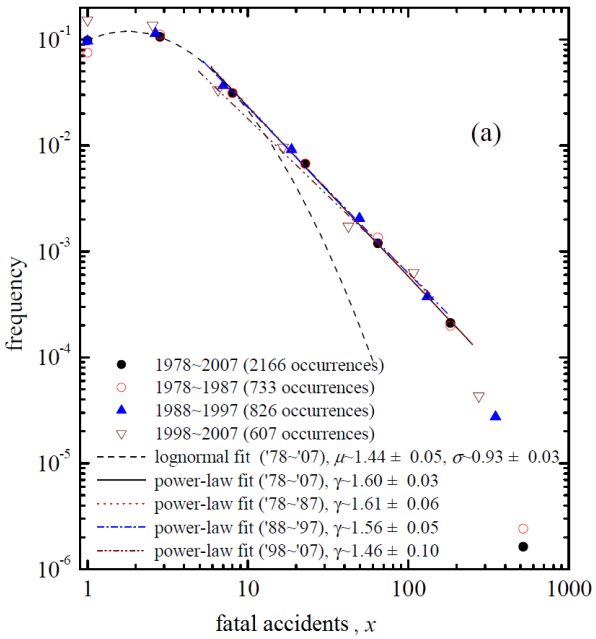


Figure 6.27: Empirical data showing that aircraft accidents occur with a heavy-tailed distribution, from Tseng [142].

moments. This mathematical property has so far discouraged the application of heavy-tailed distributions to physical problems.

6.4.1 Transients in Gas Turbine Cycle Missions

The details of the subsequent engineering model are explained in full in Montomoli [91] and are briefly summarised in the following. Running gas turbines follow prescribed missions for which the occurring metal temperatures are predicted using deterministic computational models. Figure 6.28 shows a representative flight cycle for the time-dependent metal temperature in a region of the turbine stator well, as it occurs in the transition process from a lower shaft rotation speed ($U = 0.226U_M$) to a higher speed ($U = U_M$). The transition phase is marked by two acceleration ramps: A slower (first) and faster (second) ramp. The slower ramp takes about 270s and the faster 9s. The temperature curves in the right graph of Figure 6.28 were obtained from a coupled CFD/FEM simulation for the solid 3D model. It is based on a real aircraft engine and the boundary conditions were assigned matching the engine's instantaneous operating conditions. Several locations are used for monitoring the temperature as described by Amirante [3] and as shown in the right graph of Figure 6.28. The temperature has been non-dimensionalised using the maximum value of annulus temperature T_M pertaining to the high speed following

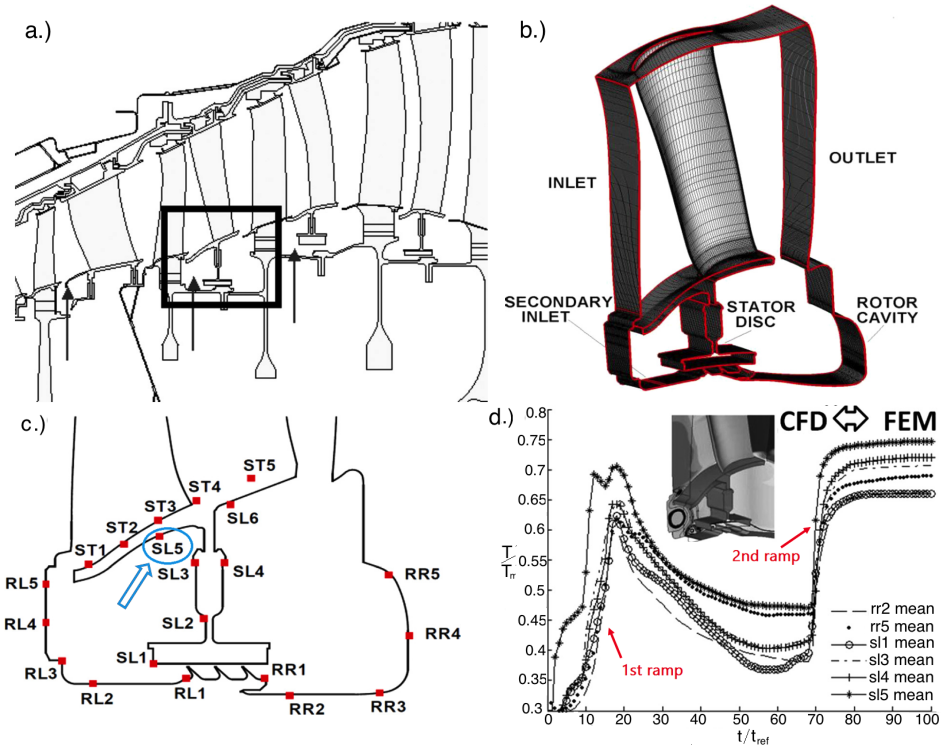


Figure 6.28: a.) Physical domain b.) Computational model c.) Temperature measurement point d.) Temperatures curves at different rotor cavity locations

the cycle and the time is non-dimensionalised with the overall mission time of the transition $t_{ref} = 4000s$. Under operating conditions, the mission cycle is exposed to stochastic variations. In the test bed, this variation was shown to be in the order of several seconds. In Montomoli [91] it was approximated with a Gaussian distribution with standard deviation $\sigma = 1.125t_{ref}\%$. Under real flight conditions, the uncertainty is expected to be even higher. The importance of considering time variation was shown in several technical reports, see for example Section 3.7 in Agard [37] or [85]. For example, a fast shut-down can lead to an engine life reduction of 200 hours [91]. In a similar way, high transient temperature gradients can also reduce the engine life and a metal variation of 20K changes the life about 50%.

To test the sensitivity of the model towards fatter tails SAMBA is used here. For this case, benchmark results are available from Montomoli [91], who showed that a small variation in the tail of the probability distribution describing mission time uncertainty, namely from Gaussian to Student-t with three degrees of freedom, is most strongly reflected where the temperature gradients are highest. The effect of fat tails was simulated by performing a Monte Carlo Method (MCM) with a Student-t temperature distribution in sensor $SL5$. Because a direct standard MCM was infeasible due to the complexity of the CFD model, Montomoli calculated a response surface and sampled this surface 10^8 times instead of running the entire model that often. To avoid non-physical time variations, negative time or infinitely large time,

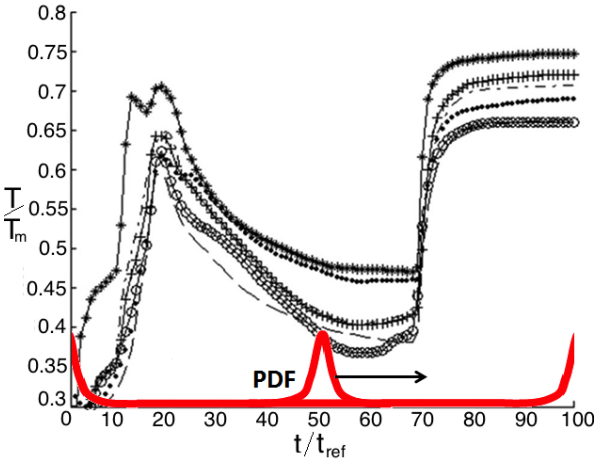


Figure 6.29: Illustration of how the PDF 'moves' through the truncated time domain to simulate timing uncertainty.

the time variations are simulated by making random draws from a time array reaching from 0 to 100 t/t_{ref} using a Student-t distribution with three degrees

of freedom, such that the current time is always the mean, but the tails of the used PDF are cut-off if they reach either $t = 0$ or $t = 100$, as depicted in Figure 6.29. Exactly the same domains are used to calculate the orthogonal polynomials for SAMBA. While a meta-model approach is computationally much less expensive than a direct MCM, it comes with certain disadvantages. First, sampling the response surface requires just as many runs as a brute-force MCM: 10^8 runs, plus the recreation of the response surface required 23 CFD simulations in this case. Second, the creation of a trustworthy response surface usually requires several checks so that the overall computational effort is higher than for SAMBA. SAMBA is, therefore, more efficient and needs significantly fewer model runs to evaluate the moments, especially if considering that only 6 CFD simulations and 6 model evaluation are required in total for a fifth order approach for only one uncertain variable. Figure 6.30 shows that SAMBA

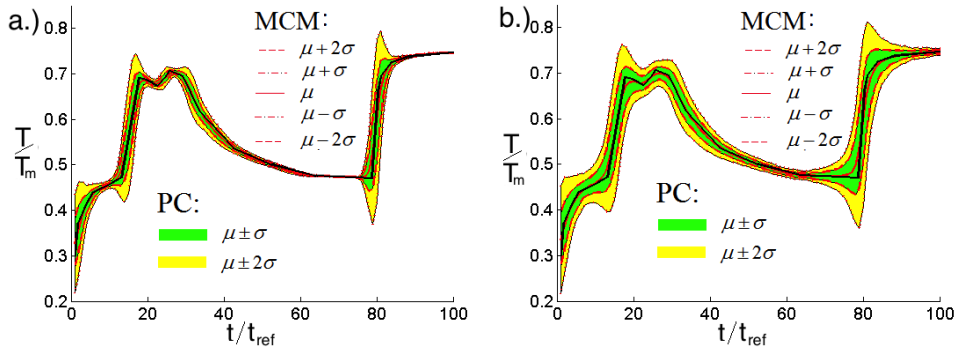


Figure 6.30: SAMBA and MCM comparison for the variability of the transient temperature based on a.) Gaussian and b.) Student-t input distribution.

provides exactly the same results with lower computational cost. Figure 6.30 a.) shows the results obtained for an input Gaussian distribution, and Figure 6.30 b.) for a Student-t distribution. The iso-contours of the mean value of the transient temperature with +/-1 and 2 standard deviations are shown against the dashed lines obtained by a Monte Carlo Simulation with a response surface.

6.4.2 Hot Gas Ingestion

One of the most important problems faced in modern gas turbine design is the ingestion of hot mainstream gas into the inter-wheel spaces between the turbine rotors and the spacer [88]. The reason is that the ingress of hot gas into the lower cavity can result in dangerous damage to the turbine rim, blade roots and discs with a reduction of component life [95]. To prevent hot gas ingestion,

a stable 'cold' purge flow is bled from the compressor and used to cool the cavity. However, being removed from the compressor, this colder air has been compressed (using mechanical energy) but has a minimum contribution to the turbine power output because it has not been energized by the combustion energy. This limits the amount of coolant that can be used because it directly affects the engine's efficiency, which leads to an increase in fuel consumption and CO_2 emissions. Therefore, a challenging task is to balance the purge flow across the spacer. Moreover, the presence of random gaps between the stator and the diaphragm complicates the problem with unwanted secondary air flow system leakages. In this work, state-of-the-art methods have been used for the CFD simulations in combination with SAMBA to calculate the statistics of the temperature in the lower cavity as well as the probability of hot gas ingestion. The most important aleatory factors influencing the probability are the sizes of gap 1 and gap 2 as displayed in Figure 6.31. How a stochastic

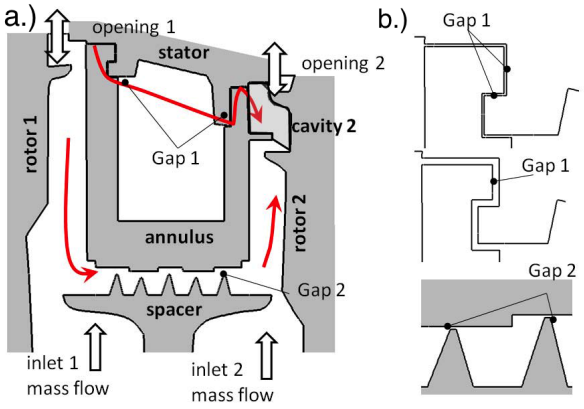


Figure 6.31: Illustration of a.) hot gas ingestion mechanism b.) the two gaps varied in the simulation.

variation of ±50% of the gap sizes influences the probability of failure was investigated by Montomoli [95] using a Meta-Model Monte Carlo Simulation based on 25 full CFD runs. The CFD simulation used in the following is the same as performed by Montomoli in the previous work: a quasi-3D simulation of the cavity region with the CFD solver CFX and the k-ω SST turbulence closure [150]. The mesh used contains 0.502M elements, which was shown to be sufficient, and it is comparable to the industrial standard in similar cases. All temperature values are normalised by the maximum temperature in the free stream gas path of a representative 32MW machine. Further details regarding the boundary conditions and the values for the reference temperatures can be found in Montomoli [95] and are based on realistic gas turbine

values and measured distributions in real gas turbines. The simulation and results presented here, however, are only an example of the application of SAMBA and the comments and conclusions do not reflect any specific gas turbine. Under nominal operating conditions, the cavity is correctly purged and the normalised temperature is 0.86, which is well below the critical value of $T_{crit} = 0.92$. However, the gap sizes can vary from their mean value due to manufacturing variation. The standard deviation of this uncertainty will be assumed to be $\sigma = 6\%$ in the following. For the application of SAMBA two histograms based on 1000 samples are drawn from Gaussian distributions for the two gaps sizes as displayed in the schematic overview in Figure 6.32. We assume that the true underlying distributions are Gaussian, although only a limited amount of samples are available. The use of a smaller set of measure-

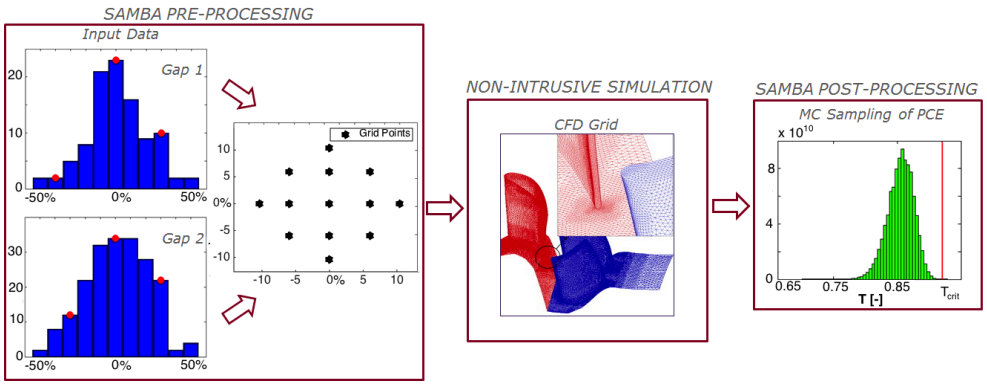


Figure 6.32: Schematic overview of how SAMBA is applied to the hot gas ingestion problem. The graph on the left shows the used histograms and the created sparse grid. The graph in the middle shows the CFD domain and mesh. The graph on the right shows the output distribution obtained when using Gaussian input PDFs.

ment data is not only more realistic but also demonstrates that SAMBA can propagate real measurement data without needing PDF fitting. The optimal collocation points for each individual input histogram are computed and combined into a nearly Gaussian looking sparse Smolyak mesh. Moreover, also two Student-t distributions are fitted to the data to perform rare event analysis with fat-tailed distributions. The grids in Figure 6.33 show that the Gaussian grid takes only the main bulk of the distributions into account, whereas the fat-tailed grids dedicate one point to explore the failure regions. The Gaussian mesh is not only shown in the schematic, but also in Figure 6.34 a.) in front of the contour lines of the response surface created through varying the two gap sizes. The response surface in Figure 6.34 b.) was set-up for the Monte

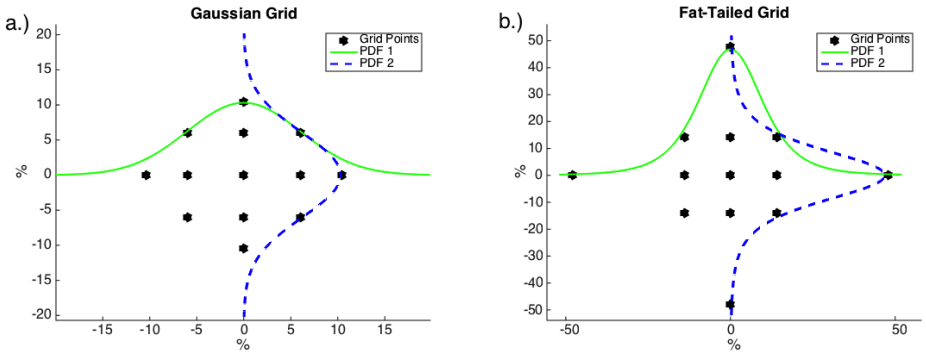


Figure 6.33: SAMBA grids obtained for a.) two Gaussian and b.) two KoBoL input distributions.

Carlo Meta-Model, which is also displayed in the figure. Most importantly, it can be seen that the manual set-up grid for the meta-model overestimates the relevant region to evaluate the effect of the inputs, while it under-samples the important region. Choosing the relevant region and an appropriate sampling rate is the main difficulty in manually designing meta-models. SAMBA, on the other hand, provides an optimal Design of Experiment (DoE) based on the moments. Practically, the resulting Smolyak grid represents a list of various gap sizes for which the CFD code needs to be executed. From the 13 CFD code results, the moments of the stochastic behaviour of the temperature can be derived. Moreover, by sampling the PCE the shape of the posterior PDF can be obtained. Since the lowest level Smolyak model only contains 5 points

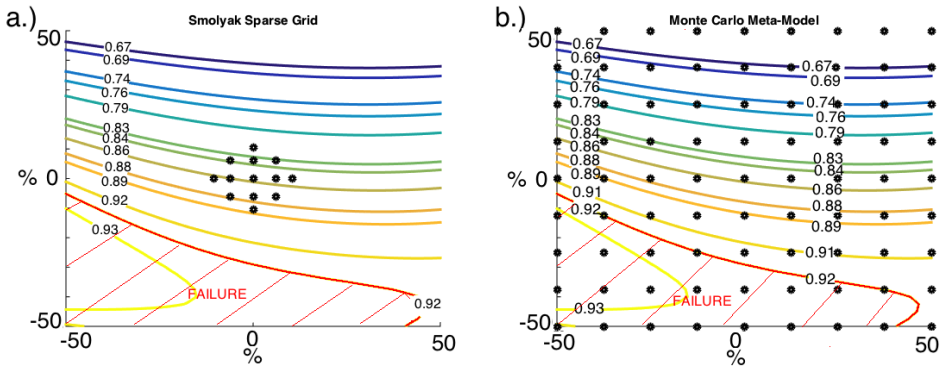


Figure 6.34: Contour plots obtained with a.) SAMBA and b.) Monte Carlo meta-model.

close to the mean, the second level Smolyak grid was used. To calculate the probability of failure, the PCE based on a full tensor with 25 simulations, a Monte Carlo with 25 simulations, direct importance sampling, and a Smolyak

sparse grid with 13 simulations were compared. To obtain estimates of the probability distributions the meta-model and the PCEs were sampled with 10^8 samples. For importance sampling, the sampling distributions for the gap sizes were changed from $N(0,6)$ to $N(37,6)$ for gap 2 and $N(17,6)$ for gap 1. This creates samples mainly in the estimated region of failure indicated by the red stripes in Figure 6.34. In this way, the probability of failure P_f of the response function of the CFD model $f(x,y)$ can be estimated orders of magnitudes more efficiently by calculating

$$P_{f,IS} = \frac{1}{N}\sum_{i=1}^{N} f\left(x_i,y_i\right)\frac{w\left(x_i,y_i\right)}{g\left(x_i,y_i\right)} = \frac{1}{N}\sum_{i=1}^{N} f\left(x_i,y_i\right)b\left(x_i,y_i\right) \tag{6.1}$$

where $w\left(x,y\right)$ and $g\left(x,y\right)$ are distributed as

$$w\left(x,y\right) \sim N\left(\left[\begin{array}{c} 0 \\ 0 \end{array}\right], \left[\begin{array}{cc} 6 & 0 \\ 0 & 6 \end{array}\right]\right) \qquad g\left(x,y\right) \sim N\left(\left[\begin{array}{c} 37 \\ 17 \end{array}\right], \left[\begin{array}{cc} 6 & 0 \\ 0 & 6 \end{array}\right]\right). \tag{6.2}$$

While Importance Sampling is of course less efficient for a mere two uncertain input variables, it becomes quickly more attractive as the dimension of the problem increases, because it is not affected by the curse of dimensionality. The results of all methods are summarised in Table 6.5. The effect that the fat

Table 6.5: Probability of critical hot gas ingestion

Method	Order	Runs	$P_{fail}\ N(\xi)$	$P_{fail}\ t(\xi)$
Importance Sampling		10^3	$2.2495 \cdot 10^{-4}$ %	3.57 %
Meta-Model MCS		25	$2.04557 \cdot 10^{-4}$ %	3.22 %
Tensor Quadrature	$p = 4$	25	$2.05017 \cdot 10^{-4}$ %	3.21 %
SAMBA PC	$l = 2$	13	$2.02901 \cdot 10^{-4}$ %	3.04 %

tails of the input distribution have on the output PDF is shown in Figure 6.35, where T_{crit} indicates the failure region. The results show that sparse grids can be used to predict rare events and probability of failures. The advantage of using SAMBA is that the lower moments of the posterior probability distribution can be obtained quite efficiently for larger problems where only small data sets and no PDFs are available.

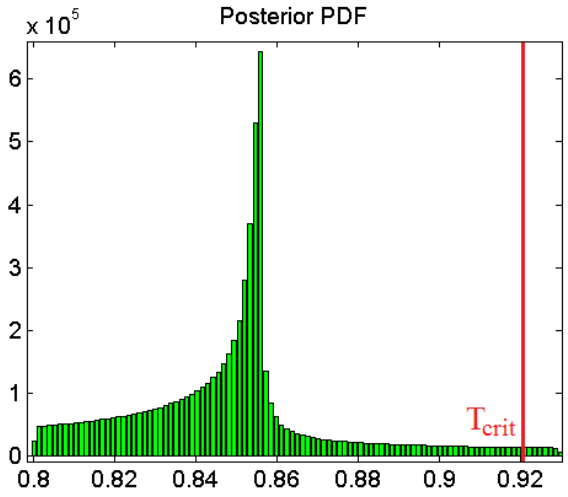


Figure 6.35: Output PDF for engine temperature when using fat-tailed input distributions. It can be seen that the critical temperature T_{crit} is clearly exceeded.

6.5 Discontinuities in Aircraft Engine Simulations

Discontinuities caused by compression shocks are encountered in many engineering applications that involve high-speed flows or flows with large pressure differences in gas turbines, internal combustion engines or supersonic aircraft. To demonstrate the multivariate Padé approximation method developed in Chapter 4, an uncertainty quantification study is performed for a realistic configuration of a cooling flow interacting with various shocks in an HP nozzle.

6.5.1 Film Cooling and Shock Interaction

An effective strategy to increase the efficiency and the specific power output of a gas turbine is to increase the Turbine Entry Temperature (TET). This, however, increases the load on the blades as shown by Salvadori et al. [126] and also the temperature on the blade surfaces, which gets closer to the melting point of the alloys used in HP nozzles. Only advanced cooling systems can make it possible for the nozzle to survive in such conditions. Film cooling is the most effective technique to control the high temperature. It controls the temperature of the flow inside the nozzle by injecting a cold flow acting as film-protection for the surface blade. The interaction of the cooling flow and the main flow in the passage leads to flow structures characterized by high

complexity, in which several parameters play a crucial role in the thermal efficiency of the cooling device.

Montis et al. [90] could connect the features of the cooling flow with the increasing of the losses through the nozzle. Andreopulos and Rodi [4] were the first to identify a mechanism of redistribution of the vorticity content in the boundary layer inside the coolant hole duct to generate two contra-rotating vortexes usually known as 'kidney vortex'. A numerical investigation by Sgarzi et al. [71] provided a detailed study of this three-dimensional structures over flat plates. Hagen and Kurosaka [56] demonstrated the presence of a core-wise, cross-flow transport in hairpin vortexes (created by coolant injection) in the laminar boundary layer. These vortexes continuously move the fluid from the wall to the free stream. These flow structures are known as 'tornado vortex'. More recently similar structures have been identified by Carnevale [16] as consequence of the interaction of the cooling flow with a shock.

The vortex structures generated by the interaction between the coolant flow and the main flow is strongly dependent on global quantities such as the main flow velocity and temperature, and on local quantity such as the coolant blowing ratio, and local pressure of the coolant region. These parameters are strongly affected by the presence of the shocks, and some geometrical parameters. Manufacturing deviation and machine wear due to phenomena such as fouling (as detailed in Casari et al. [17]) contribute to the increase of the uncertainty on the actual efficiency of the cooling device. The channel is contoured in order to reach a sonic condition in the lower region near the cooling hole. A system of shocks is generated and interacts with the cooling flow. The flow features are characterised by high complexity and the efficiency of the coolant device is strongly sensitive to small deviations in the boundary conditions and geometrical parameters.

In this example, the geometry is parametrised with two random variables: a uniformly distributed diameter of the coolant channel, and a similarly uniformly distributed diameter of the leading edge of the blade located in the middle of the channel. These two random parameters affect both the shock angle and local characteristic of the cooling flow. As output quantity, the adiabatic effectiveness η is considered. According to Montomoli et al. [93] the typical manufacture deviations associated with the coolant and the trailing edge are respectively 10% and 20%. A UQ study for geometric uncertainty for a similar arrangement of film cooling shock interaction was previously conducted by Carnevale [16] also using Padé approximations. However, the authors only used univariate approximations and could only consider the inputs separately.

This study is thus new in the sense that it can consider connections between the geometric parameters.

The CFD simulations used were developed and previously investigated by Salvadori et al. [127]. A complete description of the experimental set-up can be found in Ochs et al. [105]. Only one cylindrical hole (out of five) is considered and it is assumed that the flow is symmetric. The final control volume is shown in Figure 6.36a and includes the converging nozzle, the shaped plate, the cylindrical hole and the plenum. According to Garg [46] and Acharya [2], the shape of the plenum in film cooling simulations needs to be properly modelled because of its effect on the boundary layer development. Therefore, the plenum has been fully resolved and the inlet conditions for the coolant are imposed at the plenum inlet. Figure 6.36b shows the main flow features of the present problem. The main-flow reaches the speed of sound at the throat while the shape of the lower part of the contoured plate allows the flow to accelerate further to supersonic velocities generating an oblique shock wave. The shock starts at the trailing edge and impinges on the lower wall in a region a few diameters after the cooling hole exit position. The flow structures, shown in Figure 6.36b, are referred to as the deterministic reference case. As previ-

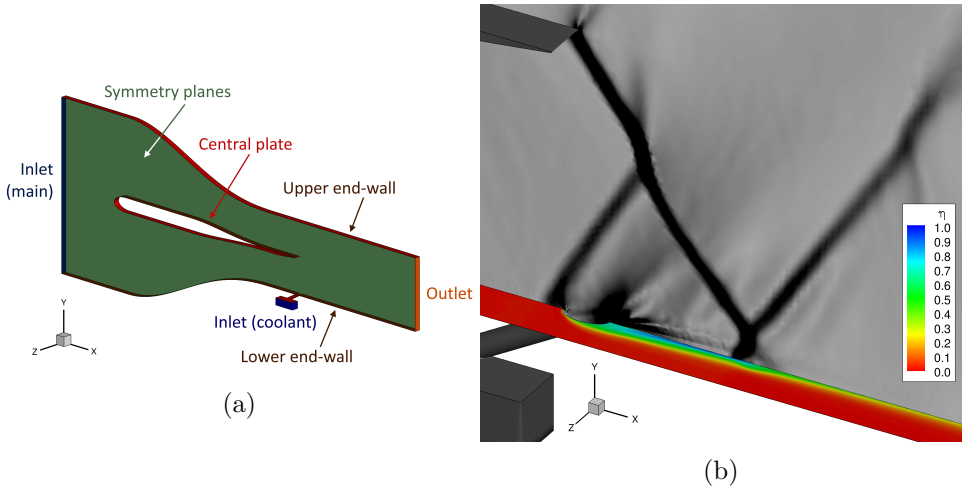


Figure 6.36: (a) Control volume (b) Flow features

ously mentioned, the test case is designed to have an injection hole positioned upstream of the shock impingement position. The original pitch-to-diameter ratio of the five holes is 4. The free stream Mach number at the shock location is around $Ma = 1.5$ (but the numerical data and the experiments suggest that it reaches higher values) and the expected blowing ratio is $M_B = 1.0$. The non-dimensional coolant total pressure is 1.01 and non-dimensional coolant

total temperature is 0.557: these values are normalised with respect to the main-flow inlet total pressure and total temperature, respectively. As can be observed in Figure 6.36b, a weak shock is generated by the interaction between the coolant and the main-flow. This shock interacts with the oblique shock shed by the plate trailing edge that impinges on the lower wall. Furthermore, a lambda shock is generated in the impingement region. It is to be expected that the adiabatic effectiveness will be modified by the adverse pressure gradient associated with such kind of flow structure. It is usually also responsible for a local growth of the boundary layer and the insurgence of vortex structures.

The efficacy of the coolant coverage is quantified by the adiabatic effectiveness η_{aw}. It is evaluated with

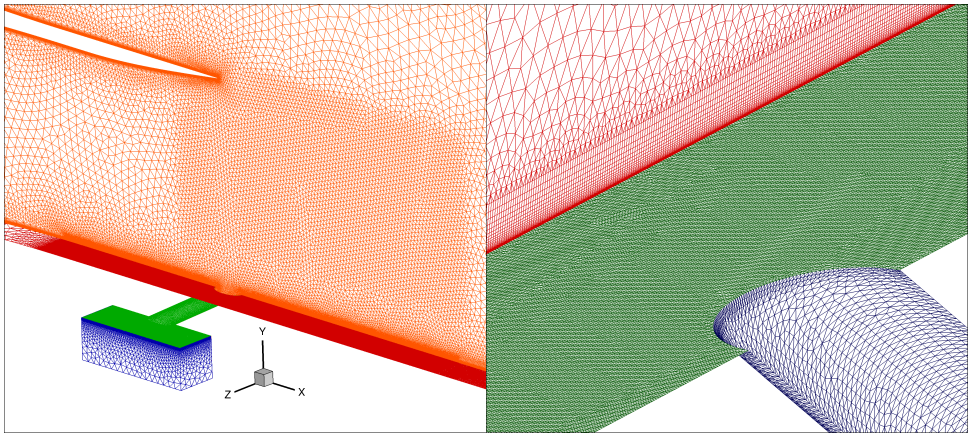
$$\eta_{aw} = \frac{T_{aw} - T_{rec,m}}{T_{0,c} - T_{rec,m}}. \tag{6.3}$$

The probability distribution of η_{aw} is the parameter considered for the UQ study described later. The adiabatic effectiveness is extracted along the centre line and compared with the experimental data for validation purposes. The distribution of the adiabatic wall temperature T_{aw} is evaluated on the cooled plate by performing adiabatic calculations. The calculation of the adiabatic effectiveness requires also the evaluation of $T_{rec,m}$:

$$T_{rec,m} = T_{0,m} \frac{1 + Pr^{0.33} \frac{\gamma-1}{2} Ma_{is}^2}{1 + \frac{\gamma-1}{2} Ma_{is}^2}. \tag{6.4}$$

Grid and Test Case Modelling

Steady Reynolds-Averaged Navier-Stokes (RANS) equations are solved with a pressure-based approach. It must be underlined that RANS methods are still widely used in industry since they are able to capture the overall phenomenology. Pressure and density equations are coupled and the calculation is fully second-order accurate. The k-ω SST model by Menter [87] is used as turbulence closure with compressibility effects and a low-Reynolds correction. Inlet turbulence level is set at 5% while a turbulence length scale of $1mm$ has been considered. Viscous walls have been set as adiabatic and viscous heating has also been considered to make the heat transfer evaluation accurate. The calculation is considered to be converged when all the residuals are constantly below 10^{-4} and the mass-flow error is below 1%. The meshes used for all the nine calculations (necessary for the UQ analysis) are generated using the commercial hybrid grid generator $Centaur^{TM}$ by Centaursoft. The mesh is denser in the region where the shock-boundary layer interaction is expected



(a) (b)

Figure 6.37: CFD grid details in coolant/shock interaction region.

(see Figure 6.37a) and the overall quality of the mesh is very high. The mean dimension of tetrahedral elements downstream of the coolant hole has been reduced to better capture the local phenomena, with a very high resolution of the lower surface (see Figure 6.37b). A y+ far below 1. has been obtained for the first prismatic layer and at least five prismatic elements are included in the laminar sub-layer.

CFD Results

Figure 6.38 shows a comparison between the numerical and experimental maps of η_{aw}. The numerical simulation greatly underestimates the lateral spreading of the cooling flow compared to the experiment by Ochs [105]. The coolant flow seems to be concentrated around the centre line in the numerical simulation, whereas the experiment shows a wider redistribution of the flow coming from the hole. Figure 6.39 compares the film cooling effectiveness distribution in stream-wise direction on the centre-line. The obtained result can be explained looking at the streamlines exiting from the coolant flow for the nominal case. In Figure 6.40 the isentropic Mach number distribution is shown on the coolant symmetry plane and lower wall and the streamlines exiting from the cooling hole. As can be observed, the cold flow is initially kept near the lower wall, coherently with the adiabatic effectiveness maps in Figure 6.38.

A few diameter lengths later (approximately $x/D = 5.0$), the adverse pressure gradient associated with the shock impingement generates a lift-off effect of the coolant that is visible in the smaller image in Figure 6.40, where the density gradient magnitude is reported. In that region, the boundary layer

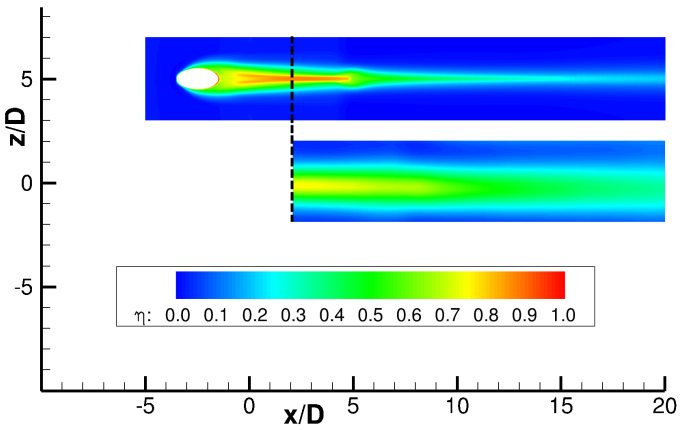


Figure 6.38: Comparison between CFD (upper) and experimental data (lower) for the deterministic reference configuration case. The numerical simulation greatly underestimates the lateral spreading of the cooling flow.

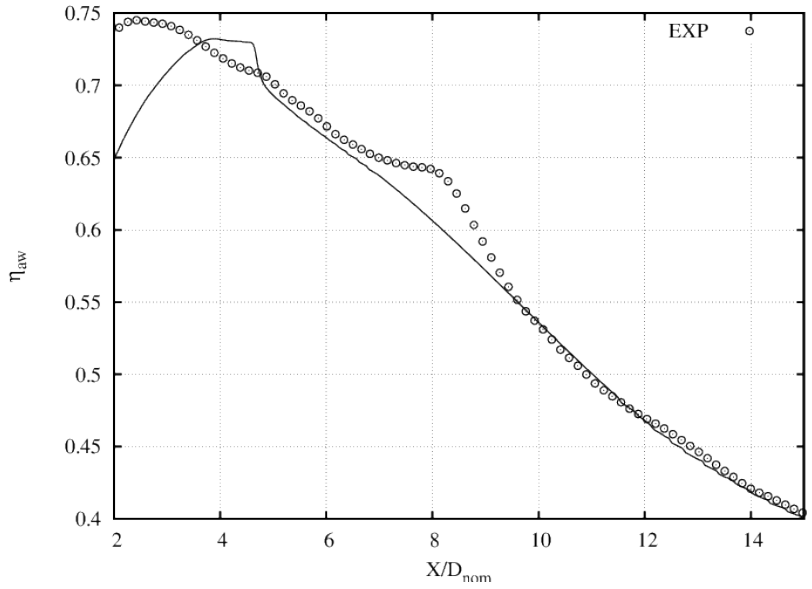


Figure 6.39: Film cooling effectiveness η_{aw} along the centreline of Figure 6.36

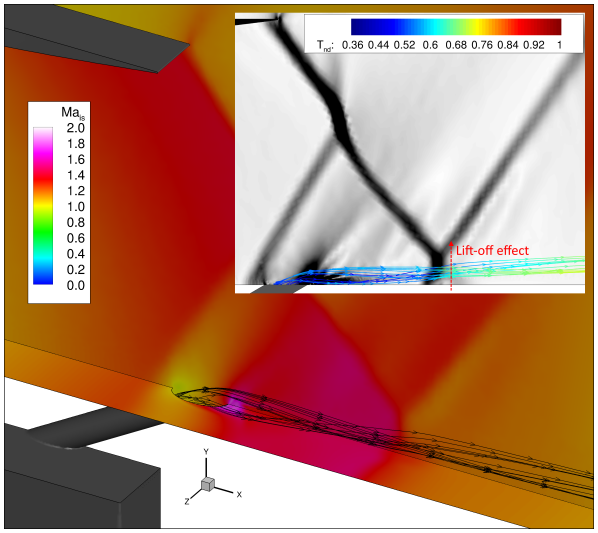


Figure 6.40: Flow streamlines and interaction with the shock.

height is increased by the shock impingement and the cooling flow is detached from the wall. After the shock interaction, an additional vortex system is generated with the axis normal to the wall as shown in Figure 6.41. This vortex is usually named "tornado vortex". As a consequence of this vortex, the coolant flow does not reattach before the end of the control volume. This explains why the cooled end-wall is so narrow in the numerically predicted adiabatic effectiveness map.

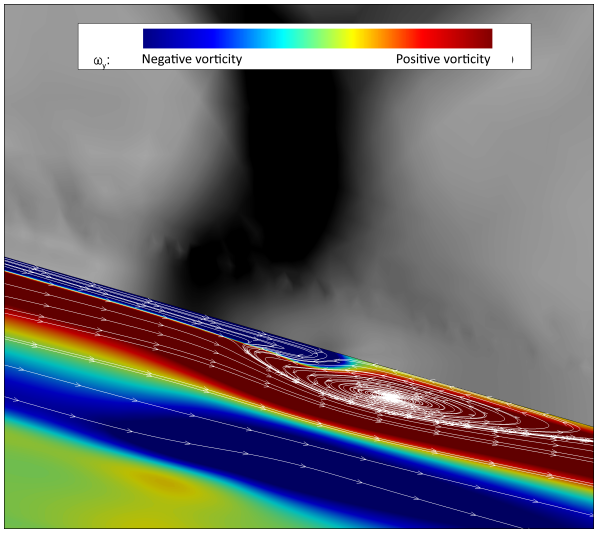
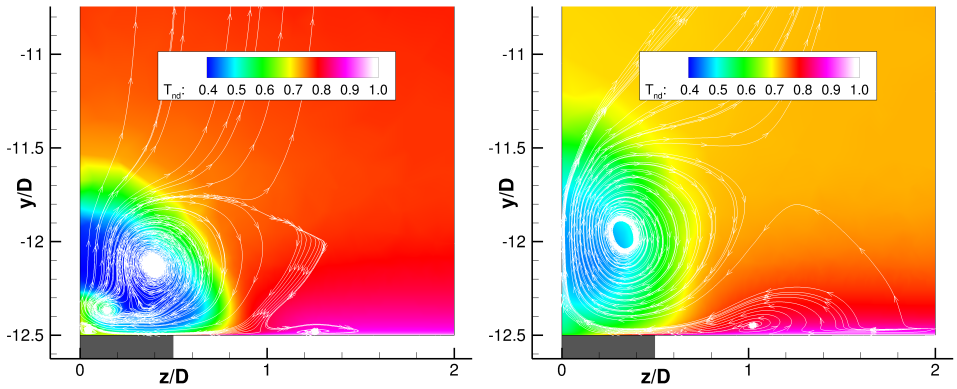


Figure 6.41: The tornado vortex after the shock.

A more detailed analysis of the local interaction between the main flow and

the cooling flow is reported in Figure 6.42, where two slices perpendicular to the end-wall are considering a region between the hole and the shock impingement position. The streamlines are superimposed on the temperature profiles to visualize the local vortex formation and the cooling flow redistribution. At the section $x/D = -1.0$ the cooling flow is initially located in a region next to the end-wall (Figure 6.42a) where is affected by a strong system of a recirculating vortex. These kinds of vortexes are generated by the interaction of the boundary layer inside the channel flow and are usually referred to as "kidney vortex". That flow structure is responsible for the rapid movement of the coolant flow far from the wall. In a downstream location such as $x/D = 2.0$, the coolant flow is completely detached from the end-wall (Figure 6.42b) and the hot main-flow substitutes the cold flow thus decreasing the adiabatic effectiveness value. This mechanism is strongly dependent on geometrical parameters such as the coolant diameter. The dimension of the kidney vortex is affected by shape and the dimension of the coolant hole. The interactions with the shocks and their angle are affected by the trailing edge of the next blade. All described flow features are coherent with the previously observed tornado vortex phenomenon, which according to Hagen and Kurosaka [56], Carnevale et al. [16] and recently Zhou et al. [160], is responsible for the local reduction of the cooling effectiveness.



(a) $x/D = -1.0$ (b) $x/D = 2.0$

Figure 6.42: Temperature maps and streamlines between hole and shock for different x locations as shown in Figure 6.38.

UQ Results

In this UQ study, the above-described CFD case is considered dependent on two random variables: a uniformly distributed diameter of the coolant channel

$\xi_1 \sim U[0.9, 0.1]$ and a similarly uniformly distributed diameter of the leading edge of the blade located in the middle of the channel $\xi_2 \sim U[0.04, 0.06]$. Using the adiabatic effectiveness η as output quantity, three-dimensional discontinuous response surfaces are constructed for the two uncertain parameters using the least squares multivariate Padé approximation method developed in Chapter 4.

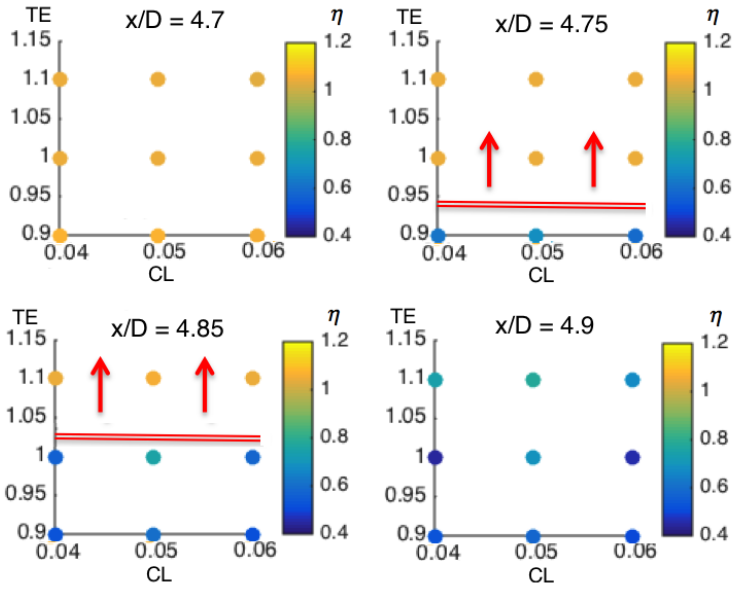


Figure 6.43: The graphs illustrate how the shock moves through the many stochastic response surfaces of the thermal efficiency. Each of the nine points represents a CFD simulation. Until x/D=4.75 the response surfaces are smooth. Then, a shock (illustrated by two red lines) moves through the domain until x/D=4.85.

A test matrix of 9 deterministically run CFD cases was already available. The top four graphs in Figure 6.43 show the test matrix at four different x-locations, which are the nine different combinations of the three values of trailing edge diameters $[0.9, 1.0, 1.1]$ and cooling hole diameters $[0.04, 0.05, 0.06]$. Naturally, an execution of the model in the optimal sub-sample quadrature point locations would have been much preferable. However, with the idea in mind to create realistic industrial working conditions, it is assumed that the maximum number of time-wise feasible simulations nine.

The four graphs in Figure 6.43 show how the discontinuity moves from the bottom to the top of the stochastic domain between $x/D = 4.7$ and $x/d = 4.9$. At $x/D = 4.7$ all values of the response function have nearly the same level. In the graph right next to it, however, at $x/D = 4.75$ a clear drop can be seen

between two domains: one levelled domain at the top and another levelled domain at the bottom at a trailing edge diameter of 0.9. The same drop in the stochastic domain can be observed for trailing edge diameter values of $0.9 - 1$ at $x/D = 4.85$. The location $x/D = 4.9$ is located after the shock so that the domain is at nearly at one level again.

Because the CFD mesh in the shock region is extra fine, hundreds of such discontinuous shock responses have to be fitted automatically and reliably. To that end, the automated multivariate Padé approximation described above is used. Various orders of M and L are automatically tested and fitted to the available CFD points. All parameter possibilities for one response surface are displayed in Figure 6.44. In the considered example a clear minimum is

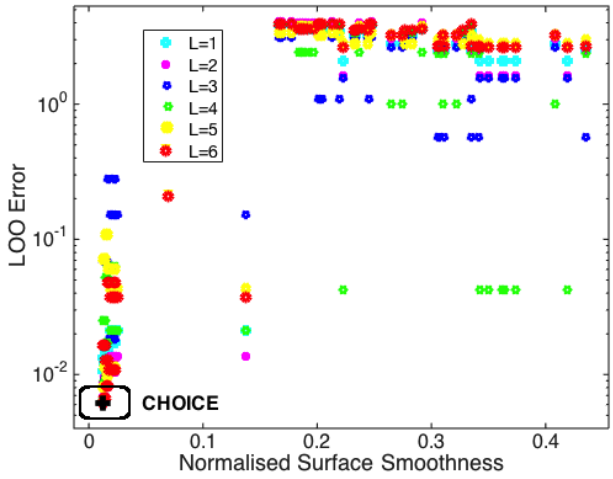


Figure 6.44: Choice of the automatic response surface as a trade-off between leave-one-out error and smoothness indicator.

found by the method. In other cases, a Pareto front is found and an equally weighted trade-off between the two objectives is made. The graphs in Figure 6.45 show three examples of surfaces that were automatically fitted by the approach. Throughout the CFD domain, not only discontinuous but in absence of shocks also smooth response surfaces occur. The smooth response surfaces before (1) and after (3) the shock in Figure 6.45 were fitted by the method with a denominator order of $L = 1$. This means that no discontinuity was found, which of course agrees with the physical data. It shows that the Padé approximation is also quite suitable for lower order approximations of continuous response surfaces. At the shock locations the value was mostly $L = 2$. Thus, only a single discontinuity was found by the method for certain values

of the trailing edge diameter. The probability distributions found from these model evaluations are shown in the bottom row of Figure 6.45.

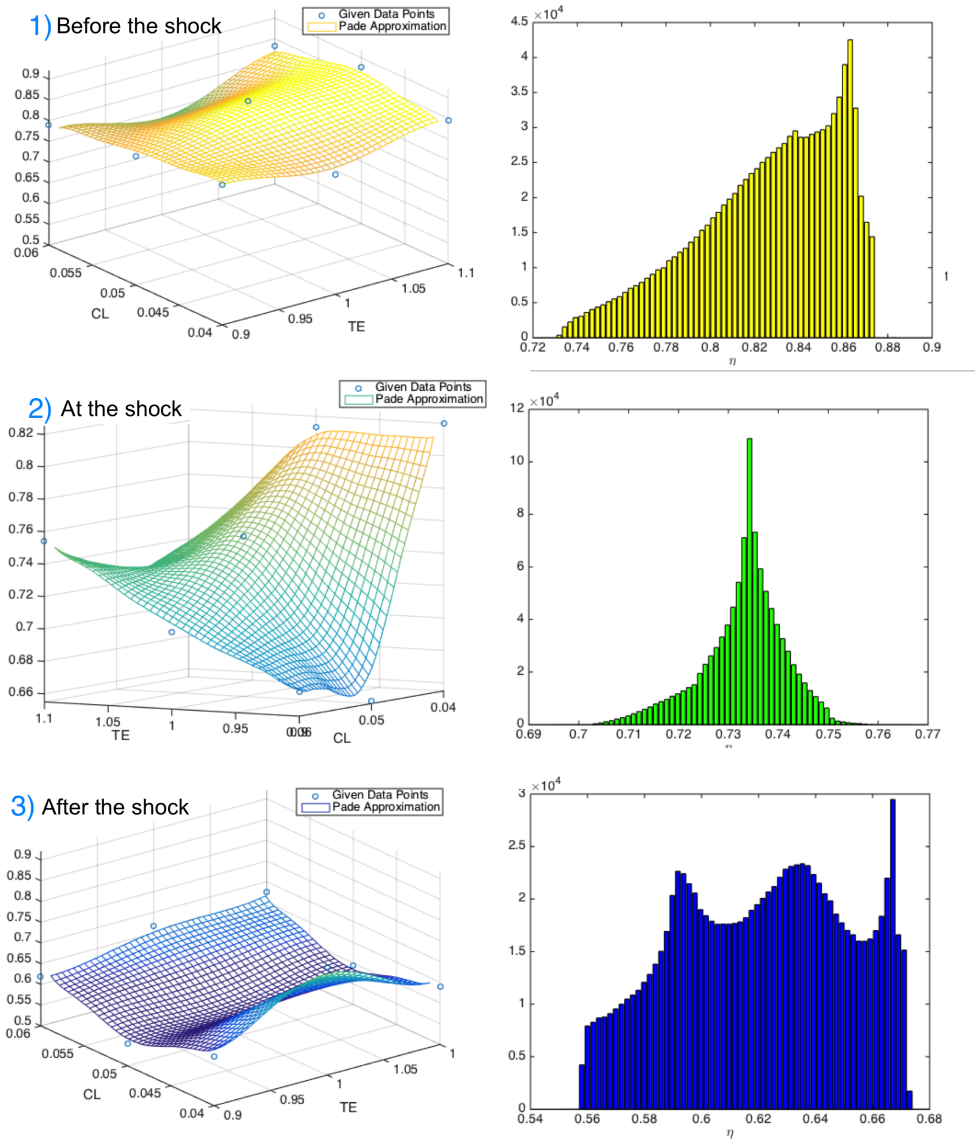


Figure 6.45: Three examples of automatically fitted Padé response surfaces before (1), at (2) and after (3) the shock location. The bottom row shows three PDFs to the x-locations before, at and after the shock.

They show the following: before the shock, uncertainty in the diameters is twice as likely to cause a higher thermal efficiency than a lower one. The shock location itself is only weakly influenced by the uncertainty. It remains stationary with a high probability. However, the resulting distribution is clearly fat-tailed. After the shock, the distribution seems to be in three states with

peaks at 0.59, 0.63 and 0.67. After fitting response surfaces to all x locations the Polynomial Chaos coefficients of the expansion are used to a) calculate mean and standard deviation as shown in the top graph of Figure 6.46 and b) the Sobol indices to understand to what proportions the diameter of the trailing edge and cooling holes contribute to the experienced uncertainty. The results, shown in the graph at the bottom of Figure 6.46, are quite coherent with the authors' expectation that it should be the cooling hole diameter which is more significant. An interesting observation made from this sensitivity analysis is that the uncertainty in the output is significantly reduced at the shock location. The bottom graph of Figure 6.46 shows that at $x/D = 4.8$ the standard deviation goes almost down to zero. Another interesting observation is that the cooling effect is strongly influenced by the shock. It sharply decreases before the occurrence of the shock and increases afterwards. The trailing edge diameter, on the other hand, seems to have more of a constant effect on the standard deviation.

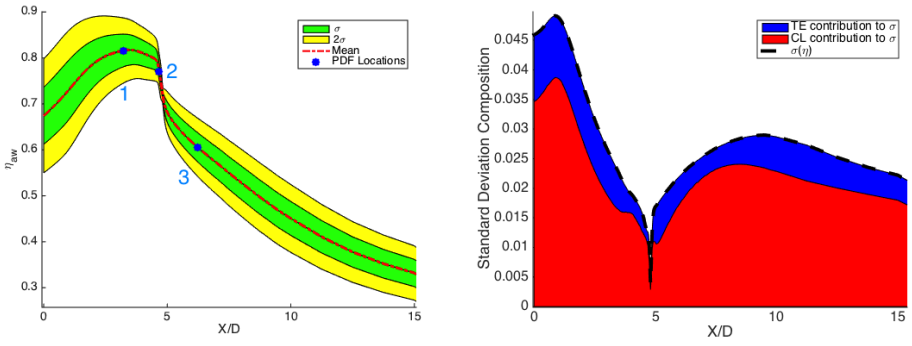


Figure 6.46: Mean and standard deviation of the film cooling effectiveness (left) and sensitivity analysis of input parameters using Sobol Indices (right)

7 Conclusion

This thesis presented a range of novel mathematical and computational Uncertainty Quantification (UQ) techniques to overcome problems in industry to efficiently integrate manufacturing, experimental or operational data into engineering design simulations. In particular, five areas of interest for industrial engineers were investigated:

1. Dealing with data scarcity and limited availability of trained statisticians in **Chapter 2**.
2. Mitigating the growth of computational cost with increasing number of variables (the curse of dimensionality) in **Chapter 3**.
3. Quantifying the likelihood of rare events with catastrophic consequences, the so-called Black Swans, in **Chapter 3.3**.
4. Dealing with discontinuities in models in **Chapter 4**.
5. Dealing with epistemic model-form uncertainty in **Chapter 5**.

Various practical examples from the fields of gas turbine, race car and rocket engineering were performed in **Chapter 6** to demonstrate that data-driven UQ methods provide a variety of advantages for practical engineering applications. In summary, these advantages are: UQ can help engineers to understand trends and intricacies in complicated designs; it can help make better decisions without requiring detailed physical knowledge of a problem; it can assist with identifying or localising problems; it can quantify the probability of failure of engineering components; it can provide a reliable meta-model that can be used for robust design.

In the future, UQ methods will play a vital role in the design of more realistic **Digital Twins**, digital replica that combine data and simulations to mirror their real-life counterparts [50], and to reduce designing time and costs during the transition to **Industry 4.0**.

7.1 Summary of Contributions

This PhD dissertation contributed a new algorithm for data-driven Polynomial Chaos and extensively demonstrated its advantages for a list of realistic engineering applications. It was shown how the method can be used to accurately and efficiently incorporate pure statistical random data into computational models without the need to alter statistical information by fitting probability distributions to the data. The method was shown to be particularly useful for applications that are characterised by scarce data and high safety requirements, as it is the case in Formula 1 race car engineering or gas turbine design.

Moreover, data-driven Polynomial Chaos was extended to multiple dimensions for the first time in this work. This sparse extension was denoted as SAMBA and applied to various engineering examples from Formula 1 to gas turbine engines including simulations containing complex shock interaction, and to the structural dynamics of a NASA rocket. These examples demonstrated that non-intrusive uncertainty quantification methods can indeed be applied to all kinds of computational models and that it provides a powerful tool to predict real-life performance volatility, estimate reliability or failure, prevent rare events or optimise performance under uncertainty (robust design).

By extending the mathematical theory of data-driven Polynomial Chaos and applying it to a series of complex and realistic engineering cases, the following advantages of SAMBA were demonstrated in this work:

- It provides the highest accuracy compared to all Polynomial Chaos methods for problems characterised by scarce data or random data that does not follow parametric probability distributions.
  This property is particularly useful for race car design in Formula 1 because race cars are produced in very low quantities.

- It can create anisotropic and adaptive Gaussian quadrature rules with the same complexity that other methods need to create isotropic uniform grids.
  These quadrature rules can be used for any mathematical problem that requires multi-dimensional integration.

- It provides a method to perform highly efficient global sensitivity analysis for large, complex or non-linear problems.
  This property is particularly useful for very large and complicated engineering models like the NASA Space Launch System.

- It can be used as an efficient method to perform rare event simulation

with a single global model that represents both mean and tail probabilities accurately by using truncated fat-tailed distributions or Kobol distributions.
This property can be used to prevent hot gas ingestion or harmful temperature transients in aircraft engines.

Further accomplishments of this work are two methods to complement the data-driven Polynomial Chaos framework to better deal with special cases:

- Multivariate Padé approximations were stabilised, automated, and modified in order to work with arbitrary points. The result is the first autonomous global and discontinuous non-intrusive Polynomial Chaos method.

- A multi-fidelity framework was tested for DNS and RANS for the first time to combine high-fidelity and low-fidelity simulations to reduce model-form uncertainty in CFD.

7.2 Future Research

The work in this thesis provides several potential future research opportunities. The biggest opportunity is a better combination of the data-driven Polynomial Chaos and **Machine Learning**.

On the one hand, machine learning techniques could benefit from data-driven Polynomial Chaos. The focus in this work was mostly on the advantages for applications with scarce data because one of biggest challenges in all practical engineering applications is the lack of experimental data. However, data-driven Polynomial Chaos expansions also have significant benefits for big data sets as they are encountered in machine learning or predictive maintenance applications. Data sets in typical machine learning applications are generally large, do not follow parametric probability distributions and can be multi-modal. This makes it challenging to use parametric probabilistic models. Since data-driven Polynomial Chaos can be regarded as a data compression algorithm for arbitrary data sets, it has a unique advantage to efficiently describe such non-parametric data sets. Future work could consequently investigate SAMBA as a **Big Data** approach for predictive maintenance. Sudret [139] already identified a need for better data-driven orthogonal polynomials with regard to introducing Polynomial Chaos as a new machine learning regression technique.

On the other hand, data-driven Polynomial Chaos could benefit from new methods in machine learning. There are many synergies between Design of Experiment (DoE) in UQ (like the optimal quadrature points provided by a SAMBA sparse grids in Chapter 3) and **Active Learning**, incremental learning or semi-supervised learning. In both cases, DoE and Active Learning, the incentive is to keep the computational cost low for models with long run times. As shown in Chapter 4, a UQ method can save computational time by using statistical learning. Instead of designing an experiment up front, the method iteratively chose the next simulation configuration to execute, based on available data, and stopped when further simulations no longer added information. Active learning deals with exactly the same scenario, only that the computational model refers to a user whom the model queries to obtain more information. Similar to UQ, this helps to keep the number of samples low for cases where labelling a data point is expensive. In addition, many methods in current UQ focus on the structure of available random data to determine the most efficient points to run the model. This has two good reasons. One, the data shape is often the only information available a priori and two, exploiting the data structure can definitely be used to reduce computational costs. However, future research should place more importance on the shape of the physical model, because it strongly affects the convergence behaviour of all UQ methods. If there are high gradients, periodicity or discontinuities in its response the computational cost needed for convergence will grow significantly with the wrong technique. This is another reason, why adaptive and intelligent learning methods should be developed to complement existing UQ methods.

I hope that this work will help uncertainty quantification methods to become a more integral part of engineering designs in the future, and that future research will further unite data-based machine learning with physical simulations to lower designing costs and improve the safety in engineering.

Abbreviations

Statistical Terms

UQ Uncertainty Quantification

PDF Probability Density Function

CDF Cumulative Distribution Function

QoI Quantity of Interest

DoE Design of Experiment

DoF Degrees of Freedom

FFT Fast Fourier Transformation

FRF Frequency Response Function

MCS Monte Carlo Simulation

MCM Monte Carlo Method

PC Polynomial Chaos

PCE Polynomial Chaos Expansion

gPC generalised Polynomial Chaos

aPC arbitrary Polynomial Chaos

NIPC Non-Intrusive Polynomial Chaos

SAMBA Sparse Approximation of Moment-Based Arbitrary (Polynomial Chaos)

LAR Least Angle Regression

Engineering Terms

TET Turbine Entry Temperature

LPT Low Pressure Turbine

IPT Intermediate Pressure Turbine

HPT High Pressure Turbine

KSI Kinetic Energy Loss Coefficient

UTE Upper Trailing Edge

LTE Lower Trailing Edge

SLS Space Launch System

SRB Solid Rocket Booster

RANS Reynolds-Averaged Navier-Stokes

DNS Direct Numerical Simulation

Bibliography

[1] Milton Abramowitz and Irene A Stegun. *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*. Number 55. Courier Corporation, 1964.

[2] Sumanta Acharya and David Houston Leedom. Large eddy simulations of discrete hole film cooling with plenum inflow orientation effects. *Journal of Heat Transfer*, 135(1):11010, 2013.

[3] D. Amirante, N.J. Hills, and C.J. Barnes. A moving mesh algorithm for aero-thermo-mechanical modeling in turbomachinery. *Wiley International Journal for Numerical Methods in Fluids*, 70(9):1–14, 2012.

[4] Jannis Andreopoulos and Wolfgang Rodi. Experimental investigation of jets in a crossflow. *Journal of Fluid Mechanics*, 138:93–127, 1984.

[5] Meisam Ashraf, Sergey Oladyshkin, and Wolfgang Nowak. Geological Storage of CO2: Application, Feasibility and Efficiency of Global Sensitivity Analysis and Risk Assessment Using the Arbitrary Polynomial Chaos. *International Journal of Greenhouse Gas Control*, 19:704–719, nov 2013.

[6] R Askey and J Wilson. Some Basic Hypergeometric Polynomials That Generalize Jacobi Polynomials. *Memoirs of the American Mathematical Society*, 319, 1985.

[7] Volker Barthelmann, Erich Novak, and Klaus Ritter. High dimensional polynomial interpolation on sparse grids. *Advances in Computational Mathematics*, 12(4):273–288, 2000.

[8] Géraud Blatman and Bruno Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230(6):2345–2367, 2011.

[9] Dieter Bohn and Robert Krewinkel. The Effects of Unintentional Deviations Due to Manufacturing of Cooling Holes and Operation on Tran-

spiration Cooling Efficiency. In *ASME Turbo Expo 2006*, pages 689–699, Barcelona, 2006.

[10] Stefan Brack and Yannick Muller. Probabilistic Analysis of the Secondary Air System of a Low-Pressure Turbine. *Journal of Engineering for Gas Turbines and Power*, 137(2):22602, 2015.

[11] M Breuer, N Peller, Ch Rapp, and M Manhart. Flow over periodic hills - numerical and experimental study in a wide range of Reynolds numbers. *Computers and Fluids*, 38(2):433–457, 2009.

[12] Hans-Joachim Bungartz and Michael Griebel. Sparse grids. *Acta numerica*, 13:147–269, 2004.

[13] Ronald S Bunker. The Effects of Manufacturing Tolerances on Gas Turbine Cooling. *Journal of Turbomachinery*, 131(4):041018, jul 2009.

[14] Dan G Cacuci, Mihaela Ionescu-Bujor, and Ionel Michael Navon. *Sensitivity and uncertainty analysis*, volume 2. CRC press, 2005.

[15] M. Carnevale, F. Montomoli, A. D'Ammaro, S. Salvadori, and F. Martelli. Uncertainty Quantification: A Stochastic Method for Heat Transfer Prediction Using LES. *Journal of Turbomachinery*, 135(5):051021, jun 2013.

[16] Mauro Carnevale, Antonio D'Ammaro, Francesco Montomoli, and Simone Salvadori. Film Cooling and Shock Interaction: An Uncertainty Quantification Analysis With Transonic Flows. In *ASME Turbo Expo 2014: Turbine Technical Conference and Exposition*, pages V05BT13A001–V05BT13A001. American Society of Mechanical Engineers, 2014.

[17] Nicola Casari, Michele Pinelli, Alessio Suman, Luca di Mare, and Francesco Montomoli. An energy based fouling model for gas turbines: EBFOG. *Volume 2B: Turbomachinery*, 139(February):V02BT38A057, 2016.

[18] Vinod Chandavari and Sanjeev Palekar. Diffuser Angle Control To Avoid Flow Separation. *Technical Research and Applications*, 2(5):16–21, 2014.

[19] T. Chantrasmi, A. Doostan, and G. Iaccarino. Padé-Legendre approximants for uncertainty analysis with discontinuous response surfaces. *Journal of Computational Physics*, 228(19):7159–7180, 2009.

[20] Tonkid Chantrasmi, Gianluca Iaccarino, Habib Nasri Najm, and George Papanicolaou. *Padé-Legendre method for uncertainty quantification with fluid dynamics applications*. Stanford University, 2011.

[21] Yen-Chi Chen. A tutorial on kernel kensity estimation and recent advances. *arXiv preprint arXiv:1704.03924*, 2017.

[22] G Chiandussi, M Codegone, S Ferrero, and F E Varesio. Comparison of multi-objective optimization methodologies for engineering applications. *Computers & Mathematics with Applications*, 63(5):912–942, 2012.

[23] Jim S Chow, Gregory G Zilliac, and Peter Bradshaw. Mean and Turbulence Measurements in the Near Field of a Wingtip Vortex. *AIAA Journal*, 35(10):1561–1567, oct 1997.

[24] P.M. Congedo, G Iaccarino, and G. Geraci. On the Use of Higher Order Statistics in Robust Design Optimization. In *European Conference on Computational Fluid Dynamics*, pages 1–12, 2014.

[25] P Conrad and Y Marzouk. Adaptive smolyak pseudospectral approximations. *SIAM Journal on Scientific Computing*, 35(6):2643–2670, jan 2013.

[26] Paul G. Constantine, Michael S. Eldred, and Eric T. Phipps. Sparse pseudospectral approximation method. *Computer Methods in Applied Mechanics and Engineering*, 229-232(2012):1–12, jul 2012.

[27] M G Cox and P M Harris. Best Practice Guide No. 6 Uncertainty Evaluation, 2006.

[28] Thierry Crestaux, Olivier Le Maitre, and Jean-Marc Martinez. Polynomial chaos expansion for sensitivity analysis. *Reliability Engineering & System Safety*, 94(7):1161–1172, 2009.

[29] Teryn DalBello, Vance Dippold, and Nicholas J Georgiadis. *Computational study of separating flow in a planar subsonic diffuser*. National Aeronautics and Space Administration, Glenn Research Center, 2005.

[30] John D Denton. The calculation of three-dimensional viscous flow through multistage turbomachines. *ASME J. Turbomach*, 114(1):18–26, 1992.

[31] John D Denton. TBLOCK Manual. Technical report, Turbostream Ltd, 2002.

[32] David L Donoho. For most large underdetermined systems of linear equations the minimal l1-norm solution is also the sparsest solution. *Communications on pure and applied mathematics*, 59(6):797–829, 2006.

[33] Alireza Doostan, Gianluca Geraci, and Gianluca Iaccarino. A bi-fidelity approach for uncertainty quantification of heat transfer in a rectangular ribbed channel. In *ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition*, pages V02CT45A031—-V02CT45A031. American Society of Mechanical Engineers, 2016.

[34] E Dow and Q Wang. Uncertainty quantification of structural uncertainties in rans simulations of complex flows. In *20th AIAA Computational Fluid Dynamics Conference*, pages 1–16, 2011.

[35] Eric a. Dow and Qiqi Wang. The Implications of Tolerance Optimization on Compressor Blade Design. *Journal of Turbomachinery*, 137(10):1–7, 2015.

[36] B Dubrulle and J.-P. Laval. Truncated Levy Laws and 2D Turbulence. *The European Physical Journal B - Condensed Matter and Complex Systems*, 4(2):143–146, 1998.

[37] Earl H. Dudgeon. Guide to the Measurement of the Transient Performance of Aircraft Turbine Engines and Components (Report No. AR-320). Technical report, Advisory Group for Aerospace Research and Development, Neuilly-sur-Siene, France, 1994.

[38] John D Duffner. *The effects of manufacturing variability on turbine vane performance*. PhD thesis, Massachusetts Institute of Technology, 2008.

[39] W. N. Edeling, R. P. Dwight, and P. Cinnella. Simplex-stochastic collocation method with improved scalability. *Journal of Computational Physics*, 310:301–328, 2016.

[40] M S Eldred and J Burkardt. Comparison of Non-Intrusive Polynomial Chaos and Stochastic Collocation Methods for Uncertainty Quantification. In *Proceedings of the 47th AIAA Aerospace Sciences Meeting*, pages 1–20, 2009.

[41] Michael Eldred. Recent Advances in Non-Intrusive Polynomial Chaos and Stochastic Collocation Methods for Uncertainty Analysis and Design. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural*

Dynamics, and Materials Conference, number May, pages 1 – 37, Reston, Virigina, may 2009. American Institute of Aeronautics and Astronautics.

[42] Oliver G Ernst, Antje Mugler, Hans-Jörg Starkloff, and Elisabeth Ullmann. On the Convergence of Generalized Polynomial Chaos Expansions. *Mathematical Modelling and Numerical Analysis*, 46(02):317–339, 2012.

[43] Michael Havbro Faber. Risk and safety in civil engineering. *Lecture Notes. Swiss Federal Institute of Technology, Zurich*, 2007.

[44] W Feller. *An introduction to probability theory and Its applications*. Russian Academy of Sciences, Branch of Mathematical Sciences, 1968.

[45] Nicolas Fischer, Alexandre Allard, B Markus, Elena Barton, and Johan Bunde. Best practice guide to uncertainty evaluation for computationaly expensive methods. Technical report, 2015.

[46] Vijay K Garg and David L Rigby. Heat transfer on a film-cooled blade–effect of hole physics. *International Journal of Heat and Fluid Flow*, 20(1):10–25, 1999.

[47] Victor E Garzon. *Probabilistic aerothermal design of compressor airfoils*. PhD thesis, Massachusetts Institute of Technology, 2003.

[48] Walter Gautschi. Algorithm 726: ORTHPOL- a package of routines for generating orthogonal polynomials and Gauss-type quadrature rules. *ACM Transactions on Mathematical Software (TOMS)*, 20(1):21–62, 1994.

[49] Michael B Giles. Multilevel monte carlo path simulation. *Operations Research*, 56(3):607–617, 2008.

[50] Edward H Glaessgen and David Stargel. The Digital Twin paradigm for future NASA and US Air Force vehicles. In *53rd Struct. Dyn. Mater. Conf. Special Session: Digital Twin, Honolulu, HI, US*, pages 1–14, 2012.

[51] Gene H. Golub and John H. Welsch. Calculation of Gauss Quadrature Rules. *Mathematics of Computation*, 23(106):221–221, 1968.

[52] Naveen Prasad Gopinathrao, Christophe Mabilat, and Sohail Alizadeh. Non-deterministic thermo-fluid analysis of a compressor rotor-stator cavity. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dy-*

namics, and Materials Conference 17th AIAA/ASME/AHS Adaptive Structures Conference 11th AIAA No, page 2278, 2009.

[53] Alex Gorodetsky and Youssef Marzouk. Efficient Localization of Discontinuities in Complex Computational Simulations. *SIAM Journal on Scientific Computing*, 36(6):A2584–A2610, 2014.

[54] Henning Grafe. *Model updating of large structural dynamics models using measured response functions*. PhD thesis, University of London, 1999.

[55] Philippe Guillaume, Alain Huard, and Vincent Robin. Generalized multivariate Padé approximants. *Journal of Approximation Theory*, 95(2):203–214, 1998.

[56] J P Hagen and M Kurosaka. Corewise cross-flow transport in hairpin vortices - The tornado effect'. *Physics of Fluids A: Fluid Dynamics (1989-1993)*, 5(12):3167–3174, 1993.

[57] F. Hayot and L. Wagner. Robustness of Vortex Streets. *Physical Review E*, 49(1):470–473, jan 1994.

[58] A. Hellsten. Description of the Computations for Periodic Flow over a 2D Hill (Testcase 9.2). In *9th ERCOFTAC/IAHR workshop on refined turbulence modelling*, pages 1–5, 2001.

[59] L. G. Horta, M. C. Reaves, R. D. Buehrle, J. D. Templeton, D. R. Lazor, J. L. Gaspar, R. A. Parks, and P. A. Bartolotta. Finite Element Model Calibration for Ares I-X Flight Vehicle. *Experimental Mechanics*, 51(8):1251–1263, 2011.

[60] Serhat Hosder and Robert W Walters. Non-intrusive polynomial chaos methods for uncertainty quantification in fluid dynamics. In *48th AIAA Aerospace Sciences Meeting*, pages 1–16, 2010.

[61] Serhat Hosder, Robert W Walters, and Michael Balch. Efficient sampling for non-intrusive polynomial chaos applications with multiple uncertain input variables. *Collection of Technical Papers - AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 3(April):2946–2961, 2007.

[62] Gianluca Iaccarino. Predictions of a turbulent separated flow using commercial CFD codes. *Journal of Fluids Engineering*, 123(4):819–828, 2001.

[63] Kenneth L. Judd, Lilia Maliar, Serguei Maliar, and Rafael Valero. Smolyak Method for Solving Dynamic Economic Models: Lagrange Interpolation, Anisotropic Grid and Adaptive Domain. *Journal of Economic Dynamics and Control*, 44:92–123, jul 2014.

[64] Ismo Koponen. Analytical Approach to the Problem of Convergence of Truncated Levy Flights Towards the Gaussian Stochastic Process. *Physical Review E*, 52(1):1197–1199, 1995.

[65] Sotiris B Kotsiantis, I Zaharakis, and P Pintelas. Supervised machine learning: A review of classification techniques, 2007.

[66] Sylvain Laizet and Eric Lamballais. High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics*, 228(16):5989–6015, 2009.

[67] Sylvain Laizet and Ning Li. Incompact3d: A powerful tool to tackle turbulence problems with up to O (105) computational cores. *International Journal for Numerical Methods in Fluids*, 67(11):1735–1757, 2011.

[68] B E Launder and B I Sharma. Application of the energy-dissipation model of turbulence to the calculation of flow near a spinning disc. *Letters in heat and mass transfer*, 1(2):131–137, 1974.

[69] O P Le Maitre, O M Knio, H N Najm, and R G Ghanem. Uncertainty propagation using Wiener–Haar expansions. *Journal of computational Physics*, 197(1):28–57, 2004.

[70] Olivier Le Maitre and Omar M Knio. *Spectral methods for uncertainty quantification: with applications to computational fluid dynamics*. Springer Science & Business Media, 2010.

[71] Francis Leboeuf and Olivier Sgarzi. The detailed structure and behavior of discrete cooling jets in a turbine. *Annals of the New York Academy of Sciences*, 934(1):95–109, 2001.

[72] S H Lee, W Chen, and B M Kwak. Robust design with arbitrary distributions using Gauss-type quadrature formula. *Structural and Multidisciplinary Optimization*, 39(3):227–243, 2009.

[73] Sanjiva K Lele. Compact finite difference schemes with spectral-like resolution. *Journal of computational physics*, 103(1):16–42, 1992.

[74] G Lin and A M Tartakovsky. An efficient, high-order multi-element probabilistic collocation method on sparse grids for three-dimensional flow in random porous media. In *AGU Fall Meeting Abstracts*, volume 1, page 1318, 2007.

[75] Alex Loeven and Hester Bijl. An efficient framework for uncertainty quantification in CFD using probabilistic collocation. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, number May, pages 1–14, 2009.

[76] G. J. A. Loeven. *Efficient uncertainty quantification in computational fluid dynamics*. PhD thesis, Delft University of Technology, 2010.

[77] G J A Loeven and Hester Bijl. The application of the probabilistic collocation method to a transonic axial flow compressor. In *51st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2010.

[78] G.J.A. Loeven, J.A.S. Witteven, and H. Bijl. Airfoil analysis with uncertain heometry using the probabilistic collocation method. In *AIAA 2008*, pages 1–11, 2008.

[79] Jean-Eloi W Lombard, David Moxey, Spencer J Sherwin, Julien F A Hoessler, Sridar Dhandapani, and Mark J Taylor. Implicit Large-Eddy Simulation of a Wingtip Vortex. *AIAA Journal*, pages 1–13, nov 2015.

[80] Didier Lucor, Jeroen a.S. Witteveen, P Constantine, D Schiavazzi, and Gianluca Iaccarino. Comparison of adaptive uncertainty quantification approaches for shock wave-dominated flows. *Center For Turbulence Research, Prooceedings of the Summer Program 2012*, pages 219–228, 2012.

[81] Didier Lucor, Chau-Hsing Su, and George Em Karniadakis. Generalized polynomial chaos and random oscillators. *International Journal for Numerical Methods in Engineering*, 60(3):571–596, 2004.

[82] R N Mantegna and H Eugene Stanley. Stochastic Process with Ultrasound Convergence to a Gaussian Truncated Levy Flight. *Physical Review Letters*, 73(22):2946, 1994.

[83] M Marquillie, J.-P. Laval, and R Dolganov. Direct numerical simulation of a separated channel flow with a smooth profile. *Journal of Turbulence*, 9:1–24, 2008.

[84] Cosimo Mazzoni, Richard Ahlfeld, Budimir Rosic, and Francesco Montomoli. Uncertainty Quantification of Leakages in a Multistage Simulation and Comparison with Experiments. *Journal of Fluids Engineering*, 140(2), sep 2017.

[85] C. B. Meher-Homji and G. Gabriles. Gas Turbines Blades Failures: Causes, Avoidance and Troubleshooting. In *27th Turbomachinery Symposium*, pages 129–180, Houston, Texas, 1998.

[86] F R Menter, M Kuntz, and R Langtry. Ten years of industrial experience with the SST turbulence model. *Turbulence, heat and mass transfer*, 4(1):625–632, 2003.

[87] F R Menter, R Langtry, and S Völker. Transition modelling for general purpose CFD codes. *Flow, turbulence and combustion*, 77(1-4):277–303, 2006.

[88] J. Michael Owen. Theoretical Modelling of Hot Gas Ingestion Through Turbine Rim Seals. *Propulsion and Power Research*, 1(1):1–11, dec 2012.

[89] Cheon-Hong Min, Sup Hong, Soo-Yong Park, and Dong-Cheon Park. Sensitivity-based finite element model updating with natural frequencies and zero frequencies for damped beam structures. *International Journal of Naval Architecture and Ocean Engineering*, 6(4):904–921, 2014.

[90] M Montis, R Ciorciari, S Salvadori, M Carnevale, and R Niehuis. Numerical prediction of cooling losses in a high-pressure gas turbine airfoil. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, page 0957650914542630, 2014.

[91] F. Montomoli, D. Amirante, N. Hills, S. Shahpar, and M. Massini. Uncertainty Quantification, Rare Events, and Mission Optimization: Stochastic Variations of Metal Temperature During a Transient. *Journal of Engineering for Gas Turbines and Power*, 137:42–101, 2014.

[92] F Montomoli, A Ammaro, and S Uchida. Uncertainty quantification and conjugate heat transfer: A stochastic analysis. *Journal of Turbomachinery*, 135(3):31014, 2013.

[93] Francesco Montomoli, Mauro Carnevale, Antonio Ammaro, Michela Massini, and Simone Salvadori. *Uncertainty Quantification in Computational Fluid Dynamics and Aircraft Engines*. Springer, Berlin, 2015.

[94] Francesco Montomoli, Howard Hodson, and Frank Haselbach. Effect of Roughness and Unsteadiness on the Performance of a New Low Pressure Turbine Blade at Low Reynolds Numbers. *Journal of Turbomachinery*, 132(3):31018, 2010.

[95] Francesco Montomoli and Michela Masini. Gas Turbine and Uncertainty Quantification: Impact of PDF Tails on UQ Predictions, The Black Swan. In *Proceedings of ASME Conference GT2013*, pages 1–9, 2013.

[96] Francesco Montomoli, Michela Massini, Simone Salvadori, and Francesco Martelli. Geometrical Uncertainty and Film Cooling: Fillet Radii. *Journal of Turbomachinery*, 134(1):011019, 2012.

[97] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.

[98] I P Mysovskikh. On the Construction of Cubature Formulas with Fewest Nodes. *Doklady Akademii Nauk SSSR*, 178(6):1252–1254, 1968.

[99] Habib N Najm. Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics. *Annual review of fluid mechanics*, 41:35–52, 2009.

[100] Hiroya Nakao. Multi-Scaling Properties of Truncated Levy Flights. *Physics Letters, Section A: General, Atomic and Solid State Physics*, 266:282–289, 2000.

[101] M Navarro, J Witteveen, and J Blom. Polynomial Chaos Expansion for general multivariate distributions with correlated variables. *ArXiv e-prints, http://arxiv.org/pdf/1406.5483.pdf*, pages 1 – 24, jun 2014.

[102] L. W. T. Ng and M. S. Eldred. Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, pages 1–17, 2012.

[103] Leo W T Ng and Karen E Willcox. Multifidelity approaches for optimization under uncertainty. *International Journal for Numerical Methods in Engineering*, 100(10):746–772, 2014.

[104] William L Oberkampf, Jon C Helton, Cliff A Joslyn, Steven F Wojtkiewicz, and Scott Ferson. Challenge problems: uncertainty in system response given uncertain parameters. *Reliability Engineering & System Safety*, 85(1):11–19, 2004.

[105] M Ochs, A Schulz, and H-J Bauer. Investigation of the influence of trailing edge shock waves on film cooling performance of gas turbine airfoils. In *ASME Turbo Expo 2007: Power for Land, Sea, and Air*, pages 465–474. American Society of Mechanical Engineers, 2007.

[106] S Oladyshkin, H Class, R Helmig, and W Nowak. A concept for data-driven uncertainty quantification and its application to carbon dioxide storage in geological formations. *Advances in Water Resources*, 34(11):1508–1518, nov 2011.

[107] S. Oladyshkin and W. Nowak. Data-Driven Uncertainty Quantification Using the Arbitrary Polynomial Chaos Expansion. *Reliability Engineering and System Safety*, 106:179–190, oct 2012.

[108] Sergey Oladyshkin, Holger Class, Rainer Helmig, and Wolfgang Nowak. An integrative approach to robust design and probabilistic risk assessment for CO2 storage in geological formations. *Computational Geosciences*, 15(3):565–577, 2011.

[109] Sergey Oladyshkin, Patrick Schröder, Holger Class, and Wolfgang Nowak. Chaos Expansion based Bootstrap Filter to Calibrate CO2 Injection Models. *Energy Procedia*, 40:398–407, 2013.

[110] G Onorato, G J A Loeven, G Ghorbaniasl, H Bijl, and C Lacor. Comparison of intrusive and non-intrusive polynomial chaos methods for CFD applications in aeronautics. In *V European Conference on Computational Fluid Dynamics ECCOMAS, Lisbon, Portugal*, pages 14–17, 2010.

[111] Pramudita Satria Palar, Takeshi Tsuchiya, and Geoffrey Thomas Parks. Multi-fidelity non-intrusive polynomial chaos based on regression. *Computer Methods in Applied Mechanics and Engineering*, 305:579–606, 2016.

[112] Andrea Panizza, Alessio Bonini, and Luca Innocenti. Uncertainty Quantification of Hot Gas Ingestion for a Gas Turbine Using Polynomial Chaos. In *ASME Turbo Expo 2015*, pages 1–12, 2015.

[113] Rene Pecnik, Jeroen Witteveen, and Gianluca Iaccarino. Uncertainty Quantification for Laminar-Turbulent Transition Prediction in RANS Turbomachinery Applications. In *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, Aerospace Sciences Meetings. American Institute of Aeronautics and Astronautics, jan 2011.

[114] Giovanni Petrone, C D Nicola, Domenico Quagliarella, Jeroen Witteveen, and Gianluca Iaccarino. Wind turbine performance analysis under uncertainty. *AIAA Paper No. AIAA*, 544, 2011.

[115] V F Pisarenko, Mikhail Vladimirovich Rodkin, and Mikhail Vladimirovich Rodkin. *Heavy-Tailed Distributions in Disaster Analysis*. Springer Science and Business Media, Berlin, 2010.

[116] Stephen B Pope. Turbulent flows, 2001.

[117] P Prempraneerach, Franz S Hover, Michael S Triantafyllou, and George E Karniadakis. Uncertainty quantification in simulations of power systems: Multi-element polynomial chaos methods. *Reliability Engineering and System Safety*, 95(6):632–646, 2010.

[118] P Prinos and D Sofialidis. Test Case 9.2: Periodic Flow over a 2D Hill. In *9th ERCOFTAC/IAHR workshop on refined turbulence modelling*, pages 1–3, 2001.

[119] Ruihong Qin, Yaping Ju, Yazhou Wang, and Chuhua Zhang. Flow Analysis and Uncertainty Quantification of a 2D Compressor Cascade With Dirty Blades. In *ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition*, pages V02CT45A017—-V02CT45A017. American Society of Mechanical Engineers, 2016.

[120] Hassan Raiesi, Ugo Piomelli, and Andrew Pollard. Evaluation of turbulence models using direct numerical and large-eddy simulation data. *Journal of Fluids Engineering*, 133(2):21203, 2011.

[121] Budimir Rosic. *Control of shroud leakage flows in low aspect ratio multistage turbines*. PhD thesis, University of Cambridge, 2006.

[122] Budimir Rosic, John D Denton, and Eric M Curtis. The influence of shroud and cavity geometry on turbine performance: an experimental and computational study - part I: shroud geometry. *Journal of Turbomachinery*, 130(4):41001, 2008.

[123] Budimir Rosic, John D Denton, Eric M Curtis, and Ashley T Peterson. The Influence of Shroud and Cavity Geometry on Turbine Performance: An Experimental and Computational Study - Part II: Exit Cavity Geometry. *Journal of Turbomachinery*, 130(4):41002, 2008.

[124] H Rutishauser. On a Modification of the QD-Algorithm with Graeffe-Type Convergence. In *Zeitschrift für angewandte Mathematik und Physik*, volume 13, pages 493–496, 1963.

[125] Cosmin Safta, Bert J Debusschere, Habib N Najm, and Khachik Sargsyan. Advanced methods for uncertainty quantification in tail regions of climate model predictions. Technical report, Sandia National Laboratories, 2010.

[126] S Salvadori, F Montomoli, F Martelli, P Adami, K S Chana, and L Castillon. Aerothermal study of the unsteady flow field in a transonic gas turbine with inlet temperature distortions. *Journal of Turbomachinery*, 133(3):31030, 2011.

[127] Simone Salvadori, Francesco Montomoli, and Francesco Martelli. Film-cooling performance in supersonic flows: effect of shock impingement. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, pages 295–305, 2013.

[128] Khachik Sargsyan, Cosmin Safta, Bert Debusschere, and Habib Najm. Uncertainty Quantification given discontinuous model response and a limited number of model response evaluations. *SIAM Journal on Scientific Computing*, 34(1):44–64, 2012.

[129] Sunetra Sarkar and Jeroen A S Witteveen. *Uncertainty Quantification in Computational Science: Theory and Application in Fluids and Structural Mechanics*. World Scientific, 2016.

[130] Daniele Schiavazzi, Alireza Doostan, and Gianluca Iaccarino. Sparse Multiresolution Regression for Uncertainty Propagation. *International Journal for Uncertainty Quantification*, 4(4):303–331, 2014.

[131] Rainer Schnell, T Lengyel-Kampmann, and Eberhard Nicke. On the impact of geometric variability on fan aerodynamic performance, unsteady blade row interaction, and its mechanical characteristics. *Journal of Turbomachinery*, 136(9):91005, 2014.

[132] Roland Schöbi and Bruno Sudret. Uncertainty propagation of p-boxes using sparse polynomial chaos expansions. *Journal of Computational Physics*, 339:307–327, 2017.

[133] Hyun Seok Seo and Byung Man Kwak. Efficient statistical tolerance analysis for general distributions using three-point information. *International Journal of Production Research*, 40(4):931–944, 2002.

[134] Pranay Seshadri, Akil Narayan, and Sankaran Mahadevan. Optimal Quadrature Subsampling for Least Squares Polynomial Approximations. *arXiv preprint arXiv:1601.05470*, 2016.

[135] Sergey A Smolyak. Quadrature and Interpolation Formulas for Tensor Products of Certain Classes of Functions. *Doklady Akademii Nauk SSSR*, 4(240-243):123, 1963.

[136] Christian Soize. *Uncertainty Quantification: An Accelerated Course with Advanced Applications in Computational Engineering*, volume 47. Springer, 2017.

[137] Christian Soize and Roger Ghanem. Physical Systems with Random Uncertainties: Chaos Representations with Arbitrary Probability Measure. *SIAM Journal on Scientific Computing*, 26(2):395–410, 2004.

[138] Bruno Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964–979, 2008.

[139] Bruno Sudret, S Marelli, and C Lataniotis. Sparse Polynomial Chaos Expansions as a Machine Learning Regression Technique. In *International Symposium on Big Data and Predictive Computational Modeling*, 2015.

[140] N. N. Taleb. *The Black Swan: The Impact of the Highly Improbable*. Random House and Penguin, New York, USA, 2007.

[141] Nikita Thakur, Andy Keane, and Prasanth B Nair. Capture of manufacturing uncertainty in turbine blades through probabilistic techniques. In *Association for Structural and Multidisciplinary Optimization in the UK*, pages 1–10, 2008.

[142] J.-J. Tseng, M.-J. Lee, and S.-P. Li. Heavy-Tailed Distributions in Fatal Traffic Accidents:. Role of Human Activities. *International Journal of Modern Physics C*, 20:1281–1290, 2009.

[143] Xiaoliang Wan and George Em Karniadakis. An adaptive multi-element generalized polynomial chaos method for stochastic differential equations. *Journal of Computational Physics*, 209(2):617–642, 2005.

[144] Xiaoliang Wan and George Em Karniadakis. Beyond Wiener-Askey Expansions: Handling Arbitrary PDFs. *Journal of Scientific Computing*, 27(1-3):455–465, 2006.

[145] Xiaoliang Wan and George Em Karniadakis. Multi-element generalized polynomial chaos for arbitrary probability measures. *SIAM Journal on Scientific Computing*, 28(3):901–928, 2006.

[146] Xiaoliang Wan and George Em Karniadakis. Error control in multi-element generalized polynomial chaos method for elliptic problems with random coefficients. *Communications in Computational physics*, 5(2-4):793–820, 2009.

[147] Jian-Xun Wang, Jin-Long Wu, and Heng Xiao. Physics-Informed Machine Learning for Predictive Turbulence Modeling: Using Data to Improve RANS Modeled Reynolds Stresses. *arXiv preprint arXiv:1606.07987*, 2016.

[148] G.W. Wasilkowski and H. Wozniakowski. Explicit Cost Bounds of Algorithms for Multivariate Tensor Product Problems. *Journal of Complexity*, 11(1):1–56, mar 1995.

[149] Martin Weitzman. Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change. *Review of Environmental Economics and Policy*, 5:275–292, 2011.

[150] D C Wilcox. *Turbulence Modelling for CFD*. DCW Industries, Palm, 1993.

[151] David C Wilcox. Formulation of the kw turbulence model revisited. *AIAA journal*, 46(11):2823–2838, 2008.

[152] Jeroen A S Witteveen and Hester Bijl. Modeling Arbitrary Uncertainties Using Gram-Schmidt Polynomial Chaos. In *44th AIAA Aerospace Sciences Meeting and Exhibit*, Aerospace Sciences Meetings, pages 1–17. American Institute of Aeronautics and Astronautics, jan 2006.

[153] Jeroen A S Witteveen and Gianluca Iaccarino. Simplex stochastic collocation with random sampling and extrapolation for nonhypercube probability spaces. *SIAM Journal on Scientific Computing*, 34(2):A814—-A838, 2012.

[154] Jeroen a S Witteveen, Sunetra Sarkar, and Hester Bijl. Modeling Physical Uncertainties in Dynamic Stall Induced Fluid-Structure Interaction of Turbine Blades Using Arbitrary Polynomial Chaos. *Computers and Structures*, 85(11-14):866–878, 2007.

[155] J L Wu, J X Wang, and H Xiao. A Bayesian calibration-prediction method for reducing model-form uncertainties with application in RANS simulations, submitted, 2015.

[156] Jin-Long Wu, Jian-Xun Wang, Heng Xiao, and Julia Ling. Physics-informed machine learning for predictive turbulence modeling: A priori assessment of prediction confidence. *arXiv preprint arXiv:1607.04563*, 2016.

[157] Heng Xiao, Jian-Xun Wang, and Roger G Ghanem. A random matrix approach for quantifying model-form uncertainties in turbulence modeling. *Computer Methods in Applied Mechanics and Engineering*, 313:941–965, 2017.

[158] Dongbin Xiu and George Em Karniadakis. The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations. *SIAM journal on scientific computing*, 24(2):619–644, 2002.

[159] Dongbin Xiu and George Em Karniadakis. Modeling uncertainty in flow simulations via generalized polynomial chaos. *Journal of Computational Physics*, 187(1):137–167, 2003.

[160] Chao Zhou. Thermal Performance of Transonic Cooled Tips in a Turbine Cascade. *Journal of Propulsion and Power*, 31(5):1268–1280, 2015.