



UNIVERSITÀ DEGLI STUDI DI TRIESTE

Dipartimento di Ingegneria e Architettura

XXXIII Ciclo del Dottorato di Ricerca in

Ingegneria Industriale e dell'Informazione

**MULTI-FIDELITY SURROGATE-ASSISTED
DESIGN OPTIMISATION UNDER
UNCERTAINTY FOR COMPUTATIONALLY
EXPENSIVE AEROSPACE APPLICATIONS**

In partial fulfilment of the requirements for the degree of
Doctor of Philosophy

in

Industrial Engineering and Information Technology

Author:

Péter Zénó Korondi

Supervisor:

Prof. Carlo Poloni

Co-Supervisor:

Prof. Lucia Parussini

Co-Supervisor:

Dr. Mariapia Marchi

ANNO ACCADEMICO 2019-2020

Preface

I decided to pursue a PhD mostly for the sake of curiosity and challenge. I was always curious what lies beyond my current understanding and interpretation of the world. The exploration of new perspectives and knowledge have always allured me. The journey of learning is a time demanding challenge. Certainly, the path to a PhD is full of obstacles of various nature. During this path, I had to face the obstacles and challenges alone many times. Nevertheless, I also faced challenges which I could not have overcome without the support of my supervisors, colleagues, friends and my family.

With these few lines, I would like to express my gratitude to my supervisors Mariapia Marchi, Lucia Parussini and Carlo Poloni who were ready to provide guidance at any time. I had the pleasure to conduct my research as a member of the UTOPIAE network¹. I wish to thank to each member of the network for the great team we have become over the past 3 years. Particularly, I was hosted by ESTECO and its family. Thank you for being great colleagues and friends. I also would like to thank to Gregor Papa and Domenico Quagliarella alongside the colleagues of Jožef Stefan Institute (JSI) and Italian Aerospace Research Centre (CIRA) for the kind hospitality and collaboration.

Last but not least, I would like to address my utmost gratitude to my family and friends who supported me along this path.

PÉTER ZÉNÓ KORONDI
Trieste
November 2020

¹This work was partially supported by the H2020-MSCA-ITN-2016 UTOPIAE, grant agreement 722734.

Abstract

Virtual design analysis has become an indispensable component in most engineering disciplines. Despite the immense developments and availability of computational resources, the relative computational cost of high-fidelity simulations is getting more and more expensive. This opened the chapter of multi-fidelity learning techniques in the field of automated design optimisation.

This work presents a novel multi-fidelity surrogate-assisted design optimisation approach for computationally expensive aerospace applications under uncertainty.

The proposed optimisation framework overcomes the challenges of probabilistic design optimisation of computationally expensive problems and is capable of finding designs with optimal statistical performance for both single- and multi-objective problems, as well as constrained problems. Our approach performs the design optimisation with a limited computational budget using integrated multi-fidelity surrogates for design exploration and uncertainty quantification. The design optimisation is realised following the principles of Bayesian optimisation. The acquisition function balances exploration and exploitation of the design space and allocates the available budget efficiently considering the cost and accuracy of the fidelity levels.

To validate the proposed optimisation framework, available multi-fidelity test functions were tailored for benchmarking problems under uncertainty. The benchmarks showed that it is profitable to use multi-fidelity surrogates when the computational budget is too limited to allow for the construction of an accurate surrogate with high-fidelity simulations but is large enough to generate a great number of low-fidelity data.

The applicability of the proposed optimisation framework for aerospace applications is presented through optimisation studies of a propeller blade airfoil and a 3D propeller blade.

Contents

List of Figures	ix
List of Tables	xiii
List of Acronyms	xvii
List of Symbols	xix
1 Introduction	1
1.1 Motivation	1
1.2 Automated design optimisation	4
1.3 Uncertainties in design optimisation	9
2 Multi-fidelity information fusion	11
2.1 Information fusion	11
2.2 Gaussian process regression	12
2.3 Parameter estimation of Gaussian process regression	14
2.4 Multi-fidelity Gaussian process regression	15
2.5 Parameter estimation of multi-fidelity Gaussian process regression	17
3 Uncertainty quantification	19
3.1 Statistical measures	19
3.2 Surrogate-assisted uncertainty propagation	21
3.3 An academic example for surrogate-assisted uncer- tainty propagation	21
3.4 Implications for design optimisation of real-world systems	25

4	Optimisation framework	27
4.1	Surrogate-assisted design optimisation	27
4.2	Proposed single-objective framework	30
4.3	Single-objective acquisition function	33
4.4	Proposed multi-objective framework	35
4.5	Multi-objective acquisition function	38
5	Multi-fidelity benchmark functions	41
5.1	Multi-fidelity benchmark functions	41
5.2	Multi-fidelity test problems under uncertainty . . .	43
5.3	Comparison of single- and multi-fidelity surrogates	45
5.4	Acquisition step of a multi-fidelity surrogate-based optimisation	48
5.5	Benchmarking the single-objective multi-fidelity surrogate-based optimisation under uncertainty . .	51
5.6	Benchmarking the multi-objective multi-fidelity surrogate- based optimisation under uncertainty	53
6	Airfoil shape optimisation under uncertainty	57
6.1	Introduction to airfoil design	57
6.2	Airfoil shape optimisation: uncertainties, reliability- based risk measure and other details	59
6.3	Solvers for aerodynamic force prediction of an airfoil	60
6.3.1	Parametric geometry definition of airfoils . .	62
6.3.2	Mesh generation of airfoils	63
6.3.3	Far-field drag prediction of coarse grids . .	65
6.4	Multi-objective airfoil optimisation under uncer- tainty with XFOIL and SU2	67
6.4.1	Workflow for multi-objective airfoil design optimisation	68
6.4.2	Results of multi-objective airfoil design op- timisation	69
6.4.3	Conclusion of multi-objective airfoil design optimisation under uncertainty	77
6.5	Constrained single-objective optimisation of an air- foil under uncertainty	80
6.5.1	Constrained single-objective optimisation workflow for airfoil design under uncertainty	81
6.5.2	Results	86

6.5.3	Conclusion of the constrained single-objective airfoil design optimisation	93
7	Shape optimisation of a ducted propeller blade	95
7.1	Introduction to propeller blade optimisation	95
7.2	Uncertainty quantification with local probabilistic model	96
7.3	Aerodynamic performance prediction of a propeller	100
7.3.1	Parametric geometry definition of propeller blades	100
7.3.2	Low-fidelity predictor: Blade element mo- mentum theory	103
7.3.3	High-fidelity predictor: Reynolds Averaged Navier-Stokes solver	106
7.3.3.1	Ducted propeller geometry	107
7.3.3.2	Parametric mesh generation . . .	108
7.3.3.3	Boundary conditions	112
7.4	Multi-objective optimisation under uncertainty with XROTOR and SU2	113
7.4.1	Computational chain for multi-objective pro- peller blade design optimisation under un- certainty	114
7.4.2	Results of multi-objective propeller optimi- sation under uncertainty	116
7.4.3	Conclusion of multi-objective propeller blade optimisation study	120
7.5	Constrained single-objective optimisation under un- certainty with XROTOR and SU2	123
7.5.1	Computational chain for constrained single- objective propeller blade design optimisa- tion under uncertainty	123
7.5.2	Results of constrained single-objective pro- peller optimisation under uncertainty . . .	125
7.5.3	Conclusion of the constrained single-objective propeller blade optimisation study	130
8	Conclusions	131
8.1	Conclusions	131
8.2	Possible developments	133

A	Appendix	135
A.1	An alternative Gaussian process regression formulation	135
A.1.1	Gaussian Markov random field	135
A.1.2	Gaussian Markov Random Field Regression	138
A.1.3	Multi-fidelity Gaussian Markov random field regression	141
A.1.4	Comparison of multi-fidelity Gaussian process regression and multi-fidelity Gaussian Markov random field regression	143
A.2	Space systems resilience optimisation under episodic uncertainty	145
	Bibliography	149

List of Figures

3.1	Comparison of Monte Carlo sampling and virtual sampling with <i>Polynomial Chaos Expansion</i> (PCE)	23
3.2	Superquantile calculation	24
4.1	Proposed single-objective surrogate-assisted design optimisation framework for expensive probabilistic problems.	31
4.2	Workflow of multi-objective multi-fidelity surrogate-based optimisation under uncertainty	37
4.3	Illustration of the <i>Hypervolume Improvement</i> (HVI) and <i>Hypervolume Improvement of the Lower Confidence Bound</i> (HVI _{LCB}) for a bi-objective problem. . .	39
5.1	Landscapes of the single- and multi-fidelity surrogate models. (In the plots for the sake of better visibility the low-fidelity curves are shifted up by 1. The true high- and low-fidelity functions are drawn with solid and dashed black lines respectively. The high- and low-fidelity model predictions are drawn with red and blue respectively.)	47
5.2	Comparison of single- and multi-fidelity GPR for single-objective deterministic case. The lines represent the absolute errors of the surrogate models calculated at 2000 locations sampled by latin hypercube sampling. The black dashed line corresponds to the GPR model; the three differently shaded red lines correspond to the three error levels of the low-fidelity model.	48
5.3	Iteration step of multi-fidelity EGO	50

5.4	Comparison of single- and multi-fidelity surrogate-assisted optimisation convergence. (The convergence histories were obtained by a trial run of TU2 with a budget of 4000.)	52
6.1	Aerodynamic forces of an airfoil	58
6.2	Comparison of lift and drag coefficient curves of MH114 calculated with XFOIL and SU2.	61
6.3	Baseline design and the deformation modes of the airfoil.	63
6.4	MH114 airfoil mesh	64
6.5	Structured meshes of different size.	66
6.6	Multi-objective probabilistic optimisation workflow for aerodynamic shape optimisation of an airfoil with <i>Multi-Fidelity Gaussian Process Regression</i> (MFGPR).	70
6.7	History of fidelity selection throughout the multi-objective design optimisation of the airfoil.	71
6.8	Comparing relative prediction errors of MFGPR and <i>Gaussian Process Regression</i> (GPR) in every iterations when a HF evaluation was performed	72
6.9	Correlation history of objectives over all iterations	73
6.10	Comparison of high- and low-fidelity Pareto fronts of the multi-objective airfoil design optimisation problem.	74
6.11	Comparison of the baseline and Pareto optimal designs with uncertainty prediction	76
6.12	Comparison of Pareto optimal design geometries based on various selection criteria.	77
6.13	Comparison of the pressure distribution of Pareto optimal designs based on various selection criteria.	79
6.14	Computational chain of the aerodynamic forces with the probabilistic model.	82
6.15	MFGPR models of the aerodynamic force coefficients of the baseline configuration (The LF and HF probabilistic models are drawn with light blue and grey lines respectively. The LF and HF samples are displayed with black and red crosses. The virtual samples are represented by blue dots.)	84

6.16	Convergence of the risk measure value of the aerodynamic force coefficients. (The figures show the averaged value of the risk measure, where 10 independently generated virtual sample set were used at each magnitude level of the virtual sample size. The shaded area represents the approximated 95% confidence interval.)	85
6.17	Prediction of the distributions for the baseline and optimal designs.	89
6.18	Comparison of optimal airfoil geometries of the constrained single-objective problem	91
6.19	Comparison of the pressure distributions of the deterministic and probabilistic optimal designs.	92
7.1	Convergence analysis of the number of low-fidelity samples of the local probabilistic model. (The figures show the averaged value of the risk measure, where 10 independently generated LF sample sets were used at each level of the low-fidelity sample size. The shaded area represents the approximated 95% confidence interval.)	98
7.2	Convergence analysis of the number of virtual samples of the local probabilistic model (The figures show the averaged value of the risk measure, where 10 independently generated virtual sample sets were used at each magnitude level of the virtual sample size. The shaded area represents the approximated 95% confidence interval.)	99
7.3	Design variables: twist and chord distributions	101
7.4	Shape modes of chord and twist distribution functions (mode 1: blue line, mode 2: orange line, mode 3: green line, mode 4: red line, mode 5: purple line)	102
7.5	Blade Element velocities and forces	103
7.6	Computational chain for calculating the propeller aerodynamic performance with XROTOR.	105
7.7	Computational chain for solving RANS equations numerically.	106
7.8	Comparison of thrust and power distributions of the baseline design calculated with XROTOR and SU2. .	107
7.9	Propeller blade geometry with double ring mounting. (The rotating parts are highlighted with turquoise colour.)	108
7.10	Outer wedge faces.	109

7.11	Near propeller faces.	109
7.12	Structured boundary layer mesh.	110
7.13	Mesh sizing	112
7.14	Computational chain of the propeller aerodynamics with reliability measure calculation for multi-objective optimisation problem	115
7.15	History of the low- and high-fidelity thrust and power coefficients, and the computational budget in the multi- objective problem. (The positive thrust coefficients are given here: $T_c^{\mathcal{S}_{95}^{left}} = -(-T_c^{\mathcal{S}_{95}^{right}})$.)	117
7.16	Pareto front of multi-objective propeller optimisation problem	118
7.17	Thrust coefficient distributions of selected Pareto opti- mal designs.	119
7.18	Power coefficient distributions of selected Pareto opti- mal designs.	119
7.19	Comparison of the sectional properties of the Pareto optimal designs.	122
7.20	Computational chain of the propeller aerodynamics with reliability measure calculation for constrained single-objective optimisation problem.	124
7.21	History of the low- and high-fidelity objectives and the computational budget in the constrained single- objective problem.	126
7.22	Comparison of the baseline and the probabilistic opti- mum	128
7.23	Comparison of the baseline and the probabilistic opti- mum of the constrained single-objective propeller blade optimisation	129
A.1	Multi-fidelity GMRF regression.	143
A.2	Comparison of the prediction error distribution of MFGPR and MF-GMRF.	145
A.3	Comparison of the margin-based and the resilience- based optimum when the transmitted data volume is at least 600 GB as published in [Paper F: Filippi et al.,2020].	147

List of Tables

1.1	Classification of optimisation algorithms based on their information usage	8
5.1	List of test cases for deterministic single-objective problems. (c_{HF} : cost of high-fidelity evaluation, c_{LF} : cost of low-fidelity evaluation)	45
5.2	Comparison of single- and multi-fidelity GPR for deterministic problems.	46
5.3	List of test cases for uncertain single-objective problems	51
5.4	Comparison of single- and multi-fidelity surrogate-assisted optimisation results. The values in the table are averaged over 10 independent runs.	53
5.5	List of test cases for uncertain multi-objective problems	54
5.6	Comparison of single- and multi-fidelity surrogate-assisted multi-objective optimisation under uncertainty. The <i>Inverted Generational Distance</i> (IGD) values of the obtained Pareto fronts are calculated with a reference set of 100 points. The values in the table are averaged over 10 independent runs.	55
6.1	Design variables of the airfoil.	62
6.2	Mesh size parameters for low- and high-fidelity simulations. (N_b : number of cells on the body surface, N_w : number of cells in the wake, N_j : number of cells in far-field direction, N_{total} : total number of cells) [Paper C: Morales and Korondi et al.,2020]	65
6.3	Geometrical uncertainties of the multi-objective airfoil design problem.	67

6.4	Summary of the multi-objective airfoil design optimisation.	69
6.5	Multi-objective airfoil design optimisation: number of LF and HF samples and their costs.	71
6.6	Mean prediction error of the lift and drag reliability measures in the multi-objective airfoil design problem.	72
6.7	List of design variables and objectives of the non-dominated airfoil designs.	75
6.8	Comparison of Pareto optimal designs based on various selection criteria.	75
6.9	Comparison of multi-fidelity, single-fidelity surrogate-based and population-based optimisation results of the constrained single-objective airfoil design problem. (The cost of a single high- and low-fidelity simulations are 10 and 1 respectively.) [Paper C: Morales and Korondi et al.,2020]	87
6.10	Comparison of the prediction error of MFGPR and GPR surrogate models of the constrained single objective airfoil design problem. (Prediction error is defined as the arithmetic mean value of the relative error of the high-fidelity predictions during the course of optimisation.)	88
6.11	Comparison of the requirements considering environmental uncertainty.	90
7.1	The uncertain environmental parameters of the propeller blade optimisation problems	96
7.2	Shape modes of chord and twist distribution functions (r_R is the relative radial location)	101
7.3	Design variables of the twist and chord distributions (r_R is the relative radial location)	102
7.4	Boundary condition definition of periodic propeller mesh.	113
7.5	Computational budget spent on multi-objective propeller design optimisation. (Total computational budget is 21240 units which are equal to the cost of 120 LF and 22 HF reliability samples.)	116
7.6	Design variables and objective values of the Pareto optimal designs.	121

7.7	Computational budget spent on constrained single-objective propeller design optimisation. (Total computational budget is 26910 units which are equal to the cost of 120 LF and 37 HF samples.)	127
7.8	Comparison of the baseline and the optimal designs of the constrained single-objective problem	128
A.1	Comparison of MFGPR and multi-fidelity GMRF regression. (The mean absolute error is calculated with 100 instances.)	144

List of Acronyms

ADT	Actuator Disk Theory
BEMT	Blade Element Momentum Theory
BET	Blade Element Theory
CAD	Computer-Aided Design
CAE	Computer-Aided Engineering
CDF	Cumulative Distribution Function
cEI	Constrained Expected Improvement
CFD	Computational Fluid Dynamics
DoE	Design of Experiments
ECDF	Empirical Cumulative Distribution Function
EGO	Efficient Global Optimisation
EI	Expected Improvement
GMRF	Gaussian Markov Random Field
GPR	Gaussian Process Regression
HF	High-Fidelity
HV	Hypervolume
HVI	Hypervolume Improvement

HVI_{LCB}	Hypervolume Improvement of the Lower Confidence Bound
IGD	Inverted Generational Distance
LE	Leading Edge
LF	Low-Fidelity
LHS	Latin Hypercube Sampling
MCS	Monte Carlo Sampling
MFGPR	Multi-Fidelity Gaussian Process Regression
MLE	Maximum Likelihood Estimation
PCE	Polynomial Chaos Expansion
PDF	Probability Density Function
QoI	Quantity of Interest
RANS	Reynolds-averaged Navier–Stokes
SEVR	Scaled Expected Variance Reduction
SHVR	Scaled Hypervariance Reduction
SPDE	Stochastic Partial Differential Equation
TE	Trailing Edge
UQ	Uncertainty Quantification

List of Symbols

C_d	Airfoil drag coefficient
C_l	Airfoil lift coefficient
c_l	Cost of fidelity level
Σ	Covariance matrix
x_i	Design variable
\mathbb{E}	Expected value
I	Identity matrix
l	Level of fidelity
\mathcal{L}	Likelihood
\hat{m}	Mean model prediction
θ	Model parameter
\hat{s}^2	Model variance
\mathcal{N}	Normal distribution
\mathcal{P}	Pareto set
P_c	Power coefficient
P	Probability
q_ξ	Quantile
\mathcal{R}	Risk measure
\mathcal{S}	Superpercentile
\bar{q}_ξ	Superquantile
T_c	Thrust coefficient
L	Top fidelity level
u_i	Uncertain variable

“Nothing in life is to be feared, it is only to be understood. Now is the time to understand more, so that we may fear less.”

— Marie Skłodowska Curie

1.1 Motivation

With the recent technological improvements, electrical aviation has become an attainable goal for our near future. To reach this goal the research community has to address various issues ranging from storing energy efficiently to redefining maintenance protocols. One important factor in accomplishing this challenge is to have more efficient propeller designs. The efficient performance must be maintained even in uncertain conditions. Propeller efficiency earns even higher significance in distributed propulsion applications where energy losses are multiplied by the number of propellers. This motivates the development of more advanced design optimisation techniques.

In the past few decades, we have witnessed the rapid development of *Computer-Aided Engineering* (CAE) technologies. Computer-based design analysis has become an indispensable component in most engineering disciplines. The development of advanced software tools furthered the automation of the engineering design optimisation process combining the automated evaluation of design configurations with the use of optimisation techniques. Despite the immense developments and availability

of computational resources, the relative computational cost of high-fidelity simulations is getting more and more expensive. This opened the chapter of multi-fidelity learning techniques in the field of automated design optimisation. The main challenge here is how to combine information stemming from various fidelity sources in a design optimisation process in such a way that the available computational resources are efficiently allocated.

Multi-fidelity learning techniques are gaining increasing popularity in the aerospace sector. The aerodynamic design optimisation of an aerospace application involves solving the Navier-Stokes partial differential equations. In the absence of any generally applicable analytical solution, the Navier-Stokes equation has to be solved numerically, which is computationally demanding. This computational cost can render standard automated design optimisation techniques impractical. The issue of expensive simulations becomes an even more relevant factor when the underlying problem is solved by taking uncertainties into account. Typically, *Uncertainty Quantification* (UQ) requires a large set of samples to calculate the risk measure of interest. This further motivates the incorporation of computationally inexpensive aerodynamic predictions into the design optimisation process.

In the scope of this dissertation, a multi-fidelity surrogate-assisted design optimisation strategy is proposed for solving computationally expensive aerospace applications.

The key findings and implications of this dissertation lies in:

- Presenting a novel multi-fidelity surrogate-assisted optimisation strategy for both single and multi-objective probabilistic optimisation problems.
- Proposing a fidelity selection strategy for multi-fidelity Bayesian optimisation.
- Providing a set of benchmark functions for multi-fidelity optimisation under uncertainty.
- Providing simplified industrial case studies of solving a propeller airfoil design optimisation problem under uncertainty with multi-fidelity learning.

- Providing simplified industrial case studies of solving a 3D propeller blade optimisation problem under uncertainty with multi-fidelity learning.

The dissertation includes the following manuscripts:

- [Paper A: Korondi et al.,2020a]:
Korondi, P. Z., Marchi, M., Parussini, L., and Poloni, C. (2020a). Multi-fidelity design optimisation strategy under uncertainty with limited computational budget. *Optimization and Engineering*
- [Paper B: Korondi et al.,2020b]:
Korondi, P. Z., Marchi, M., Parussini, L., Quagliarella, D., and Poloni, C. (2020b). Multi-objective design optimisation of an airfoil with geometrical uncertainties leveraging multi-fidelity gaussian process regression. In *UQOP:International Conference on Uncertainty Quantification & Optimisation*
- [Paper C: Morales and Korondi et al.,2020]¹:
Morales, E., Korondi, P. Z., Quagliarella, D., Tognaccini, R., Marchi, M., Parussini, L., and Poloni, C. (2020). Multi-fidelity surrogate assisted design optimisation of an airfoil under uncertainty using far-field drag approximation. In *UQOP:International Conference on Uncertainty Quantification & Optimisation*
- [Paper D: Korondi et al.,2019]:
Korondi, P. Z., Parussini, L., Marchi, M., and Poloni, C. (2019b). Reliability-based design optimisation of a ducted propeller through multi-fidelity learning. In *UNCECOMP: 3rd International Conference on Uncertainty Quantification in Computational Sciences and Engineering Proceedings*

¹Paper C is written in collaboration with Elisa Morales and her supervisors. The dissertation adopts only the author's contribution in the collaboration work. The contribution of the collaborating authors of Paper C is properly referenced.

- [Paper E: Korondi et al.,2020c]:
Korondi, P. Z., Marchi, M., and Poloni, C. (2020c). *Optimization Under Uncertainty with Applications to Aerospace Engineering: Response Surface Methodolgy (Ch.12)*. Springer
- [Paper F: Filippi et al.,2020]²:
Filippi, G., Vasile, M., Krpelik, D., Korondi, P. Z., Marchi, M., and Poloni, C. (2019). Space systems resilience optimisation under epistemic uncertainty. *Acta Astronautica*, 165:195–210

The dissertation is organised as follows. In this chapter (Chapter 1), automated design optimisation and uncertainties in engineering are introduced. In Chapter 2, the multi-fidelity information fusion is discussed by presenting the *Multi-Fidelity Gaussian Process Regression* (MFGPR) surrogate technique. Chapter 3 goes into the depth of uncertainty quantification techniques and introduces the superpercentile risk measure. A novel optimisation strategy is proposed in Chapter 4. A set of benchmark functions is introduced in Chapter 5 where also a series of properties of the proposed strategy is investigated. The proposed strategy is employed for solving airfoil design optimisation problems under uncertainty in Chapter 6. A simplified industrial propeller blade optimisation case study is detailed in Chapter 7. Finally, the dissertation is concluded in Chapter 8. The dissertation is complemented with an alternative multi-fidelity Gaussian process regression formulation in Appendix A.1 and a brief summary of a collaboration work on resilience-based optimisation of a space systems under epistemic uncertainty in Appendix A.2.

1.2 Automated design optimisation

The primary goal of engineering design optimisation is to find the best possible solution which satisfies all the underlying requirements. Full automation of the design optimisation process

²Paper F is written in collaboration with Gianluca Filippi, Daniel Krpelik and our supervisors. Paper F discusses how to treat epistemic uncertainty during the early design stages of a space system. To keep the coherent structure of the dissertation, the findings of Paper F are summarised in the Appendix.

is a long-awaited goal in engineering [Bandler, 1969]. Therefore, various techniques were developed to tackle the highly diverse problems of engineering [Papalambros, 2002]. Automated design optimisation techniques are a cleverly organised set of instructions, which represent the different rationales of how to reach the best design. As the *No Free Lunch theorem* dictates, there is no strategy which outperforms all the other strategies for every optimisation problem [Wolpert and Macready, 1997]. Nevertheless, preferable strategies have been identified over the years for solving different classes of optimisation problems.

The optimisation problems can be grouped according to the number of their requirements (or objectives). Consequently, we differentiate single- and multi-objective problems.³

Single-objective optimisation focuses on improving the design considering a single requirement (or a single objective function) which might have a single global optimal solution. However, formulating a real problem as single-objective usually implies a decision making on the preferences of various requirements to aggregate the requirements into a single objective function.

In *multi-objective* problems, the preferences of various requirements are decided after a set of Pareto optimal solutions is found [Deb, 2014]. Consequently, in the context of multi-objective design optimisation, we are searching for a set of designs which is Pareto optimal for our multi-objective problem formulation. This set is called the Pareto set and its projection into the objective space is the Pareto front. When the objectives are conflicting, the Pareto front contains more than one solution that cannot be improved in any of the objectives without degrading at least one of the others. Solutions that are optimal in respect of all objectives simultaneously are typically non-existing.

This work focuses on solving computationally expensive aerodynamic design optimisation problems. The topology is predefined and only the shape parameters are considered as variables. A possible classification of optimisation strategies is provided in Table 1.1. This classification breaks down the optimisation algorithms into four classes according to their information usage.

³Since recently, the multi-objective optimisation has started referring to problems with only 2 or 3 objectives and problems with 4 or more objectives are referred to as many-objective problems.

Global Search techniques do not exploit any relations between design parameters and response values to drive the search for optimal solutions. The design parameters are explored in a large predetermined distribution whether random or fully structured. Consequently, the required number of design evaluations makes Global Search techniques impractical for optimising expensive aerospace problems.

Evolutionary Algorithms consider the reflection of design parameters in the response values on the individual level. The fitness level of designs is calculated based on their response values. Inspired by evolution, a population of designs evolves towards the optimum through selection, recombination and mutation. The applicability of such a method is also impractical for expensive problems due to the immense cost of evaluating populations over the iterations.

Local Search techniques use the information of a handful of design-response values. Based on the few evaluated design-response values local properties of the problem are assumed which drives the optimisation. Probably the most prominent members of this class are the gradient-based techniques. They have already been successfully applied for many expensive optimisation problems. However, the computational demand significantly increases when no analytical gradients are available. Therefore, they are typically employed when adjoint solvers or automatic differentiation techniques are applicable.

Surrogate-based optimisation essentially assumes a well determinable relation between the design parameters and response values. In the context of machine learning, surrogate-based optimisation is called supervised learning. Alternatively, in engineering, the *Response Surface Methodology* is also used to label this kind of approaches [Korondi et al., 2020c]. The available design-response values are used for training a statistical model of the relationship. This is the so-called *surrogate* model which can be exploited to find optimal designs. Depending on the chosen surrogate model, this technique is highly applicable for optimising expensive problems as this technique requires the least number of observations. However, the quality of the trained surrogate model is vital, which can limit the applicability of this approach.

For the sake of completeness, it must be noted that the combination of the above-mentioned strategies is also possible. Indeed,

some of the most successful optimisation strategies combine the advantages of the listed techniques such as CMA-ES [Hansen, 2016], memetic algorithms [Turco, 2011; Vasile and Zuiani, 2011] and other hybrid strategies [Korondi, 2016; Sabater et al., 2020].

In this work, a *Bayesian optimisation* strategy is discussed [Frazier, 2018]. The surrogate model does not only define the relationship between design parameters and response values, but also calculates the uncertainty of the predicted response values. This allows the optimisation strategy to balance between searching for the best design parameters and improving the surrogate model quality during the optimisation.

Furthermore, this work assumes a highly limited computational budget. Therefore, the employed surrogate model incorporates information from aerodynamic solvers with various levels of fidelity.

Table 1.1: Classification of optimisation algorithms based on their information usage

	Global Search	Evolutionary Algorithm	Local Search	Surrogate-based optimisation
information usage	no response information used, only the distribution of design parameters	information of individual design-response values	multiple structured design-response values	global surrogate model of the design-response relation
number of design variables	high	high	high	low-moderate
number of samples	very high	high	low (pattern dependent, numerical derivation can be expensive)	low (scales linearly with the number of design variables)
advantages	finds global optimum, applicable for any cheap problem (combinatorial, discontinuous, multi-modal)	tends to find global optimum, applicable for any cheap problem (combinatorial, discontinuous, multi-modal)	consecutive improvement of designs, good convergence rate	finds global optimum, provides a surrogate model of the problem
disadvantages	expensive, slow convergence	still expensive, better convergence	finds local optima, relies on only local information (gradient, pattern, etc.)	highly depends on the surrogate model quality
examples	Monte Carlo method [Rubinstein and Kroese, 2016], Full-factorial	NSGA2 [Deb et al., 2000], MOGA-II [Poloni and Pediroda, 1997]	SLSQP [Kraft and Schnepfer, 1989], Nelder-Mead [Singer and Nelder, 2009]	EGO [Jones et al., 1998], SVM [Martin, 2002]
applicability for expensive problems	impractical	impractical	applicable (if economic gradient information is available)	applicable

1.3 Uncertainties in design optimisation

Probably, the most important paradox of science is contained in the well-known quote: "All models are wrong but some are useful." [Box, 1976]. This motivates the scientific community to come up with new models and update the existing ones even though we will never reach a perfect model. Nevertheless, these imperfect models provide us useful information to understand, explain or predict the reality around us. The imperfection of our models cause uncertainty (i.e. doubt) about the validity of the information obtained with the model. Therefore, we develop additional models and techniques to understand, explain or predict the degree of uncertainty.

Based on their assumptions, uncertainty modelling techniques can be labelled as *aleatory* or *epistemic* [Beyer and Sendhoff, 2007; Helton et al., 2010]⁴. The aleatory uncertainty is the inherent randomness of the system under investigation. For example, the wind directions during a flight can be considered as aleatory uncertainty. The complete history of weather conditions can be collected, and the distribution of wind directions can be accurately predicted. However, regardless of the amount of the collected data the uncertainty of the wind direction will never reduce to zero.

On the other hand, epistemic uncertainty stems from the lack of knowledge which can be reduced by gathering new information. The wind directions can be treated as a source of epistemic uncertainty as well. Imagine, only limited data of the weather condition history is available. The probability distribution of the wind directions cannot be determined with great certainty. Instead, the distribution of the wind directions is enveloped by a lower bounding belief and an upper bounding plausibility distribution. This type of uncertainty can be reduced by gathering new information.

In this dissertation, the uncertainty is treated probabilistically considering only aleatory uncertainty. We assume that all uncertainty can be represented by a random variable with a cor-

⁴Beyer and Sendhoff [2007] provide another possible classification by grouping the uncertainty modelling techniques into three categories: deterministic, probabilistic and possibilistic. In this classification the aleatory uncertainties belong to the probabilistic type and the epistemic uncertainties are divided into deterministic and possibilistic types.

responding probability distribution function (*Probability Density Function* (PDF) or *Cumulative Distribution Function* (CDF)).

In engineering, the designer can have control over only a limited number of parameters (or variables). Nevertheless, the variability of ignored or uncontrolled parameters can result in the variability of the response value. Therefore, the variations of the response function are taken into account to ensure the robust and reliable performance of real-world systems. Even when the probability distributions of the input parameters are known, the response probability distribution is usually unknown a priori or too costly to determine directly. Therefore, probabilistic UQ techniques are applied to obtain the PDF and CDF curves of the response values of the designs.

During the design optimisation process, the probability distributions of the response values are transformed into a risk-measure value. This scalar value can be used to optimise the design parameters.

“Young man, in mathematics you don’t understand things. You just get used to them.”

— John von Neumann

2.1 Information fusion

Surrogate-assisted optimisation is a commonly used technique for the design optimisation of computationally expensive problems [Forrester et al., 2008]. However, the number of attainable observations can be so sparse that standard surrogate techniques are not able to provide an accurate approximation. In this case, information from cheaper low-fidelity design analyses can be extracted and fused with the information of high-fidelity design analyses [Alexandrov et al., 2000]. Depending on whether the fidelities can be ranked or not, one can do non-hierarchical information fusion [Poloczek et al., 2017] or hierarchical information fusion [Kennedy and O’Hagan, 2000; Le Gratiet, 2013].

In *Computer-Aided Engineering* (CAE) design optimisation, hierarchical information fusion gained popularity due to the fact that many numerical solvers inherently provide an approximation of their numerical error [Forrester et al., 2007; Han et al., 2013]. This motivated the bi-fidelity hierarchical formulation of the *Gaussian Process Regression* (GPR) which is also called Co-Kriging [Myers, 1982; Xiao et al., 2018]. Co-Kriging was extended to fuse information from multiple levels of fidelity [Kennedy and O’Hagan, 2000] which is called *Multi-Fidelity Gaussian Process*

Regression (MFGPR). The method in [Kennedy and O’Hagan, 2000] requires composing and invert a large cross-correlation matrix containing the correlation information of all available fidelity levels. To avoid the construction and inversion of this large matrix, MFGPR was reformulated in a recursive form in [Le Gratiet, 2013]. The recursive formulation allows combining any levels of fidelities by a sequential and hierarchical model construction. Instead of building the model by combining the different levels of fidelities in a single step, the multi-fidelity surrogate model is built by sequentially adding the fidelity levels to the model starting from the lowest fidelity to the highest. In this work, the recursive MFGPR technique is employed as described in [Le Gratiet, 2013].

2.2 Gaussian process regression

In design optimisation, the GPR technique (or Kriging) is often used to predict the performance of a design. The performance of a d -dimensional design variable x can be calculated by a function $f(x)$ which is often unknown. Therefore, the performance function f must be modelled based on some (empirical) observations z . GPR models f as a random process \tilde{f} (i.e. random field¹). Accordingly, the performance of a design can be given by $\tilde{f}(x)$:

$$\tilde{f}(x) = m(x) + \tilde{\delta}(x), \quad (2.1)$$

where $m(x)$ is the mean trend and represents the global variation of the process. The mean trend is formulated as a least squares regression $m(x) = h(x)\beta$, where $h(x)$ is the vector of regression functions and β is the vector of regression coefficients. The $\tilde{\delta}$ is called local departure and represents local variations of the process [Rasmussen, 2003]. It is modelled as a Gaussian Process (GP) with zero mean and a corresponding covariance function ς , $\tilde{\delta} \sim \mathcal{GP}(0, \varsigma)$, so that $\tilde{\delta}(x) \sim \mathcal{N}(0, \Sigma(x, x'))$. The random process \tilde{f} can be observed with some noise conducting experiments (in our case computer-based simulations):

$$z(x) = \tilde{f}(x) + \epsilon = m(x) + \tilde{\delta}(x) + \epsilon, \quad (2.2)$$

¹In this dissertation, the terms *random process* and *random field* are interchangeable. Nevertheless, some mathematical details of the slight difference between the two terms are given in [Adler and Taylor, 2009].

where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is an independent and identically distributed Gaussian noise with zero mean and σ_ϵ variance. This independent noise term is included in our model for numerical purposes (i.e. to reduce the condition number of the covariance matrix) and not for modelling the uncertainty of the performance observations due to the variability of uncontrolled parameters.

Given n observations, the covariance matrix Σ is an n -by- n symmetric matrix, with entries $\Sigma_{ij} = \varsigma(x_i, x_j) = \sigma^2 R(x_i, x_j; \theta)$, where σ^2 is the global process variance, and R is the parametric correlation function with hyperparameters θ .

The covariance function describes our prior knowledge of the problem. As this knowledge is rarely complete, the covariance function is described by a parametric function. Its parameters are tuned based on our observations. The hyperparameter tuning of the covariance function is described in Section 2.3.

Given $z^{(n)}$ observations at design locations $x^{(n)}$, the joint probability distribution of the available observations and the prediction $\tilde{f}(x^*)$ at a new location x^* can be given in form:

$$\begin{pmatrix} \tilde{f}(x^*) \\ z^{(n)} \end{pmatrix} \sim \mathcal{N} \left(\begin{bmatrix} h(x^*)\beta \\ h(z^{(n)})\beta \end{bmatrix}, \begin{bmatrix} \varsigma(x^*, x^*) & \varsigma(x^*, x^{(n)})^T \\ \varsigma(x^*, x^{(n)}) & \Sigma + \sigma_\epsilon^2 I \end{bmatrix} \right), \quad (2.3)$$

where $\varsigma(x^*, x^{(n)})$ is the covariance vector of the new design and the evaluated designs and $\varsigma(x^*, x^*)$ is the value of the variance of the new design.

Then the joint distribution in Eq. (2.3) can be reformulated as a conditional probability distribution using the *Schur complement* of the joint distribution's covariance matrix [Von Mises, 2014]. The conditional posterior distribution is also a Gaussian distribution:

$$p(\tilde{f}(x^*) | z^{(n)}, \beta, \sigma^2, \sigma_\epsilon^2, \theta) \sim \mathcal{N}(\hat{m}_z(x^*), \hat{s}_z^2(x^*)), \quad (2.4)$$

where the mean $\hat{m}_z(x^*)$ and variance $\hat{s}_z^2(x^*)$ are:

$$\hat{m}_z(x^*) = h(x^*)\beta + \varsigma(x^*, x^{(n)})^T S^{-1} (z^{(n)} - H\beta), \quad (2.5)$$

$$\hat{s}_z^2(x^*) = \varsigma(x^*, x^*) - \varsigma(x^*, x^{(n)})^T S^{-1} \varsigma(x^*, x^{(n)}), \quad (2.6)$$

where $H = h(x^{(n)})$ is the matrix of regression functions evaluated at sample locations. $S = \Sigma + \sigma_\epsilon^2 I$ is the covariance matrix of the observations.

In surrogate-assisted design optimisation, the GPR mean \hat{m}_z is calculated in every point of a finite grid of the design space to obtain a prediction of the performance landscape. The prediction error of this landscape can be given by the GPR variance \hat{s}_z^2 which provides useful information about how much confidence we can have in our prediction.

2.3 Parameter estimation of Gaussian process regression

Rasmussen [2003] and Stein [2012] provide suggestions on choosing the family of the covariance function. The *squared exponential* is a popular choice due to its strong smoothness (i.e. infinitely differentiable). However, for modelling physical processes the *Matérn* covariance family might be a more realistic choice:

$$\varsigma(x_i, x_j) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{\theta} \|x_i - x_j\| \right)^\nu \mathcal{K}_\nu \left(\frac{\sqrt{2\nu}}{\theta} \|x_i - x_j\| \right), \quad (2.7)$$

where $\|x_i - x_j\|$ is the Euclidean distance of the design variables. Γ is the *gamma function* and \mathcal{K}_ν is the modified *Bessel function* of the second kind. ν is the smoothness parameter (here $\nu = 3/2$ is used). This covariance function is stationary (i.e. the covariance of two designs depends only on their distance measure and not on their positions in the design space).

The hyperparameters of the covariance function can be tuned by maximising their likelihood conditioned on the observations. Assuming that \tilde{f} is a multivariate normal distribution the GPR parameters $(\beta, \sigma^2, \sigma_\epsilon^2, \theta)$ can be determined by the *Maximum Likelihood Estimation* (MLE), where the likelihood function \mathcal{L} has the form:

$$\mathcal{L}(\beta, \sigma^2, \sigma_\epsilon^2, \theta | z^{(n)}) = \frac{1}{(2\pi)^{\frac{n}{2}} \sqrt{|S|}} \exp \left(-\frac{1}{2} (z^{(n)} - H\beta)^T S^{-1} (z^{(n)} - H\beta) \right), \quad (2.8)$$

where $S = \Sigma + \sigma_\epsilon^2 I$ is the covariance matrix of the observations. The generalised least squares estimate of β is given by:

$$\hat{\beta} = (H^T S^{-1} H)^{-1} H^T S^{-1} z^{(n)}. \quad (2.9)$$

The other parameters: σ^2 , σ_ϵ^2 and hyperparameters θ are obtained by minimising the negative log-likelihood function:

$$\sigma^2, \sigma_\epsilon^2, \theta = \underset{\sigma^2, \sigma_\epsilon^2, \theta}{\operatorname{argmin}} -\log \mathcal{L}, \quad (2.10)$$

where the log-likelihood function $\log \mathcal{L}$ is:

$$\begin{aligned} \log \mathcal{L}(\sigma^2, \sigma_\epsilon^2, \theta) = & -\frac{n}{2} \log(2\pi) - \frac{1}{2} \log(|(\Sigma + \sigma_\epsilon^2 I)|) \\ & - \frac{1}{2} (z^{(n)} - H\hat{\beta})^T (\Sigma + \sigma_\epsilon^2 I)^{-1} (z^{(n)} - H\hat{\beta}). \end{aligned} \quad (2.11)$$

2.4 Multi-fidelity Gaussian process regression

MFGPR can be framed in a recursive manner [Le Gratiet, 2013]. Information is ordered according to its fidelity level, where the rank of the lowest and highest fidelity are respectively $l = 1$ and $l = L$. The lowest fidelity information is modelled by a GPR and denoted by \tilde{f}_1 . Starting from the 2nd level, an auto-regressive model is then built at each fidelity level:

$$\begin{cases} \tilde{f}_l(x) = \rho_{l-1}(x) \tilde{f}_{l-1}(x) + h_l^T(x) \beta_l + \tilde{\delta}_l(x) + \epsilon_l, \\ \tilde{f}_{l-1}(x) \perp \tilde{\delta}_l(x), \\ \rho_{l-1}(x) = g_{l-1}^T(x) \beta_{\rho_{l-1}}, \end{cases} \quad (2.12)$$

where $\tilde{\delta}_l(x)$ is a Gaussian process with zero mean and covariance function $\sigma_l^2 R_l(x, x')$, and $\rho_{l-1}(x)$ is the adjustment parameter between $\tilde{f}_l(x)$ and $\tilde{f}_{l-1}(x)$. $g_{l-1}(x)$ and $h_l(x)$ are vectors of regression functions and $\beta_{\rho_{l-1}}$ and β_l are the vectors of coefficients. The ϵ_l is an independent Gaussian noise term used for numerical regularisation. The symbol \perp means that $\tilde{f}_{l-1}(x)$ and $\tilde{\delta}_l(x)$ are assumed to be independent. This assumption and a nested design set structure are required to derive the Bayesian parameter estimation for the recursive MFGPR [Le Gratiet, 2013]. *SMT: Surrogate Modeling Toolbox* (v0.3.4) [Bouhlel et al., 2019] relaxes

the requirement of a nested design set by using low-fidelity model predictions when a design evaluation is not available at lower fidelity levels².

The prediction at the l^{th} -level is given by a conditional probability distribution:

$$p(\tilde{f}_l(x^*)|\tilde{f}_{l-1}, z_l^{(n_l)}, \beta_l, \rho_{l-1}, \sigma_l^2, \sigma_{\epsilon_l}^2, \theta_l) \sim \mathcal{N}(\hat{m}_l(x^*), \hat{s}_l^2(x^*)), \quad (2.13)$$

where $z_l^{(n_l)}$ is the observation vector of n_l observations at level l . The MFGPR mean $\hat{m}_l(x^*)$ is:

$$\hat{m}_l(x^*) = \rho_{l-1}(x^*)\hat{m}_{l-1}(x^*) + h_l^T(x^*)\beta_l + \varsigma_l^T S_l^{-1} K_l, \quad (2.14)$$

$$K_l = z_l^{(n_l)} - \rho_{l-1}(x_l^{(n_l)}) \odot \tilde{f}_{l-1}(x_l^{(n_l)}) - H_l \beta_l, \quad (2.15)$$

where $S_l = (\Sigma_l + \sigma_{\epsilon_l}^2 I)$ is the covariance matrix of the observations at level l and \odot is the *Hadamard product* (i.e. element-wise product). The MFGPR variance is:

$$\begin{aligned} \hat{s}_l^2(x^*) &= \rho_{l-1}^2(x^*)\hat{s}_{l-1}^2(x^*) + \\ &+ \varsigma_l(x^*, x^*) - \varsigma_l(x^*, x_l^{(n_l)})^T S_l^{-1} \varsigma_l(x^*, x_l^{(n_l)}). \end{aligned} \quad (2.16)$$

By comparing Eqs.(2.5)-(2.6) and Eqs.(2.14)-(2.16), we can see that the auto-regressive model solves a simple GPR problem at each level. The random process is the sum of the adjusted model of the previous level and a difference term (i.e. linear multi-fidelity information fusion [Perdikaris et al., 2017]). Accordingly, the mean and the variance of the model are also the sums of the adjusted mean and variance of the previous level model and the mean and variance of the difference term.

²Recently, *SMT: Surrogate Modeling Toolbox* was updated. Following the version update (v0.4.2) on 20.02.2020, the code suggests to respect the requirement of a nested design set. In this dissertation, the results are generated with the older version (v0.3.4) employing a non-nested design structure. This introduce an additional (unintended) numerical error in the MFGPR model predictions.

2.5 Parameter estimation of multi-fidelity Gaussian process regression

Similarly to the single level GPR, the parameters can be determined by the MLE. The likelihood function has the form:

$$\begin{aligned} \mathcal{L} \left(\beta_l, \rho_{l-1}, \sigma_l^2, \sigma_{\epsilon_l}^2, \theta_l \mid z_l^{(n_l)}, \right) &= \\ &= \frac{1}{(2\pi)^{\frac{n_l}{2}} \sqrt{|S_l|}} \exp \left(-\frac{1}{2} K_l^T S_l^{-1} K_l \right). \end{aligned} \quad (2.17)$$

The generalised least squares estimate of β and ρ_{l-1} is given by:

$$\begin{aligned} \begin{bmatrix} \hat{\rho}_{l-1} \\ \hat{\beta}_l \end{bmatrix} &= (A_l^T S_l^{-1} A_l)^{-1} \left(A_l^T S_l^{-1} \left[\tilde{f}_{l-1} \left(x_l^{(n_l)} \right) \right] \right), \\ A_l &= \begin{bmatrix} \tilde{f}_{l-1}(x_l^{(n_l)}) & H_l \end{bmatrix}. \end{aligned} \quad (2.18)$$

The other parameters: σ_l^2 , $\sigma_{\epsilon_l}^2$ and hyperparameters θ_l are obtained by minimising the negative log-likelihood function:

$$\sigma_l^2, \sigma_{\epsilon_l}^2, \theta_l = \underset{\sigma_l^2, \sigma_{\epsilon_l}^2, \theta_l}{\operatorname{argmin}} -\log \mathcal{L}, \quad (2.19)$$

where the log-likelihood function $\log \mathcal{L}$ is:

$$\begin{aligned} \log \mathcal{L} (\sigma_l^2, \sigma_{\epsilon_l}^2, \theta_l) &= \\ &= -\frac{n_l}{2} \log(2\pi) - \frac{1}{2} \log (|(\Sigma_l + \sigma_{\epsilon_l}^2 I)|) - \frac{1}{2} \hat{K}_l^T (\Sigma_l + \sigma_{\epsilon_l}^2 I)^{-1} \hat{K}_l, \end{aligned} \quad (2.20)$$

where:

$$\hat{K}_l = z_l^{(n_l)} - \hat{\rho}_{l-1}(x_l^{(n_l)}) \odot \tilde{f}_{l-1}(x_l^{(n_l)}) - H_l \hat{\beta}_l. \quad (2.21)$$

“The miracle of the appropriateness of the language of mathematics for the formulation of the laws of physics is a wonderful gift which we neither understand nor deserve.”

— Eugene Wigner

3.1 Statistical measures

In probability theory, random values are represented by probability distributions [Maleki and Do, 2000]. Consequently, the response values are also characterised by a probability distribution in probabilistic design optimisation. In the scope of this dissertation, the *Cumulative Distribution Function* (CDF) and *Probability Density Function* (PDF) are used to describe random variables. The CDF, denoted with Φ , is defined as:

$$\Phi(y) = P[\tilde{y} \leq y]. \quad (3.1)$$

The PDF, denoted with ϕ , is defined as:

$$\phi(y) = \frac{d\Phi(y)}{dy}. \quad (3.2)$$

The decision making is difficult based on probability distributions due to their multitude of possible realisations [Rockafellar and Royset, 2015]. Therefore, various properties of the distribution, or risk measures \mathcal{R} , are used to quantify the uncertainty.

Mean Expectation is probably the most used statistical measure. It represents the average value of a large number of realisations. It is the most intuitive outcome expectation of a random variable.

$$\mu = \mathbb{E} [\tilde{y}] = \int y\phi(y)dy. \quad (3.3)$$

The probability distribution can be also characterised by its *quantile* (or Value-at-Risk) function. It is defined as the inverse of the CDF.

$$q_\zeta = \underset{y}{\operatorname{argmin}} [y | \Phi(y) \geq \zeta], \quad (3.4)$$

where q_ζ reads as the ζ -quantile and represents the value of the random variable such that lesser or equal realisations than the q_ζ value occur with ζ probability, where $\zeta \in \mathbb{R}$ and $0 \leq \zeta \leq 1$. Specialised quantile values split the domain of possible realisations into equally probable sub-domains. For example, the domain can be split into a hundred sub-domains. Then the values separating the sub-domains are called 1st, 2nd, ..., 99th percentile. Essentially, the ξ^{th} percentile corresponds to the $q_{\xi/100}$ value, where $\xi \in \mathbb{N}$ and $0 \leq \xi \leq 100$. It follows that the q_0 is the minimum (i.e. infimum) and q_1 is the maximum (i.e. supremum) of the possible realisations of the random variable.

The *superquantile* measure is commonly called as conditional Value-at-Risk as well. The superquantile calculates the expected value of the random variable when the realisation of the random variable is higher than a certain quantile value. It is defined as [Rockafellar and Royset, 2015]:

$$\mathcal{R} = \bar{q}_\zeta = \mathbb{E} [\tilde{y} | \tilde{y} \geq q_\zeta(\tilde{y})] = \frac{1}{1-\zeta} \int_\zeta^1 q_\tau(\tilde{y})d\tau, \quad (3.5)$$

where q_ζ is the quantile value below which the realisations are ignored. For instance, $\zeta = 0$ corresponds to the mean expected value and $\zeta = 1$ corresponds to the worst-case scenario. In engineering practice, the superquantile is chosen over the quantile when the shape of the probability distribution tail is relevant as in case of reliability-based optimisation [Rockafellar and Uryasev, 2002]. Additionally, the superquantile is a coherent and regular measure and does not depend on scaling which properties are preferred from a mathematical point of view [Quagliarella et al., 2014; Rockafellar and Royset, 2015].

The analytical computation of any risk measures is typically intractable apart from some benign problems. Therefore, the calculation of most risk measures requires a statistically significant sample, which is impractical if the performance evaluation is computationally expensive. For this reason, surrogate-assisted uncertainty quantification can be used to replace the expensive performance evaluations with cheap evaluations of the surrogate model of the probabilistic space. In this way, a statistically significant sample can be achieved at low cost.

3.2 Surrogate-assisted uncertainty propagation

Ideally, the design and uncertainty space should be modelled together. However, the construction of a surrogate directly on the coupled space of design variables and uncertainty variables is computationally challenging due to the high dimensionality of the problem. Therefore, the design and probability spaces are separated. At each design sample location, we sample the uncoupled probabilistic space by varying only the uncertain variables and train a surrogate model of the probabilistic space. These local probabilistic models can be used to evaluate the high number of samples required for the computation of the risk measure \mathcal{R} . With the above-described procedure, one can calculate the risk measure for each design at a relatively low computational cost.

The domain separation is a critical step as it forces us to train separate surrogates for the probability space at each design sample location and an additional surrogate over the risk measure \mathcal{R} in the design space. This introduces additional numerical error due to the nested approximations but allows us to optimise a computationally heavy design problem under uncertainty.

3.3 An academic example for surrogate-assisted uncertainty propagation

In this section the surrogate-assisted uncertainty quantification is described by employing it on an academic example. Let us denote the input uncertainty with \tilde{u} and the probabilistic output with \tilde{y} .

The performance function is denoted by g . Then the response of a design x can be given by:

$$\tilde{y} = g(x, \tilde{u}), \quad (3.6)$$

Usually, the performance function is not known. However, in our academic example we can neglect the dependency on the deterministic design variables and imagine that:

$$\tilde{y} = \sinh(\tilde{u}). \quad (3.7)$$

The distribution of the input is considered to be standard normal Gaussian $\mathcal{N}(0, 1)$. In order to accurately approximate the distribution of the output \tilde{y} , we have to employ an uncertainty quantification method. The most straightforward solution would be to use *Monte Carlo Sampling* (MCS) [Lemieux, 2009]. However, the MCS technique cannot be applied directly when the evaluation of g is computationally expensive. Therefore, a surrogate is trained on a sparse set of samples and a large number of samples are obtained virtually by replacing evaluations of g with evaluations of the surrogate model.

The surrogate model type can be chosen deliberately depending on the problem. For example, the probability space can be modelled by a *Polynomial Chaos Expansion* (PCE)¹. This surrogate can model the probability space efficiently for a range of probability distributions depending on the polynomial family of the regression functions [Choi et al., 2004]. Here, in correspondence with the Gaussian uncertainty, the Hermite polynomial family is employed up to third-order functions.

The comparison between the direct sampling technique and surrogate-assisted uncertainty quantification is displayed in Figure 3.1. We can see that PCE introduces some approximation errors but it allows us to draw a large number of samples to approximate the distribution of \tilde{y} .

After obtaining the distribution of the random variable \tilde{y} an *Empirical Cumulative Distribution Function* (ECDF) can be calculated, see Figure 3.2(a). By inverting the ECDF, we can

¹Polynomial Chaos Expansion can be considered as a particular subclass of GPR with particular regression functions and Kronecker delta covariance function [Schobi et al., 2015].

3.3. An academic example for surrogate-assisted uncertainty propagation

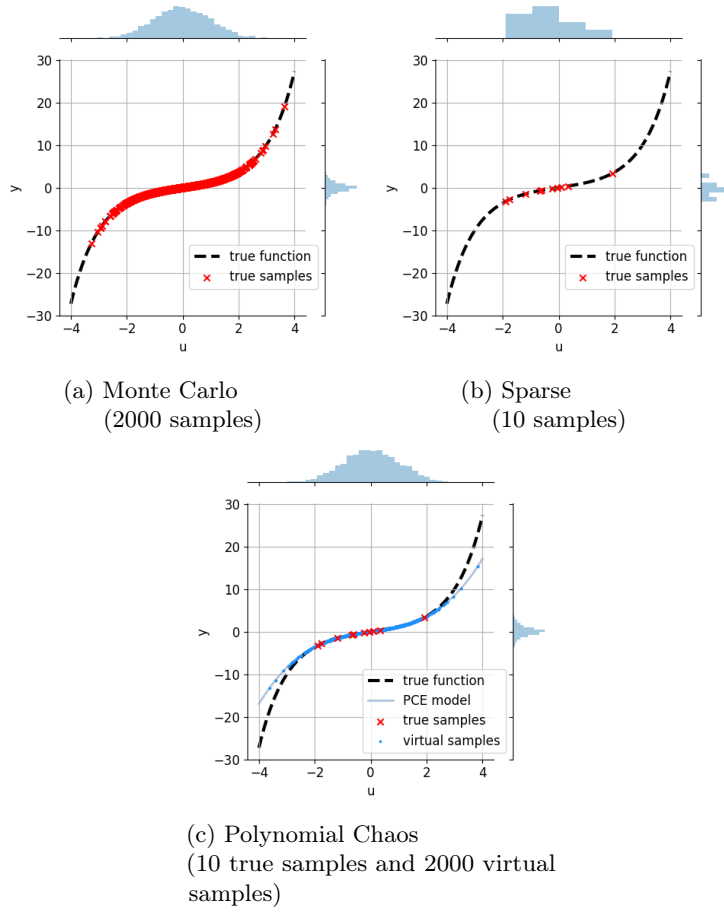


FIGURE 3.1: Comparison of Monte Carlo sampling and virtual sampling with PCE

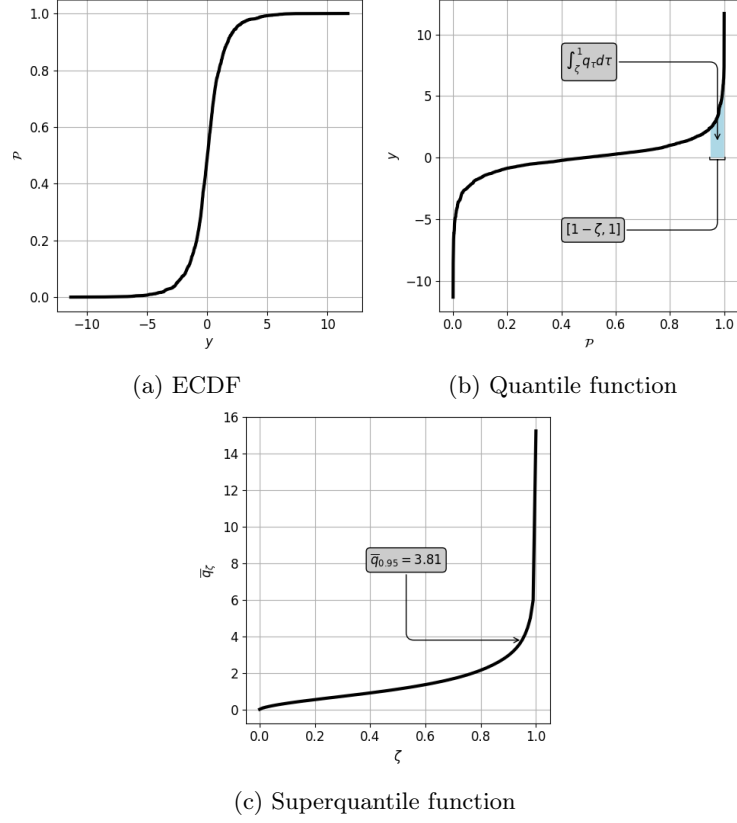


FIGURE 3.2: Superquantile calculation

obtain the quantile values of the random variable. Then the superquantile \bar{q}_ζ can be obtained by Eq. (3.5). Imagine that we are interested in the $\bar{q}_{0.95}$ value. This requires the integration of the quantile curve between 0.95 and 1, as depicted in Figure 3.2(b). The superquantile values corresponding to each ζ level is given in Figure 3.2(c) and the superquantile corresponding to level $\bar{q}_{0.95}$ is 3.81.

3.4 Implications for design optimisation of real-world systems

In the context of design optimisation, we are interested in one or more system responses or *Quantities of Interest* (QoIs). The uncertainties of input variables and environmental parameters can result in multiple uncertain system responses \tilde{y}_j (with $j = 1, \dots, n_q$), i.e. in multiple uncertain QoIs. These uncertainties must be considered in the design optimisation process otherwise the obtained optimum will be biased. Thus, optimisation under uncertainty does not usually consider, as the deterministic, QoIs to be optimised, but rather risk measures \mathcal{R}_k (with $k = 1, \dots, n_r$) of the QoIs. Note that the number of QoIs n_q is not the same as the number of uncertainty measures n_r in the sense that various risk measures can be calculated from the distribution of a QoI. For example, one might be interested in optimising both the mean \bar{q}_0 and the $\bar{q}_{0.95}$ of a random valued QoI. This results in a multi-dimensional objective space. However, in regard the surrogate training, the objectives are treated separately and for each objective a surrogate model is trained.²

In the following, the $\bar{q}_{\xi/100}$ will be called as the ξ^{th} *superper-centile* and will be denoted by \mathcal{S}_ξ . In this dissertation, the 95th superpercentile is employed for reliability-based optimisation.

The \mathcal{S}_{95} risk measure is defined for minimisation. Therefore, it provides an estimate of the expected value of the right-tail of the distribution. To distinguish the right-tail and left-tail superpercentile values, the $\mathcal{S}_{95}^{\text{right}}$ and $\mathcal{S}_{95}^{\text{left}}$ notations are applied respectively.

²It is also possible to train a multi-output GPR. However, the training of a multi-output GPR is challenging since a kernel should be found which can capture the correlation of both the input and output space [Press et al., 2007].

“Intelligence is the ability to adapt to change.”

— Stephen Hawking

4.1 Surrogate-assisted design optimisation

Surrogate-assisted design optimisation has three cornerstones: data, surrogate model and model update strategy. When the data is loaded with uncertainty an *Uncertainty Quantification* (UQ) technique also complements the optimisation process. In this dissertation, the main source of uncertainty stems from the variability of ignored or uncontrolled parameters of an aerospace application. Due to the computationally demanding aerodynamic simulations surrogate-based UQ techniques are used as was discussed in Chapter 3. The surrogate is a *Multi-Fidelity Gaussian Process Regression* (MFGPR) model which is capable to predict the performance of a design at various fidelity levels as was detailed in Chapter 2. The model update strategy has two main goals: evaluate designs with optimal performance and reduce the uncertainty of the model particularly in the neighbourhood of the optimum. These two main goals are also called *exploitation* and *exploration* respectively.

The MFGPR model is a conditional posterior distribution and exploits the properties of the *Bayes’ theorem*. A surrogate-assisted design optimisation technique with such surrogate runs under the name of *Bayesian optimisation*. In Bayesian optimisation, the model update strategy is labelled as *acquisition function* which usually also exploits the properties of the conditional distribution.

For single-objective problems, *Efficient Global Optimisation* (EGO) realises a Bayesian optimisation by updating the trained *Gaussian Process Regression* (GPR) model at locations where the *Expected Improvement* (EI) acquisition function is maximal [Jones et al., 1998]. When EGO is employed for multi-objective optimisation, a separate GPR is trained for each objective and the objectives are aggregated into a single scalar measure. Then similarly to the single-objective case, the acquisition function can identify the location of the new design candidate. For instance, ParEGO [Knowles, 2006] employs the *Chebyshev scalarisation* and MEGO [Montrone and Rigoni, 2018] uses the *Hypervolume Improvement* (HVI).

By generating the surrogate on a multi-fidelity data set, the MFGPR model can predict the responses and their model uncertainty at various levels of fidelity. The information obtained from each fidelity level can be utilised in the acquisition function of a multi-fidelity surrogate-based optimisation strategy. Additionally, the acquisition function must select the fidelity of the new design evaluation as well.

In the past years, numerous multi-fidelity surrogate-based strategies were proposed by the research community. Ariyarat and Kanazaki [2017] and Bonfiglio et al. [2018] propose the usage of multi-fidelity surrogates in a Bayesian optimisation context; however, the model updates are only realised with the high-fidelity level. The acquisition function in [Ariyarat and Kanazaki, 2017; Bonfiglio et al., 2018] considers only the top-level model prediction and variance.

Huang et al. [2006] propose to select the new design location and fidelity level simultaneously by maximising an augmented EI. Their augmented EI function contains heuristic terms accounting for the correlation between the fidelity levels, the data error and cost of various fidelity levels. Sacher et al. [2020] extend the augmented EI by taking the top-level model variance into account as well. Similarly, Zhang et al. [2018] provide an augmented EI which is also complemented with a constraint treatment.

In [Swersky et al., 2013; Zhang et al., 2017; Takeno et al., 2019] various formulations of multi-fidelity *entropy search* acquisition functions are detailed.

Kandasamy et al. [2016] present an in-depth discussion on

multi-fidelity *bandit*¹ problems where a two-stage acquisition function is given. The location of the new sample is calculated by maximising the *upper confidence bound* of the model prediction in the first stage. Then the fidelity is selected following intuition-based sequential threshold criteria.

Meliani et al. [2019] also propose a two-stage approach. The most promising candidate is selected with the EI acquisition function. Then a model uncertainty reduction criterion selects the fidelity exploiting the nested structure of multi-fidelity designs. This strategy translates the suggestion of Le Gratiet [2013] for sequentially improving the MFGPR model into a design optimisation context. To this end, Serafino et al. [2020] adopt the above strategy with an additional preference towards the low-fidelity calculation. High-fidelity calculations are selected only if the model uncertainty of the low-fidelity is below a certain threshold level.

For reliability-based optimisation, Chaudhuri et al. [2019] propose to select the new infill location by maximising the expected *feasibility function* and the fidelity selection is advised by calculating the information gain with the Kullback-Leibler divergence of the model prediction and the hypothetical model prediction after the update step.

The *No Free Lunch theorem* [Wolpert and Macready, 1997] holds also for multi-fidelity surrogate-based strategies and no method showed clear superiority to the other. Nevertheless, the advantages and disadvantages of the published works helped to develop a multi-fidelity surrogate-based optimisation strategy for computationally expensive problems. With the suggestions of the above-mentioned papers in mind, the proposed strategy of this dissertation aimed to:

- find optimal designs evaluated with the highest fidelity;
- tackle uncertain problems;
- tackle computationally expensive problems;
- consider the surrogate quality of all fidelity levels during the new design acquisition;

¹The fidelity selection in a multi-fidelity surrogate-based optimisation problem can be seen as a bandit problem [Gittins and Jones, 1979].

- consider the cost and fidelity of the newly acquired designs;
- anticipate future adaptations of the workflow by means of modularity.

Therefore, a multi-fidelity surrogate-assisted approach is developed with a two-stage acquisition function to facilitate modularity. The sample location selection and the fidelity selection are decoupled in such way that the underlying sub-problems can be treated with different rationales. The new sampling location is determined by balancing exploration and exploitation. Then the fidelity level is selected by comparing the amount of model uncertainty reduction against the cost of the design evaluation. Furthermore, the designs are evaluated with a nested surrogate-assisted UQ process to tackle uncertain problems computationally efficiently. An in-depth discussion of the proposed strategy is provided in the following sections.

4.2 Proposed single-objective framework

The single-objective surrogate-assisted design optimisation under uncertainty can be broken down to four steps: sample location determination, design evaluation, risk measure evaluation and surrogate training. The proposed strategy repeats these steps until the computational budget is exhausted. The proposed framework is illustrated in Figure 4.1. The framework is developed for the MFGPR surrogate model. In the initial iteration, the sample locations are determined with a *Design of Experiments* (DoE) technique at each fidelity level. In this dissertation, the uniform *Latin Hypercube Sampling* (LHS) is employed [McKay et al., 2000]. The uncertain responses of the designs are quantified with a surrogate-assisted UQ technique calculating the superpercentile risk measure introduced in Section 3.4.

The pseudo-code of the proposed single-objective strategy is given in Algorithm 1. The DoE samples are generated for each fidelity level $x_l^{(n_l)}$, where n_l is the number of samples generated at fidelity l . Moreover, the uncertain space is sampled as well. For each design sample $x_{l,i}$ an additional sample is generated in the probabilistic space $u_{l,i}^{(m_l)}$.

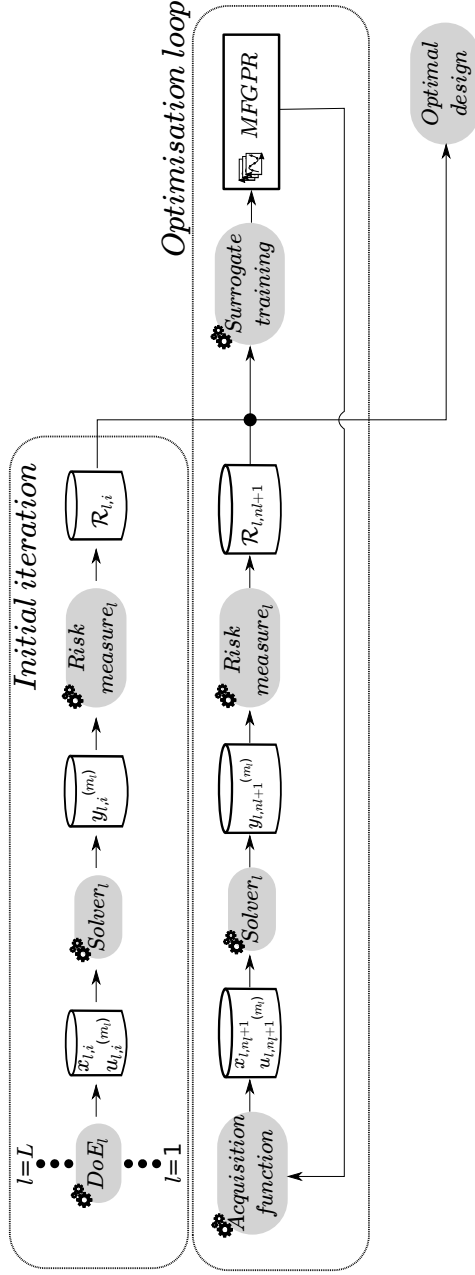


FIGURE 4.1: Proposed single-objective surrogate-assisted design optimisation framework for expensive probabilistic problems.

Algorithm 1 Single-objective multi-fidelity surrogate-based design optimisation strategy under uncertainty. (L : top-level fidelity, n_l : number of design samples at level l , m_l : number of probabilistic samples)

Input: d, L, n_l, m_l \triangleright Problem and model properties
Output: $x_{L,optimal}, \mathcal{R}_{L,optimal}, \text{MFGPR}$

```

1: for  $l=1:L$  do
2:    $x_l \leftarrow \text{generateDesignSample}()$ 
3:   for  $i=1:n_l$  do
4:      $u_{l,i} \leftarrow \text{generateProbabilisticSample}(x_{l,i})$ 
5:      $y_{l,i} \leftarrow \text{evaluateDesignPerformance}(x_{l,i}, u_{l,i})$ 
6:      $\mathcal{R}_{l,i} \leftarrow \text{evaluateRiskMeasure}(y_{l,i})$ 
7:   end for
8: end for
9:  $\text{MFGPR} \leftarrow \text{trainMFGPR}(\mathcal{R}_{1:L,1:n_l})$ 
10:  $flag \leftarrow True$ 
11: while  $flag$  do
12:    $x_{l,n_l+1} \leftarrow \text{acquisitionFunction}(\text{MFGPR})$ 
13:    $u_{l,n_l+1} \leftarrow \text{generateProbabilisticSample}(x_{l,n_l+1})$ 
14:    $y_{l,n_l+1} \leftarrow \text{evaluateDesignPerformance}(x_{l,n_l+1})$ 
15:    $\mathcal{R}_{l,n_l+1} \leftarrow \text{evaluateRiskMeasure}(y_{l,n_l+1})$ 
16:    $\text{MFGPR}_k \leftarrow \text{trainMFGPR}([\mathcal{R}_{1:L,1:n_l}, \mathcal{R}_{l,n_l+1}])$ 
17:   if  $\text{budgetExhausted}$  then
18:      $flag \leftarrow False$ 
19:      $x_{L,optimal}, \mathcal{R}_{L,optimal} \leftarrow \text{findBest}(x_{L,:}, \mathcal{R}_{L,:})$ 
20:   end if
21: end while
22: return  $x_{L,optimal}, \mathcal{R}_{L,optimal}, \text{MFGPR}$ 

```

At each fidelity level, every design is evaluated m_l times resulting in a performance matrix with entries of $y_{l,i}^{(m_l)}$, where i is the corresponding design index. To obtain the objective, the corresponding risk measure $\mathcal{R}_{l,i}$ is calculated on the performance values $y_{l,i}^{(m_l)}$.

At this point, the performance information coming from different fidelities is fused together resulting in the MFGPR model. The surrogate model provides cheap-to-evaluate approximations of our underlying problem and helps us to identify promising design

candidates. When high fidelity evaluations are involved, the time required for training the surrogate model is negligible compared to the time required to evaluate the performance of a single design.

The acquisition function updates the pool of design samples with a single design $x_{l,n_l+1}, u_{l,n_l+1}^{(m_l)}$. For the new design, the corresponding objective can be calculated as previously described resulting in \mathcal{R}_{l,n_l+1} . Then the surrogate model of the objective is retrained on the augmented risk measure samples. We iterate the model update process until the computational budget is exhausted and then return the best design of all the evaluated designs.

4.3 Single-objective acquisition function

For single-objective problems, the *Expected Improvement* (EI) acquisition function is employed due to its proven success for various kind of problems. The EI calculates the expected value of the improvement on the current best evaluation by adding a new observation as described in [Keane, 2006]. The magnitude of the improvement is:

$$\tilde{I}(x_{n+1}) = \max(0, z_{min} - \tilde{z}(x_{n+1})), \quad (4.1)$$

where $z_{min} = \min_{x_i=1\dots n} z(x_i)$ is the current best evaluation. The $\tilde{z}(x_{n+1})$ is not known a priori. Therefore it is replaced by a Gaussian random variable with a mean $\hat{m}(x_{n+1})$ and variance $\hat{s}^2(x_{n+1})$ as defined by Eq. (2.14) and Eq. (2.16) respectively. Consequently, also the improvement \tilde{I} is a random variable. The probability of the improvement can be calculated, for example as in [Emmerich et al., 2006]:

$$P[\tilde{I}(x_{n+1})] = P[\tilde{z}(x_{n+1}) \leq z_{min}] \quad (4.2)$$

$$= \Phi\left(\frac{z_{min} - \hat{m}(x)}{\hat{s}^2(x)}\right), \quad (4.3)$$

where Φ is the cumulative distribution function of the standard normal distribution. The argument of the cumulative distribution function is called germ and denoted by λ :

$$\lambda(x) = \frac{z_{min} - \hat{m}(x)}{\hat{s}^2(x)}, \quad (4.4)$$

The expectation of the improvement $\mathbb{E}[\tilde{I}(x_{n+1})]$ can be marginalised out and given in closed form:

$$\text{EI}(x_{n+1}) = \hat{s}^2(x_{n+1}) (\lambda(x_{n+1}) \Phi(\lambda(x_{n+1})) + \phi(\lambda(x_{n+1}))), \quad (4.5)$$

where ϕ is the probability density function of the standard normal distribution. The location of the new observation can be determined by maximising the $\text{EI}(x)$.

$$x_{n+1} = \underset{x}{\operatorname{argmax}} \text{EI}(x). \quad (4.6)$$

The maximum EI acquisition function offers a great balance between evaluating designs with a potentially optimal objective value and designs with high prediction uncertainty in unexplored regions.

For problems with inequality constraint, the EI can be extended with a probability of feasibility term as in [Gardner et al., 2014]. Then, the *Constrained Expected Improvement* (cEI) can be formulated as:

$$\begin{aligned} \text{cEI} &= \mathbb{E} [\max(0, z_{\min} - \tilde{z}(x_{n+1}))] \mathbb{P} [\tilde{h}(x_{n+1}) \leq h_0] \\ &= \hat{s}^2(x_{n+1}) (\lambda(x_{n+1}) \Phi(\lambda(x_{n+1})) + \phi(\lambda(x_{n+1}))) \Phi(\lambda_h(x_{n+1})), \end{aligned} \quad (4.7)$$

where $\tilde{h}(x)$ is the constraint response function which required to be smaller or equal to a threshold value h_0 . The λ_h is the germ of the constraint function. It is defined as:

$$\lambda_h(x) = \left(\frac{h_0 - \hat{h}(x)}{\hat{s}_h^2(x)} \right), \quad (4.8)$$

where \hat{h} and \hat{s}_h^2 are the predicted mean and variance of the constraint response function.

To determine the fidelity of the new observation, we can define the amount of information which can be gained for each fidelity level by performing the observation at that level. As a gain measure, we use the predicted variance reduction similarly to [Le Gratiet, 2013; Meliani et al., 2019]. This gain can be scaled with the observation cost c_l of the fidelity level l . By comparing the scaled costs of the different fidelities, we can choose the fidelity level

which provides the highest *Scaled Expected Variance Reduction* (SEVR):

$$l = \underset{1:L}{\operatorname{argmax}} \operatorname{SEVR}_l, \quad (4.9)$$

where the variance reduction $\tilde{\sigma}_l$ can be recursively defined as:

$$\tilde{\sigma}_1 = \rho_1^2 \hat{s}_{z_1}^2(x_{new}), \quad \text{for } l = 1, \quad (4.10)$$

$$\tilde{\sigma}_l = \rho_{l-1}^2 \hat{s}_{z_{l-1}}^2(x_{new}) + \hat{s}_{\delta_l}^2(x_{new}), \quad \text{for } l \geq 2. \quad (4.11)$$

The SEVR is defined as:

$$\operatorname{SEVR}_l = \frac{\tilde{\sigma}_l}{c_l}. \quad (4.12)$$

The pseudo-code of the acquisition function for single-objective problems with fidelity level selection is given in Algorithm 2. For constraint problems, the described fidelity selection process can be applied without any modification as the cEI will provide feasible design candidates or designs with negligible constraint violation.

Algorithm 2 Acquisition function for single-objective optimisation

Input: MFGPR_L

Output: x_{l,n_l+1}

```

1:  $\hat{m}_L \leftarrow \text{predictTrend}(\text{MFGPR}_L)$ 
2:  $\hat{s}_L^2 \leftarrow \text{predictVariance}(\text{MFGPR}_L)$ 
3:  $EI(x) \leftarrow \text{calculateExpectedImprovement}(\mathcal{R}_{min}, \hat{m}_L(x), \hat{s}_L^2(x))$ 
4:  $(x_{?,n_l+1}) \leftarrow \operatorname{argmax}(EI(x))$ 
5: for  $l=1:L$  do
6:    $\tilde{\sigma}_l \leftarrow \text{predictVarianceReduction}(\text{MFGPR}_l, (x_{?,n_l+1}))$ 
7:    $c_l \leftarrow \text{getCost}(l)$ 
8:    $\operatorname{SEVR}_l \leftarrow \text{calculateSEVR}(\tilde{\sigma}_l, c_l)$ 
9: end for
10:  $l \leftarrow \operatorname{argmax}(\operatorname{SEVR}_l)$ 
11: return  $x_{l,n_l+1}$ 

```

4.4 Proposed multi-objective framework

The proposed algorithm for multi-objective problems differs from the single-objective variant in that multiple risk measure objectives

have to be calculated. Consequently, a MFGPR_k model is trained for each set of risk measure values $\mathcal{R}_{l,k,:}$ and the acquisition function is adapted for multi-objective problems. Correspondingly, the outcome of the multi-objective optimisation is not a single optimal solution but a set of Pareto optimal solutions. The multi-objective strategy is depicted in Figure 4.2 and given by Algorithm 3. The number of response values is denoted by n_q and the number of uncertain QoIs is n_r .

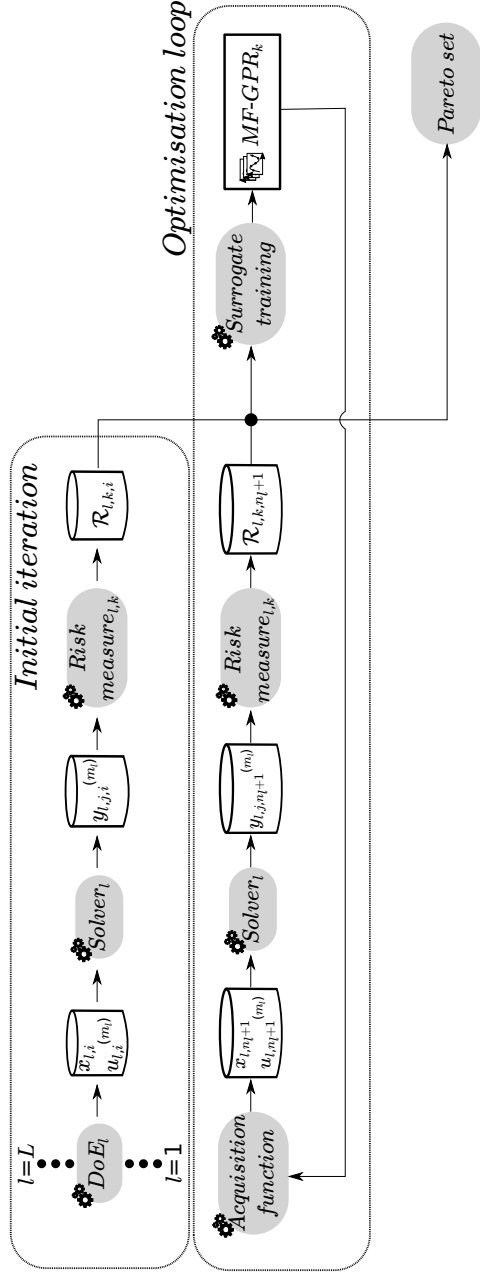


FIGURE 4.2: Workflow of multi-objective multi-fidelity surrogate-based optimisation under uncertainty

Algorithm 3 Multi-objective multi-fidelity optimisation strategy.
 (L : top-level fidelity, n_l : number of design samples at level l , n_q : number of QoIs, n_r : number of risk measures, m_l : number of probabilistic samples)

Input: d, L, n_l, n_q, n_r, m_l \triangleright Problem and model properties
Output: $x_{L,pareto}, \mathcal{R}_{L,1:n_r,pareto}, \text{MFGPR}_{1:n_r}$

```

1: for  $l=1:L$  do
2:    $x_l \leftarrow \text{generateDesignSample}()$ 
3:   for  $i=1:n_l$  do
4:      $u_{l,i} \leftarrow \text{generateProbabilisticSample}(x_{l,i})$ 
5:      $y_{l,1:n_q,i} \leftarrow \text{evaluateDesignPerformance}(x_{l,i}, u_{l,i})$ 
6:      $\mathcal{R}_{l,1:n_r,i} \leftarrow \text{evaluateRiskMeasure}(y_{l,1:n_q,i})$ 
7:   end for
8: end for
9: for  $k=1:n_r$  do
10:   $\text{MFGPR}_k \leftarrow \text{trainMFGPR}(\mathcal{R}_{1:L,k,1:n_l})$ 
11: end for
12:  $flag \leftarrow True$ 
13: while  $flag$  do
14:   $x_{l,n_l+1} \leftarrow \text{acquisitionFunction}(\text{MFGPR}_{1:n_r})$ 
15:   $u_{l,n_l+1} \leftarrow \text{generateProbabilisticSample}(x_{l,n_l+1})$ 
16:   $y_{l,1:n_q,n_l+1} \leftarrow \text{evaluateDesignPerformance}(x_{l,n_l+1})$ 
17:   $\mathcal{R}_{l,1:n_r,n_l+1} \leftarrow \text{evaluateRiskMeasure}(y_{l,n_l+1})$ 
18:  for  $k=1:n_r$  do
19:     $\text{MFGPR}_k \leftarrow \text{trainMFGPR}([\mathcal{R}_{1:L,k,1:n_l}, \mathcal{R}_{l,k,n_l+1}])$ 
20:  end for
21:  if  $\text{budgetExhausted}$  then
22:     $flag \leftarrow False$ 
23:     $x_{L,pareto}, \mathcal{R}_{L,1:n_r,pareto} \leftarrow \text{findPareto}(x_{L,:}, \mathcal{R}_{L,:,:})$ 
24:  end if
25: end while
26: return  $x_{L,pareto}, \mathcal{R}_{L,1:n_r,pareto}, \text{MFGPR}_{1:n_r}$ 

```

4.5 Multi-objective acquisition function

In multi-objective optimisation, a one-step lookahead strategy requires to aggregate the objectives into a single measure [Lam et al., 2018]. Similarly to the single-objective case, an acquisition

function is used to determine the location of the new observation. Here, the HVI is adapted to scalarise the multi-objective problem. HVI is a popular technique due to the fact that the *Hypervolume* (HV) measure is a reliable performance indicator which does not require any assumption on the true Pareto set [Emmerich et al., 2006]. HVI is defined as:

$$\text{HVI}(x) = \text{HV}(\mathcal{P} \cup \hat{m}(x_{n+1})) - \text{HV}(\mathcal{P}), \quad (4.13)$$

where the HV is a Lebesgue-measure in the objective space and \mathcal{P} denotes the current Pareto set. The next observation is made where the predicted mean value of the GPR maximises the HVI measure. However, this strategy does not consider the variance of the prediction but only exploits the input space. To introduce some exploratory behaviour into the strategy, the *Hypervolume Improvement of the Lower Confidence Bound* (HVI_{LCB}) is maximised instead as suggested by Emmerich et al. [2006]:

$$x_{n+1} = \underset{x}{\operatorname{argmax}} \text{HVI}_{\text{LCB}}(\mathcal{P} \cup \hat{m}_{\text{LCB}}(x_{n+1})), \quad (4.14)$$

$$\hat{m}_{\text{LCB}}(x) = \hat{m}(x) - \alpha \hat{s}(x), \quad (4.15)$$

where α is a positive parameter that represents the level of optimism towards model uncertainty. Here, $\alpha = 2$ as suggested in [Emmerich et al., 2006]. The HVI and HVI_{LCB} are illustrated for a bi-objective problem in Figure 4.3.

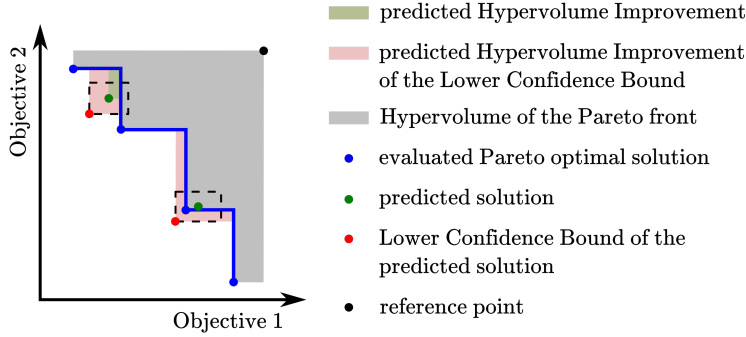


FIGURE 4.3: Illustration of the HVI and HVI_{LCB} for a bi-objective problem.

We generalised the SEVR measure to the multi-objective case and introduced the *Scaled Hypervariance Reduction* (SHVR) for the selection of the level of fidelity:

$$l = \underset{1:L}{\operatorname{argmax}} \operatorname{SHVR}_l, \quad (4.16)$$

where the SHVR is defined as:

$$\begin{aligned} \operatorname{SHVR}_1 &= \frac{1}{c_1} \prod_{j=1 \dots n_r} \rho_1^2(x_{\text{new}}) \hat{s}_{j,z_1}^2(x_{\text{new}}), \\ \text{for } l \geq 2 : \\ \operatorname{SHVR}_l &= \frac{1}{c_l} \prod_{j=1 \dots n_r} \rho_{j,l-1}^2(x_{\text{new}}) \hat{s}_{j,z_{l-1}}^2(x_{\text{new}}) + \hat{s}_{j,\delta_{z_l}}^2(x_{\text{new}}), \end{aligned} \quad (4.17)$$

The pseudo-code of the acquisition function for multi-objective problems with fidelity level selection is given in Algorithm 4.

Algorithm 4 Acquisition function for multi-objective and multi-fidelity optimisation.

Input: MFGPR_L

Output: x_{l,n_l+1}

```

1: for k=1:nr do
2:    $\hat{m}_{L,k} \leftarrow \text{predictTrend}(\text{MFGPR}_{L,k})$ 
3:    $\hat{s}_{L,k}^2 \leftarrow \text{predictVariance}(\text{MFGPR}_{L,k})$ 
4: end for
5:  $HVI_{LCB}(x) \leftarrow \text{calculateHVI}(\mathcal{P}, \hat{m}_{L,1:k}(x), \hat{s}_{L,1:k}^2(x))$ 
6:  $(x_{?,n_l+1}) \leftarrow \operatorname{argmax} HVI_{LCB}(x)$ 
7: for l=1:L do
8:    $\operatorname{SHVR}_l \leftarrow \text{calculateSHVR}(\text{MFGPR}_l, x_{?,n_l+1}, \hat{s}_l, c_l)$ 
9: end for
10:  $l \leftarrow \operatorname{argmax}(\operatorname{SHVR}_l)$ 
11: return  $x_{l,n_l+1}$ 

```

“You must be ready to give up even the most attractive ideas when experiment shows them to be wrong.”

— Alessandro Volta

5.1 Multi-fidelity benchmark functions

The mathematical formulation of the *Multi-Fidelity Gaussian Process Regression* (MFGPR) allows the combination of any level of fidelities. Here, we limit ourselves to investigate only two-level models. A test suite for multi-fidelity single-objective optimisation was proposed by Wang et al. [2017]. The test suite was extended by Habib et al. [2019] for multi-objective problems. Their work is adopted in this dissertation with minor modifications. The benchmark tests are complemented with a low-fidelity model which combines residual and stochastic errors at the same time. Moreover, the test functions are complemented with an additional probabilistic term on both high- and low-fidelity levels to emulate uncertain problems in Section 5.2.

The exact deterministic function of the underlying problem is denoted by $f(x)$. Any approximation of $f(x)$ can be generally described as the following:

$$\bar{f}(x, \phi) = f(x) + e(x, \phi). \quad (5.1)$$

where the approximation $\bar{f}(x, \phi)$ is constructed as the sum of the exact function $f(x)$ plus some error term $e(x, \phi)$. The error term depends on the accuracy level ϕ of the approximation. The error

term is the biggest when $\phi = 0$. When $\phi = 10$ the error term becomes zero and the approximation provides exact results.

The following modified Rastrigin function is considered for testing a single-objective optimisation:

$$f(x) = \left(d + \sum_{j=1}^d [(x_j - 0.5)^2 - \cos(20\pi(x_j - 0.5))] \right), \quad (5.2)$$

where the number of design variables d can be chosen arbitrarily.

For multi-objective optimisation, the DTLZ3 [Deb et al., 2005] test function is used. The number of objective functions can be chosen arbitrarily with the restriction that $n_{obj} \geq 2$ and $d \geq n_{obj}$.

$$f(x) = \text{DTLZ3}(x). \quad (5.3)$$

The error term $e(x, \phi)$ in Eq. (5.1) can take four forms: *residual*, *stochastic*, *instability error* and *residual+stochastic*. These represent the most common error types in real-life approximation models [Wang et al., 2017].

The *residual* error term takes the following form:

$$e_{res}(x, \phi) = \sum_{j=1}^d [\alpha(\phi) \cos(\omega(\phi)x_j + \beta(\phi) + \pi)], \quad (5.4)$$

$$\alpha(\phi) = (1 - 0.1\phi), \quad (5.5)$$

$$\omega(\phi) = 10\pi\alpha(\phi), \quad (5.6)$$

$$\beta(\phi) = 0.5\pi\alpha(\phi), \quad (5.7)$$

where d is the number of deterministic design variables.

The *stochastic* error term takes the following form:

$$e_{stoc}(x, \phi) = \mathcal{N}(\mu_\epsilon, \sigma_\epsilon), \quad (5.8)$$

$$\mu_\epsilon(x, \phi) = \frac{1}{10d} \nu(\phi) \gamma(x), \quad (5.9)$$

$$\sigma_\epsilon(\phi) = 0.1\nu(\phi), \quad (5.10)$$

$$\nu(\phi) = 1 - 0.1\phi, \quad (5.11)$$

$$\gamma(x) = \sum_{j=1}^d (1 - |x_j|), \quad (5.12)$$

The *instability* error term takes the following form:

$$e_{ins}(\phi) = \begin{cases} 10d, & \text{if } r \leq \sigma_\epsilon(\phi), \\ 0, & \text{if } r > \sigma_\epsilon(\phi), \end{cases} \quad (5.13)$$

where r is a random number between 0 and 1.

The *residual+stochastic* error term is simply the superposition of the *residual* and *stochastic* terms.

5.2 Multi-fidelity test problems under uncertainty

To emulate uncertain problems, the single-objective test function is extended by introducing a probabilistic term $g(x, u, \phi_u)$ for both fidelity levels:

$$f(x, u, \phi_u) = f(x) + g(x, u, \phi_u), \quad (5.14)$$

$$\bar{f}(x, u, \phi, \phi_u) = f(x) + g(x, u, \phi_u) + e(x, \phi), \quad (5.15)$$

where both the high-fidelity $f(x, u, \phi_u)$ and the low-fidelity test function $\bar{f}(x, u, \phi, \phi_u)$ depend on controllable design x and uncontrollable uncertain variables u . The level of stochasticity is determined by ϕ_u . When $\phi_u = 0$ the problem reduces to the previously described deterministic case.

The probabilistic term $g(x, u, \phi_u)$ depends on both design and uncertain variables:

$$g(x, u, \phi_u) = \sum_{i=1}^{d_u} \left[\alpha(\phi_u) \cosh \left(\frac{1}{10} (2 - \gamma(x)) \pi u_i \right) \right], \quad (5.16)$$

$$\alpha(\phi_u) = 0.1 \phi_u, \quad (5.17)$$

$$\gamma(x) = \sum_{j=1}^d (1 - |x_j - x_j^*|), \quad (5.18)$$

where d_u is the number of uncertain input parameters and $x_j^* = 0.4$ is the value of the j^{th} design vector where the uncertainty is smallest. The deterministic function $f(x)$ is the same as described previously by Eq. (5.2).

To extend the multi-objective *DTLZ3* function with uncertainty, the $g(u, x, \phi_u)$ term is added to the each objective resulting in the following equations:

$$f_1(x, u, \phi_u) = \cos(x_1 \frac{\pi}{2})(1 + \rho(x_M)) + g(x, u, \phi_u), \quad (5.19)$$

$$f_2(x, u, \phi_u) = \sin(x_1 \frac{\pi}{2})(1 + \rho(x_M)) + g(x, u, \phi_u), \quad (5.20)$$

$$\bar{f}_k(x, u, \phi, \phi_u) = f_k(x) + g(x, u, \phi_u) + e(x, \phi), \quad (5.21)$$

$$\rho(x_M) = \left(d_M + \sum_{j=1}^{d_M} [(x_j - 0.5)^2 - \cos(20\pi(x_j - 0.5))] \right), \quad (5.22)$$

where 2 objectives are considered. f_1 and f_2 are the high-fidelity objectives and \bar{f}_k is the k^{th} low-fidelity objective.

5.3 Comparison of single- and multi-fidelity surrogates

When to use multi-fidelity surrogates is still an open question [Fernández-Godino et al., 2019]. Multi-fidelity models are more complex than their single-fidelity counterparts. Hence, the comparison between the two techniques is not straightforward. The MFGPR has more parameters than the single-fidelity *Gaussian Process Regression* (GPR) and therefore the relative performance of the two techniques depends on the extra parameters of the MFGPR. In [Wang et al., 2017], the shape and correlation properties of multi-fidelity benchmarks were investigated. In this section, we discuss the surrogate quality of MFGPR on the proposed benchmark functions. The parameter settings of the test cases are described in Table 5.1.

Table 5.2 compares the accuracy of GPR and MFGPR considering the different error types of the *Low-Fidelity* (LF) level defined in Section 5.1. The accuracy measure $\text{Error}_{\text{avg}}$ is the mean absolute error calculated on 2000 points generated with a *Latin Hypercube Sampling* (LHS). In these simulations, the computational budget was fixed to 300 units, which corresponds exactly to 30 *High-Fidelity* (HF) samples. The cost of a high-fidelity evaluation is 10 and the cost of a low-fidelity evaluation is 5. Moreover,

Table 5.1: List of test cases for deterministic single-objective problems. (c_{HF} : cost of high-fidelity evaluation, c_{LF} : cost of low-fidelity evaluation)

	High-fidelity		Low-fidelity		
Test	$f(\mathbf{x})$	c_{HF}	$e(\mathbf{x}, \phi)$	ϕ	c_{LF}
TD0	Eq.(5.2)+Eq.(5.16)	10	-	-	-
TD1	Eq.(5.2)+Eq.(5.16)	10	Eq.(5.4)	5	5
TD2	Eq.(5.2)+Eq.(5.16)	10	Eq.(5.8)	5	5
TD3	Eq.(5.2)+Eq.(5.16)	10	Eq.(5.13)	5	5
TD4	Eq.(5.2)+Eq.(5.16)	10	Eq.(5.4)+Eq.(5.8)	5	5
Input					
x	$0 \leq x \leq 1, \quad d = 1$				

Table 5.2: Comparison of single- and multi-fidelity GPR for deterministic problems.

Test	Budget	LF error type	Surrogate	Samples	Error _{avg}
TD0	300	-	GPR	30(HF)	1.522e-1
TD1	300	residual	MFGPR	15(HF) 30(LF)	5.480e-2
TD2	300	stochastic	MFGPR	15(HF) 30(LF)	4.419e-2
TD3	300	instability	MFGPR	15(HF) 30(LF)	1.545e-1
TD4	300	residual + stochastic	MFGPR	15(HF) 30(LF)	5.930e-2

for the MFGPR model the computational budget was divided into two equal parts, i.e. 150 and 150, to obtain high- and low-fidelity samples.

Figure 5.1(a) shows that the landscape cannot be accurately modelled by a standard GPR because of the limited number of samples. By obtaining information from low-fidelity models a better approximation can be achieved as shown in Figure 5.1(b)-(e). The MFGPR model uses half HF samples w.r.t. classical GPR, but the high number of LF samples provides more valuable information about the landscape. We can see that the error type of the low-fidelity model can have a significant effect on surrogate quality. The MFGPR performs the best when the LF model has only stochastic errors and performs the worst when the LF model has instabilities. GPR is able to handle instabilities only up to the magnitude of the regularising term σ_ϵ in Eq. (2.5), which is also called nugget, whereas higher instabilities can cause severe approximation errors. In a MFGPR, the low-fidelity level approximation error is propagated to the higher fidelity level.

In the previous simulations, the total computational budget was fixed to 300. The size of the computational budget has a

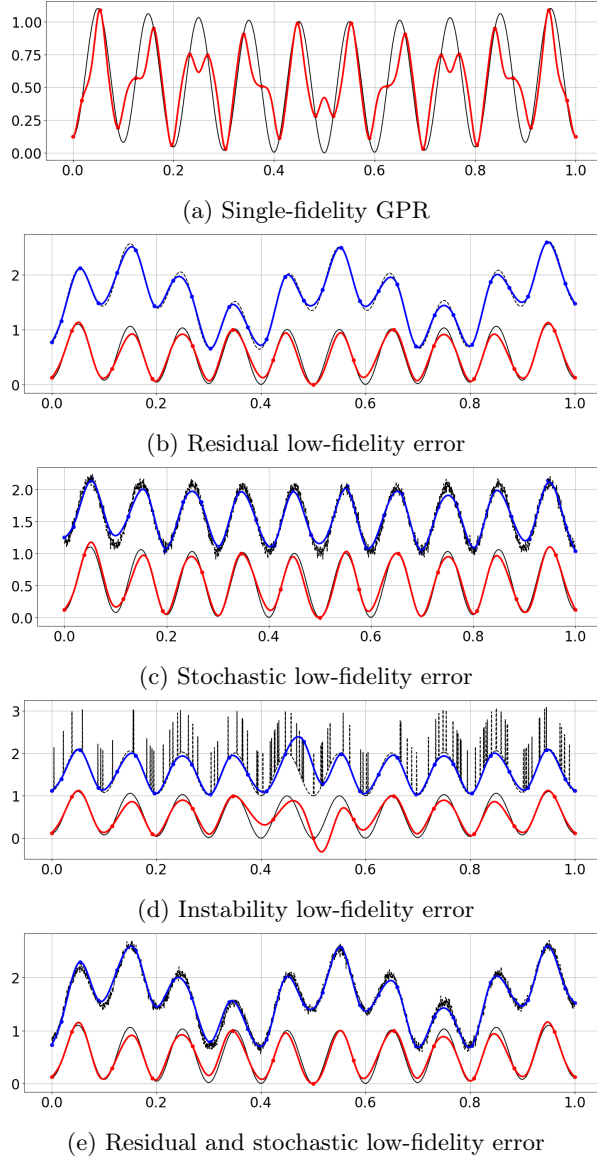


FIGURE 5.1: Landscapes of the single- and multi-fidelity surrogate models. (In the plots for the sake of better visibility the low-fidelity curves are shifted up by 1. The true high- and low-fidelity functions are drawn with solid and dashed black lines respectively. The high- and low-fidelity model predictions are drawn with red and blue respectively.)








FIGURE 5.2: Comparison of single- and multi-fidelity GPR for single-objective deterministic case. The lines represent the absolute errors of the surrogate models calculated at 2000 locations sampled by latin hypercube sampling. The black dashed line corresponds to the GPR model; the three differently shaded red lines correspond to the three error levels of the low-fidelity model.

significant effect on surrogate quality. Figure 5.2 shows that the relative performance of GPR and MFGPR changes depending on the size of the total budget. When the budget is too small, the available samples are not even enough to approximate well the low-fidelity model. This causes an even higher approximation error at the top-fidelity level. In this case, the samples from the low-fidelity model introduce a larger error than the amount of information they provide. Similarly, when the budget is big enough to have enough HF samples to accurately approximate the landscape, the low-fidelity deteriorate the multi-fidelity approximation. Consequently, in these domains, MFGPR is not useful. MFGPR provides significantly better results when the available budget is not able to provide enough HF samples. In this case, instead of a handful of HF samples, many LF samples can be generated which can provide valuable information for the landscape approximation. Consequently, MFGPR is recommended to be used for such a problem setting.

5.4 Acquisition step of a multi-fidelity surrogate-based optimisation

This section illustrates the internal processes of the proposed strategy through its application to the single-objective deterministic

multi-fidelity test-case represented by Eq. (5.2). To illustrate Algorithm 1 and Algorithm 2, the progression from iteration t to $t + 1$ is shown in Figure 5.3(a)-(e). In Figure 5.3(a) the predicted low- and high-fidelity landscapes are plotted. At iteration t the MFGPR was trained on 15 HF and 31 LF samples. The *Expected Improvement* (EI) of this prediction was calculated resulting in the black curve plotted in Figure 5.3(b). The maximum of the EI was located at $x = 0.405$. At this location, the *Scaled Expected Variance Reduction* (SEVR) of each fidelity level were calculated and compared. The low cost of the LF evaluation resulted in a higher SEVR. Therefore, a LF fidelity sample was generated and the MFGPR surrogate was retrained. The new landscape prediction at $t + 1$ is shown in Figure 5.3(c). Figure 5.3(d) and (e) shows the proportions of the total budget spent on low- and high-fidelity samples at iteration t and $t + 1$ respectively.


 (a) Landscape at iteration t

 (b) Acquisition at iteration t

 (c) Landscape at iteration $t + 1$

 (d) Budget at t

 (e) Budget at $t + 1$

FIGURE 5.3: Iteration step of multi-fidelity EGO

5.5 Benchmarking the single-objective multi-fidelity surrogate-based optimisation under uncertainty

The parameter settings of the single-objective test problems under uncertainty are summarised in Table 5.3. The *residual+stochastic* error term is considered for the study in this section.

The evolution of the objective value obtained with the classic GPR-based *Efficient Global Optimisation* (EGO) algorithm and the proposed multi-fidelity surrogate-based approach is compared on Figure 5.4. The computational budgets are equal in the two optimisation runs. This results in a higher number of HF samples for classical GPR. The higher number of HF samples provides a better starting point for the optimisation. However, the limited budget causes a premature stop of the optimisation algorithm. When MFGPR is employed in the EGO algorithm, many low-fidelity samples can be generated to explore the design landscape and then HF samples are mostly generated to exploit the identified region of interest.

The obtained objectives with different budgets confirm the findings made with the deterministic simulations of the previous subsection. Table 5.4 shows that when the budget is low (2000) then the GPR is not able to provide an accurate approximation.

Table 5.3: List of test cases for uncertain single-objective problems

Test	High-fidelity			Low-fidelity		
	$f(\mathbf{x})$	ϕ_u	c_{HF}	$e(\mathbf{x}, \phi)$	ϕ	c_{LF}
TU0	Eq.(5.2)+Eq.(5.16)	-	10	-	-	-
TU1	Eq.(5.2)+Eq.(5.16)	5	10	Eq.(5.4)+Eq.(5.8)	1	5
TU2	Eq.(5.2)+Eq.(5.16)	5	10	Eq.(5.4)+Eq.(5.8)	5	5
TU3	Eq.(5.2)+Eq.(5.16)	5	10	Eq.(5.4)+Eq.(5.8)	9	5
Input						
\mathbf{x}	$0 \leq x_j \leq 1, \quad j = 1..d, \quad d = 1$					
\mathbf{u}	$u_j \sim \mathcal{N}(0, 1), \quad j = 1..d_u, \quad d_u = 1$					
\mathcal{R}						
$\bar{q}_{0.95}$	superquantile Eq. (3.5)					




FIGURE 5.4: Comparison of single- and multi-fidelity surrogate-assisted optimisation convergence. (The convergence histories were obtained by a trial run of TU2 with a budget of 4000.)

For this reason, the strategy employing MFGPR outperforms the one with GPR. We can also notice that by increasing the fidelity of the low-level function $\phi = 5$ and $\phi = 9$, the obtained objective becomes worse. The reason for this is that the budget is fixed and by increasing the fidelity of the low-level function also its cost increases, hence the number of LF samples becomes also sparse. Therefore, the advantage of employing MFGPR is diminished; however, the total number of samples is still higher than in the GPR case resulting in a slightly better objective value.

When the budget is increased to a level that the GPR is able to provide an accurate approximation, there is no advantage to use MFGPR. In the cases of fidelity levels $\phi = 1$ and $\phi = 5$, the obtained objective values are comparable with both GPR and MFGPR. However, in the case of $\phi = 9$ the low-fidelity level with sparse samples clearly provides a bad approximation which affects the quality on the top level as well and results in a significantly worse objective value.

5.6. Benchmarking the multi-objective multi-fidelity surrogate-based optimisation under uncertainty

Table 5.4: Comparison of single- and multi-fidelity surrogate-assisted optimisation results. The values in the table are averaged over 10 independent runs.

Test	Budget	MFGPR			GPR (TU0)	
		ϕ	Samples	Obj.	Samples	Obj.
TU1	2000	1	126(<i>HF</i>) 711(<i>LF</i>)	0.360	200 (HF)	0.425
TU2	2000	5	138(<i>HF</i>) 124(<i>LF</i>)	0.38290	200 (HF)	0.425
TU3	2000	9	60(<i>HF</i>) 150(<i>LF</i>)	0.38291	200 (HF)	0.425
TU1	4000	1	229(<i>HF</i>) 1686(<i>LF</i>)	0.363	400 (HF)	0.352
TU2	4000	5	184(<i>HF</i>) 432(<i>LF</i>)	0.364	400 (HF)	0.352
TU3	4000	9	220(<i>HF</i>) 190(<i>LF</i>)	0.38247	400 (HF)	0.352

5.6 Benchmarking the multi-objective multi-fidelity surrogate-based optimisation under uncertainty

The results of the multi-objective problem are in agreement with the results obtained for the single-objective case. The parameter settings of the investigated test functions are summarised in Table 5.5.

The advantage of employing a MFGPR is significant when the approximation on the low-fidelity is sufficiently good as shown in Table 5.6. When the low-fidelity level is prone to instability error the obtained *Inverted Generational Distance* (IGD) [Zitzler et al., 2003] values are worse compared to when the low-fidelity has only residual and stochastic error. Similarly to the single-objective tests increasing the accuracy of the low-fidelity level might decrease the overall performance of the algorithm as in the

Table 5.5: List of test cases for uncertain multi-objective problems

	High-fidelity			Low-fidelity		
Test	$f(x)$	ϕ_u	c_{HF}	$e(x, \phi)$	ϕ	c_{LF}
TMU0	Eq.(5.19)- (5.20)	-	20	-	-	-
TMU1	Eq.(5.19)- (5.20)	5	20	Eq.(5.4)+Eq.(5.8)	5	2.5
TMU2	Eq.(5.19)- (5.20)	5	20	Eq.(5.13)	5	2.5
TMU3	Eq.(5.19)- (5.20)	5	20	Eq.(5.4)+Eq.(5.8)	9	5
Input						
x	$0 \leq x_j \leq 1, \quad j = 1, 2$					
u	$u_j \sim \mathcal{N}(0, 1), \quad j = 1, 2$					
\mathcal{R}						
$\bar{q}_{0.95}$	superquantile Eq. (3.5)					

TMU3 case with a computational budget of 6000 units. At any fidelity level, the accuracy can be increased typically by increasing the computational cost too, which can result in a sparse sample on the low-fidelity level. As expected, the IGD values decrease with an increasing budget. Similarly to the single-objective case, when the number of samples is high enough the GPR model approximation becomes satisfactorily good; hence, employing the MFGPR surrogate model is unnecessary.

5.6. Benchmarking the multi-objective multi-fidelity surrogate-based optimisation under uncertainty

Table 5.6: Comparison of single- and multi-fidelity surrogate-assisted multi-objective optimisation under uncertainty. The IGD values of the obtained Pareto fronts are calculated with a reference set of 100 points. The values in the table are averaged over 10 independent runs.

		MFGPR			GPR (TMU0)	
Test	Budget	ϕ	Samples	IGD	Samples	IGD
TMU1	6000	5	231(<i>HF</i>) 516(<i>LF</i>)	0.172	300 (HF)	0.303
TMU2	6000	5	230(<i>HF</i>) 540(<i>LF</i>)	0.233	300 (HF)	0.303
TMU3	6000	9	215(<i>HF</i>) 409(<i>LF</i>)	0.209	300 (HF)	0.303
TMU1	8000	5	273(<i>HF</i>) 995(<i>LF</i>)	0.164	400 (HF)	0.203
TMU2	8000	5	293(<i>HF</i>) 826(<i>LF</i>)	0.182	400 (HF)	0.203
TMU3	8000	9	264(<i>HF</i>) 658(<i>LF</i>)	0.163	400 (HF)	0.203
TMU1	10000	5	319(<i>HF</i>) 1421(<i>LF</i>)	0.175	500 (HF)	0.160

“Scientists study the world as it is, engineers create the world that never has been.”

— Theodore von Karman

6.1 Introduction to airfoil design

Shape optimisation of an airfoil is one of the most fundamental problems in aerodynamic design optimisation. The purpose of an airfoil is to generate a pressure difference in a flow so that a force is generated. The force component perpendicular to the flow direction is called lift, its magnitude and sense (in respect of the defined force reference frame) depend on the shape of the airfoil and on the flow conditions. Together with the lift, the presence of the airfoil in the flow inevitably generates a force component parallel to the flow direction, called drag. The aerodynamic forces are depicted in Figure 6.1. Most engineering applications exploit the lift, while the drag is an inevitable loss. Therefore, the shape optimisation of an airfoil aims to find the optimal airfoil shape such that the application-dependent requirements are satisfied.

Conceptually, aerodynamic design optimisation can be approached in two ways: inverse and direct [Song and Keane, 2004]. In inverse design optimisation, a desired pressure distribution is targeted, and the optimisation algorithm seeks to find the geometrical shape which produces the targeted pressure distribution or lift force. Therefore, inverse methods are commonly applied in later design phases when target values are known. A direct method, as

its name suggests, directly optimises the objective without any predefined target value. In aerodynamic shape optimisation, a maximal lift-to-drag ratio is typically desired.




FIGURE 6.1: Aerodynamic forces of an airfoil

One of the challenges of aerodynamic shape optimisation is that accurate *Computational Fluid Dynamics* (CFD) calculations are typically expensive in computational time [Poloni and Pediroda, 1997]. This issue can be tackled by employing surrogate models as in [Poloni et al., 2000; Poloni, 1995]. Expensive calculations are performed for only a handful of designs. Then a statistical model is built to approximate the aerodynamics of airfoil designs which have not been evaluated by the expensive CFD code. The accuracy of the statistical model highly depends on the number of available CFD evaluations. Consequently, sparsely sampled design landscapes are hard to approximate accurately with standard surrogate techniques. In such a case, aerodynamic calculations of lower fidelity can be used to provide sufficient information for building an accurate statistical model. The information from low-fidelity calculations can be fused together with high-fidelity data by using the *Multi-Fidelity Gaussian Process Regression* (MFGPR) model.

There is also another issue to take into account. Often the actual design, or operation point, and its performance are slightly different from the optimisation solutions because of manufacturing, wear off [Sareen et al., 2014] and other operational deformations, like icing [Arizmendi et al., 2019] and surface pollution [Bak et al., 2008]. In practice, our design problem is affected by various uncertainty sources which affect the actual aerodynamic performance. This issue can be addressed with *Uncertainty Quantification* (UQ) techniques and optimisation under uncertainty methods. UQ can be used to estimate statistical measures of the design performance

that can be in turn used as reliable or robust objectives of the optimisation under uncertainty problem (see e.g. the reviews [Beyer and Sendhoff, 2007; Schuëller and Jensen, 2008]).

6.2 Airfoil shape optimisation: uncertainties, reliability-based risk measure and other details

When simulating a flow around an airfoil only a limited number of phenomena are modelled. Therefore, the aerodynamic performance of a real airfoil might deviate from the numerical results. This motivates the construction of probabilistic models which can be appended to the design optimisation workflow to predict the variations of the aerodynamic performance.

In this study, only probabilistic uncertainties are considered. The uncertainties are propagated through the aerodynamic solver which will result in a probabilistic aerodynamic performance. The comparison of two probability distributions is not a trivial task. The possible realisations of a distribution are multiple [Rockafellar and Royset, 2015]. Therefore, only certain properties of a distribution are compared as discussed in Section 3.1. In this section, the 95th superpercentile of the C_l and C_d distributions are used for evaluating the aerodynamic performance of the airfoil. This risk measure is defined in Equation (3.5) and employed here to ensure reliability. The advantages of the superpercentile measure over other risk measures for engineering applications are discussed in [Rockafellar and Royset, 2015; Quagliarella et al., 2020].

An analytical propagation of the uncertainty is not possible due to the complex aerodynamics solvers. Therefore, the superpercentile values of the C_l and C_d distributions are calculated using empirical values obtained by sampling. To obtain a sufficient number of samples, surrogate assisted uncertainty quantification is performed. The probabilistic space is considered to be independent from the design space. For each design, a local *Gaussian Process Regression* (GPR) model is trained over a small sample randomly generated around the considered design. Then, the calculation of the superpercentile is done on a larger sample evaluated with the trained GPR model.

6.3 Solvers for aerodynamic force prediction of an airfoil


To calculate the aerodynamic forces of the airfoil two solvers are considered: XFOIL [Drela, 1989] and SU2 [Economon et al., 2016].

The former is an airfoil analysis tool based on potential flow equations (panel methods). For viscous problems, a two-equation integral boundary layer formulation is coupled with the inviscid flow solution [Drela, 2001]. The transition criterion is calculated by the e^N envelope method. XFOIL has a fairly rapid calculation time and provides sufficient accuracy for most engineering applications.


SU2 can be used to solve the compressible *Reynolds-averaged Navier-Stokes* (RANS) equations. The RANS equations are closed by Menter’s Shear Stress Transport turbulence model [Menter et al., 2003] which efficiently blends the $k-\omega$ turbulence model of the near-wall region to the $k-\epsilon$ model of the region away from any wall. Various studies have been carried out to compare the results of these two solvers, like e.g. [Vaithiyanathasamy, 2017; Barrett and Ning, 2016]. Both solvers are suitable to accurately predict the aerodynamic forces of an airfoil. For the sake of this study, SU2 is considered as a higher fidelity solver as it implements a more general form of the Navier-Stokes equations.

The aerodynamic evaluations with XFOIL and SU2 are performed with the framework software described in [Quagliarella et al., 2019; Quagliarella and Diez, 2020]. The modal shape function superposition is performed with the *wg2aer* software¹ [Quagliarella, 2016]. The modified airfoil geometry is stored in a *Selig* format which can be directly processed by XFOIL. For the CFD evaluation, the modified airfoil and its surrounding domain are discretised with the open-source *Gmsh* software which generates the mesh in *.su2* format. Finally, SU2 performs the aerodynamic analysis of the airfoil and provides the high-fidelity drag and lift predictions.

¹Software developed at the Italian Aerospace Research Centre (CIRA)



(a) Lift coefficient against angle-of-attack



(b) Drag coefficient against angle-of-attack

FIGURE 6.2: Comparison of lift and drag coefficient curves of MH114 calculated with XFOIL and SU2.

The lift and drag coefficients of the MH114 airfoil, which is taken as the baseline in our airfoil optimisation study, are plotted in Figures 6.2a and 6.2b. The calculations are carried out at Reynolds number $Re = 5 \cdot 10^6$ and Mach number $M = 0.218$ with standard sea-level conditions. SU2 considers a fully turbulent domain around the airfoil. Therefore, XFOIL was also forced to operate in fully turbulent conditions by setting the transition point location at the beginning of the lower and upper airfoil sides (XTRLO and XTRUP set to 0.01). The two solvers produce similar polar trends; however, there are some deviations in the actual values. This makes the two solvers appropriate candidates for a multi-fidelity optimisation.

6.3.1 Parametric geometry definition of airfoils

The geometrical shape of the airfoil is defined by superposing modal shape functions on the baseline geometry of MH114. Eight modal shape functions are considered. They are shown in Figure 6.3 and Table 6.1. The first two modes modify the thickness and the camber line of the airfoil. The remaining six modes introduce local shape modifications of the upper and lower side of the airfoil at the *Leading Edge* (LE), mid-span and *Trailing Edge* (TE) respectively. The design variables to be optimised are the scaling parameters (x_i) of the modal shape functions.

Table 6.1: Design variables of the airfoil.

Mode	Function type	x_i	Physical interpretation
mode 1	Polynomial	x_1	thickness
mode 2	Polynomial	x_2	camber
mode 3	Hicks-Henne	x_3	upper LE
mode 4	Hicks-Henne	x_4	lower LE
mode 5	Hicks-Henne	x_5	upper middle
mode 6	Hicks-Henne	x_6	lower middle
mode 7	Hicks-Henne	x_7	upper TE
mode 8	Hicks-Henne	x_8	lower TE

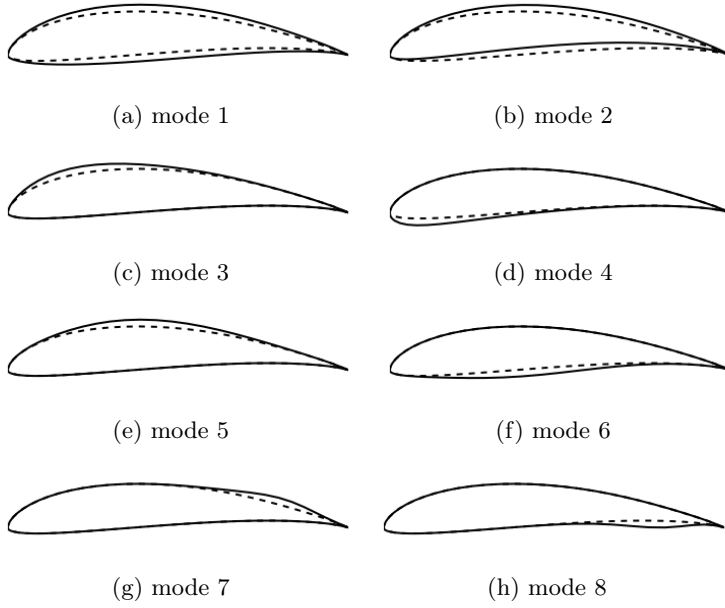
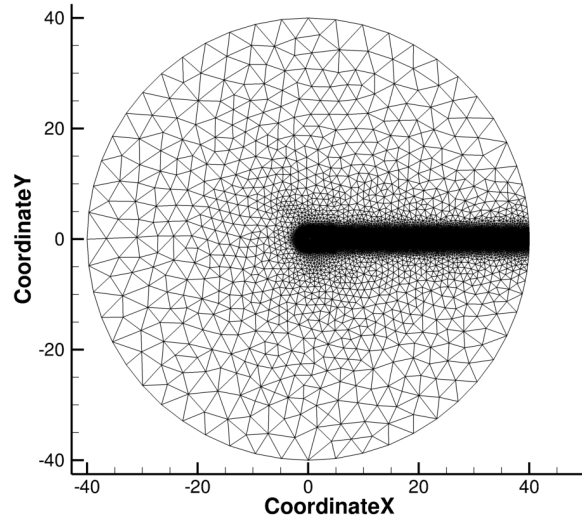



FIGURE 6.3: Baseline design and the deformation modes of the airfoil.

6.3.2 Mesh generation of airfoils


The parametric mesh of the airfoil designs is generated by the in-house software of CIRA described in [Quagliarella et al., 2019; Quagliarella and Diez, 2020]. The domain around the airfoil is discretised with a hybrid mesh. The boundary layer is meshed with a fine quadrilateral grid and, moving away from the boundary layer, the domain is tessellated with triangles of increasing size. An example is given in Figure 6.4. The depicted mesh consists of 96085 vertices and forms 87402 triangles and 48513 quadrilaterals. The parametric mesh size is determined for each case study following the suggestions of the software described in [Quagliarella et al., 2019; Quagliarella and Diez, 2020].




(a) Far-field mesh



(b) Near-field mesh (zoom on airfoil)



(c) Leading edge mesh (zoom)




(d) Trailing edge mesh (zoom)

6.3.3 Far-field drag prediction of coarse grids


Alternatively, different levels of accuracy can be achieved by using different mesh sizes for the RANS simulations. In a collaboration work [Paper C: Morales and Korondi et al.,2020], we detailed how the mesh size influences the aerodynamic force calculation. RANS simulations with coarse meshes provide drag predictions loaded with spurious drag. The numerical dissipation and truncation error become larger when coarser meshes are used. To compensate the effect of these numerical errors the drag can be predicted with the so-called far-field formula. This formula proved its efficiency for solving optimisation problems under uncertainty [Morales et al., 2019]. Nevertheless, the far-field drag prediction formula requires a fully structured grid. The structured grids are generated with the *Construct2D* open-source software [Construct2D, 2018]. The parameters of the used meshes are given in Table 6.2. The representation of a coarse and fine grid is given in Figures 6.5a and 6.5b respectively.

Table 6.2: Mesh size parameters for low- and high-fidelity simulations. (N_b : number of cells on the body surface, N_w : number of cells in the wake, N_j : number of cells in far-field direction, N_{total} : total number of cells) [Paper C: Morales and Korondi et al.,2020]

	N_b	N_w	N_j	N_{total}
Low-fidelity (LF)	96	48	48	16384
High-fidelity (HF)	512	256	256	262144



(a) Coarse structured mesh



(b) Fine structured mesh

FIGURE 6.5: Structured meshes of different size.

6.4 Multi-objective airfoil optimisation under uncertainty with XFOIL and SU2

In this section, the MH114 airfoil is optimised for a high-lift propeller. We seek for a set of Pareto optimal geometries for maximum lift (L) and minimum drag (D). Additionally, the shape of the airfoil is considered to have some geometrical uncertainties due to the manufacturing process. Therefore, each shape mode is additionally superposed on the design shape with an uncertain scaling factor (u_i). The uncertain scaling factors can be described by a Gaussian distribution with zero mean and 0.01 standard deviation. The complete list of the considered uncertain variables is reported in in Table 6.3.

As the lift and drag forces are uncertain due the geometrical uncertainties, a reliability-based multi-objective optimisation problem is considered here:

$$\min_x \mathcal{S}_{95}^{right}(-C_l(x, \tilde{u})), \quad (6.1a)$$

$$\min_x \mathcal{S}_{95}^{right}\left(\frac{C_d(x, \tilde{u})}{C_{d0}}\right), \quad (6.1b)$$

where $C_{d0} = 0.01$. The \mathcal{S}_{95}^{right} denotes the right-tail 95th super-percentile which is a reliability measure defined in Eq. (3.5) in

Table 6.3: Geometrical uncertainties of the multi-objective airfoil design problem.

Modes	Function type	u_i	Physical interpretation
mode 1	Polynomial	u_1	thickness uncertainty
mode 2	Polynomial	u_2	camber uncertainty
mode 3	Hicks-Henne	u_3	upper LE uncertainty
mode 4	Hicks-Henne	u_4	lower LE uncertainty
mode 5	Hicks-Henne	u_5	upper middle uncertainty
mode 6	Hicks-Henne	u_6	lower middle uncertainty
mode 7	Hicks-Henne	u_7	upper TE uncertainty
mode 8	Hicks-Henne	u_8	lower TE uncertainty

Section 3.1. The lift coefficient (C_l) and drag coefficient (C_d) are:

$$C_l = \frac{L(x, \tilde{u})}{1/2\rho U^2}, \quad (6.2a)$$

$$C_d = \frac{D(x, \tilde{u})}{1/2\rho U^2}, \quad (6.2b)$$

where ρ is the density and U is the free-stream velocity of the air. Compared to the traditional C_l and C_d coefficient definitions, the chord length is omitted from the above definitions as the chord length is considered as unit throughout this section.

6.4.1 Workflow for multi-objective airfoil design optimisation

Here, we employ the multi-objective optimisation workflow which was introduced in Section 4.4 and published in [Paper A: Korondi et al., 2020a]. The workflow embodies a multi-fidelity Bayesian optimisation for multi-objective problems. In this section, it is employed for the aerodynamic design optimisation of airfoils using XFOIL and SU2.

To initialise the optimisation workflow a *Design of Experiments* (DoE) of the design space is generated at each fidelity level with a uniform *Latin Hypercube Sampling* (LHS). After obtaining the 95th superpercentiles of the lift and drag of the DoE samples, two independent multi-fidelity surrogates are trained using the recursive formulation defined by Eq. (2.12). The two MFGPR models are used by the acquisition function to determine which design configuration should be evaluated in the next iteration and which solver should perform the aerodynamic calculation. Since it is a multi-objective problem, the decision on the next design location is made by maximising the *Hypervolume Improvement of the Lower Confidence Bound* (HVI_{LCB}) of the drag and lift coefficients. The selection of the solver is based on the *Scaled Hypervariance Reduction* (SHVR) values of the fidelity levels. The superpercentile values of the drag and lift of the selected design is evaluated at the selected fidelity with the surrogate-assisted UQ approach discussed in Section 3.2.

With the new superpercentile values the surrogate model of the lift and drag is retrained. The surrogate is updated iteratively with new designs until the computational budget is exhausted.

For each design configuration (during the DoE or the optimisation phase), a sample of 15 elements in the probabilistic space is generated with a uniform LHS around the selected design to train a local GPR model. The trained model is used to evaluate a larger virtual sample of 5000 designs of the local probabilistic space and compute the 95th superpercentiles. The chosen sizes seemed to be adequate after performing a few trial and error checks.

At the end of the optimisation, the set of Pareto optimal designs are presented to decision-makers. The complete optimisation workflow is illustrated in Figure 6.6.

6.4.2 Results of multi-objective airfoil design optimisation

A brief summary of the solved optimisation problem is presented in Table 6.4. The problem is bi-objective and has 8 design and 8 uncertain variables. The computational budget is set to 136500 units. The number of evaluated *Low-Fidelity* (LF) and *High-*

Table 6.4: Summary of the multi-objective airfoil design optimisation.

Number of objectives	2
Number of constraints	0
Number of design variables	8
Number of uncertain variables	8
Computational budget	136500

Fidelity (HF) samples and their cost are presented in Table 6.5. We assigned 300 cost units for running a single evaluation of aerodynamics forces with SU2 and 1 unit for evaluating the design with XFOIL. Here, we determined the cost of the fidelity levels based on the actual running times of the simulations on the used machine. The optimisation stopped when no further high-fidelity samples could be added to the surrogate training set. After the 435th HF simulation, only 4245 units remained in the budget which is not sufficient for generating the required 15 samples to

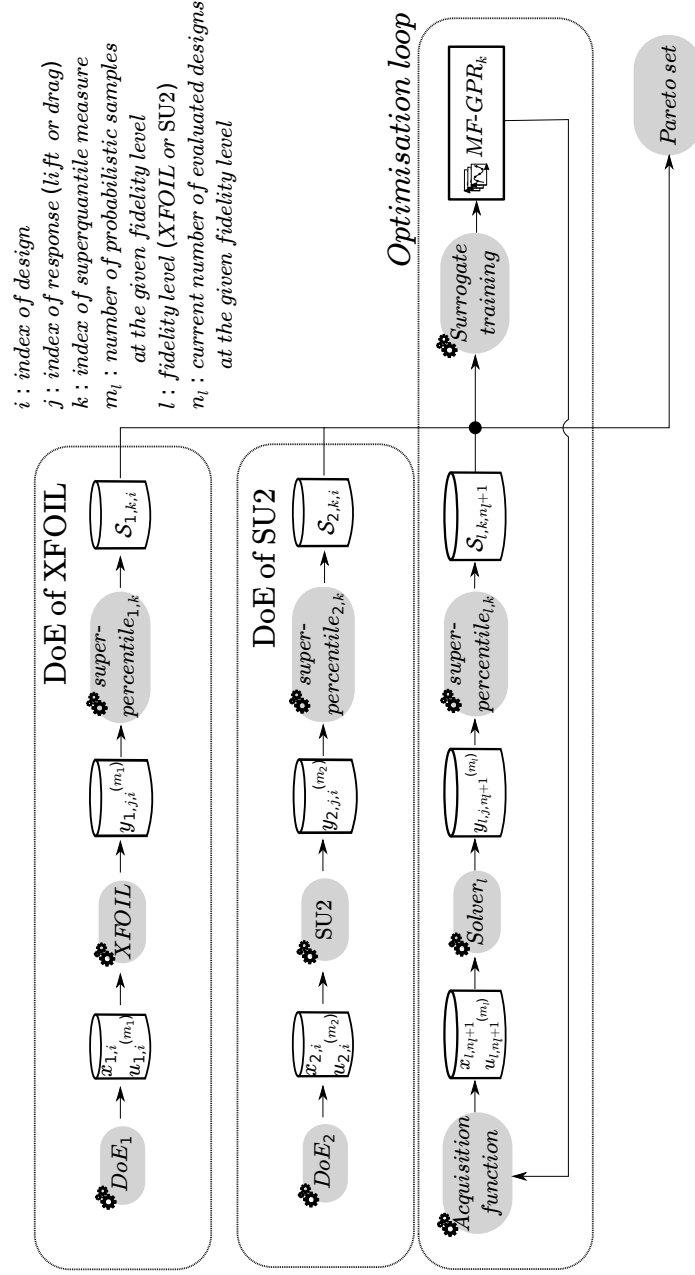


FIGURE 6.6: Multi-objective probabilistic optimisation workflow for aerodynamic shape optimisation of an airfoil with MF-GPR.

6.4. Multi-objective airfoil optimisation under uncertainty with XFOIL and SU2

Table 6.5: Multi-objective airfoil design optimisation: number of LF and HF samples and their costs.

	LF DoE	LF total	HF DoE	HF total
Solver	XFOIL	XFOIL	SU2	SU2
Evaluation cost	1	1	300	300
Total samples	450	1755	225	435
Budget spent	450	1755	67500	130500
Budget percentage	0.3%	1.2%	51.7%	95.6%

train the probabilistic model. Therefore, only 96.9 % of the budget was used.

In Figure 6.7, it is shown that the algorithm alternates the fidelity levels. The alternation stems from the fact that, in regions where the expected improvement is high due to large uncertainties, the algorithm will evaluate the new design with the low-fidelity solver. Following this step, high expected improvement values in the region are the results of promising performance prediction with low-level of uncertainty. Therefore, that region is sampled by high-fidelity simulation without risking a waste of budget.




FIGURE 6.7: History of fidelity selection throughout the multi-objective design optimisation of the airfoil.

To investigate the advantage of MFGPR over single-fidelity GPR, the relative prediction errors of the surrogates are calculated in every iteration when a HF sample is generated. The classical GPR model is trained in every iteration using only the HF samples of the actual iteration. Overall, the MFGPR provides a better prediction; however, in some iterations, GPR can temporarily be a better predictor as shown in Figure 6.8. At each iteration, the prediction error is calculated based on a single sample. When the

newly evaluated design lies in a region which can be accurately predicted by a GPR model using only HF samples, it is possible that the single-fidelity GPR model provides a slightly better prediction. However, in the majority of the iterations MFGPR outperforms the GPR model.




FIGURE 6.8: Comparing relative prediction errors of MFGPR and GPR in every iterations when a HF evaluation was performed

In Table 6.6 the mean prediction error of the single- and multi-fidelity surrogates are shown. Overall, both surrogate models can provide a relatively accurate prediction of the objectives. However, the single-fidelity surrogate model shows significantly bigger prediction errors when the newly evaluated design lies in a region which has not been explored by a sufficient number of observations.

Table 6.6: Mean prediction error of the lift and drag reliability measures in the multi-objective airfoil design problem.

	GPR	MF-GPR
Objective 1 (lift)	0.759 %	0.313 %
Objective 2 (drag)	2.191 %	0.989 %




FIGURE 6.9: Correlation history of objectives over all iterations

Throughout the iterations, the correlation between the predictions at low- and high-fidelity levels are high and steady as shown in Figure 6.9. The high correlation is expected at the beginning of the optimisation as the high number of LF samples in the training data results in a MFGPR model which predicts performances close to the LF observations. This correlation does not deteriorate by updating the model with high-fidelity samples which suggests that the initial MFGPR dominated by the LF samples provides a good approximation.

The obtained Pareto optimal solutions are depicted in Figure 6.10 (red circles, Pareto front HF). The initial HF Pareto front obtained after the DoE (dash-dotted grey line) was significantly improved. The MFGPR models of the objectives can provide accurate predictions; hence, most of the design locations suggested by the acquisition function are Pareto optimal. In the same figure, the Pareto front of the LF samples is shown. It could seem that the LF Pareto optimal solutions dominate the HF front. However, as Figures 6.2a and 6.2b also suggest, the drag and lift force are actually under-predicted by using XFOIL. Indeed, by re-evaluating the Pareto optimal LF optimal designs with the HF solver, the green circles (HF evaluations of LF Pareto front) are found. Thus, we can conclude that the gap between the LF and HF front is due to the approximation error of the LF evaluations. Therefore, the introduction of the HF samples into the surrogate model construction is beneficial for obtaining an accurate Pareto front.

The list of design variables and objectives of the non-dominated designs are shown in Table 6.7. Many design variables reached the

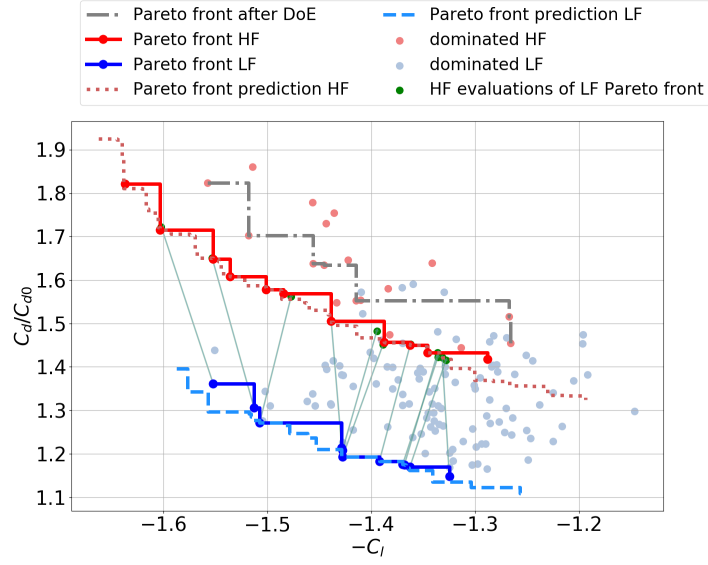


FIGURE 6.10: Comparison of high- and low-fidelity Pareto fronts of the multi-objective airfoil design optimisation problem.

boundary of the design variable limits. For instance, the design variable of the thickness mode is -1 for every non-dominated design. This is expected as thin airfoils produce significantly less drag.

Depending on some further criteria on the propeller blade, the decision-maker can choose the preferred airfoil (among the HF Pareto optimal solutions) for further analysis. For example, Table 6.8, lists three possible designs corresponding to the minimum drag, highest efficiency and maximum lift design. To compare the predicted probability distributions of their aerodynamic forces, these designs are compared with the baseline design in Figure 6.11. Violin plots are drawn based on the 5000 virtual samples mentioned in Section 6.2. By optimising the 95th superpercentile of the aerodynamic forces, the tails of the aerodynamic performance distributions are optimised. Due to the strong correlation of the lift and drag forces, the optimisation of the opposite distribution

6.4. Multi-objective airfoil optimisation under uncertainty with XFOIL and SU2

Table 6.7: List of design variables and objectives of the non-dominated airfoil designs.


x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	$\mathcal{S}_{95}(-C_l)$	$\mathcal{S}_{95}(C_d/C_{d0})$
-1.000	1.000	1.000	1.000	1.000	-1.000	0.000	-1.000	-1.637	1.822
-1.000	1.000	-0.084	1.000	-0.119	-1.000	-0.526	-1.000	-1.603	1.715
-1.000	1.000	-1.000	1.000	-1.000	-0.113	-1.000	-1.000	-1.552	1.649
-1.000	0.663	-0.927	0.654	-1.000	-1.000	-0.595	-1.000	-1.536	1.608
-1.000	0.611	-0.458	-1.000	-0.659	-0.608	-0.345	-1.000	-1.501	1.578
-1.000	0.917	0.272	-1.000	-1.000	-1.000	-1.000	-0.199	-1.485	1.569
-1.000	-0.128	-1.000	0.350	-1.000	-1.000	-0.616	-0.921	-1.439	1.505
-1.000	-0.215	1.000	-1.000	-1.000	-1.000	-1.000	-1.000	-1.388	1.457
-1.000	-0.717	-1.000	-1.000	-1.000	-1.000	0.020	-1.000	-1.363	1.450
-1.000	-1.000	-0.379	1.000	-1.000	-1.000	-1.000	-1.000	-1.346	1.433
-1.000	-1.000	1.000	-1.000	0.000	1.000	-1.000	-1.000	-1.288	1.418

Table 6.8: Comparison of Pareto optimal designs based on various selection criteria.


	C_l	C_d	η
Baseline design	1.4222	0.01646	86.4
Minimum drag design	1.2879	0.01418	90.8
Most efficient design	1.5360	0.01608	95.5
Maximum lift design	1.6371	0.01822	89.9

tails can result in more robust solutions (narrower spread of the distribution) as we can observe in the case of maximum lift and most efficient designs.


The optimisation problem defined with Eqs. (6.1a)-(6.1b) aims to optimise exclusively the drag and lift forces, other aerodynamic and structural parameters of the airfoil are neglected. This can result in airfoil designs which are sub-optimal when structural requirements and other aerodynamic parameters are considered. The shape of the airfoils selected in Table 6.8 is compared with the baseline shape in Figure 6.12. For example, the most efficient design in Figure 6.12b has a very thin trailing edge which is undesired from a structural point of view. Nevertheless, we can see that the pressure distribution in Figures 6.13a, 6.13b and 6.13c is well approximated with XFOIL. This explains the effectiveness



(a) Comparison of the baseline and minimum drag design with uncertainty prediction




(b) Comparison of the baseline and most efficient design with uncertainty prediction




(c) Comparison of the baseline and maximum lift design with uncertainty prediction

FIGURE 6.11: Comparison of the baseline and Pareto optimal designs with uncertainty prediction


of the used multi-fidelity approach.



(a) Minimum drag design



(b) Most efficient design




(c) Maximum lift design

FIGURE 6.12: Comparison of Pareto optimal design geometries based on various selection criteria.


6.4.3 Conclusion of multi-objective airfoil design optimisation under uncertainty

The optimisation of an aerodynamic shape is computationally expensive - the more so, when uncertainties are taken into account. In this study, a multi-objective aerodynamic shape optimisation under uncertainty was presented. The problem was successfully tackled by the proposed multi-fidelity surrogate-assisted optimisation framework. Combining information from XFOIL and the RANS solver of SU2 the computational load was significantly


reduced. Due to the multi-fidelity approach, the optimisation became feasible even with a limited computational budget and an accurate Pareto front approximation was found with only a limited number of high-fidelity RANS simulations. The optimal solutions found by the proposed approach display significant dominance over the baseline solution in the objective space of the reliability measures of the lift and drag.



(a) C_p distribution of minimum drag design



(b) C_p distribution of most efficient design



(c) C_p distribution of maximum lift design

FIGURE 6.13: Comparison of the pressure distribution of Pareto optimal designs based on various selection criteria.

6.5 Constrained single-objective optimisation of an airfoil under uncertainty

This section provides a case study of how to tackle expensive aerodynamic shape optimisation problems with constraints.² In this study, the drag coefficient of the MH114 airfoil is minimised considering constraints imposed on the lift and some geometrical features of the airfoil. The environmental conditions are considered uncertain. Therefore, in this study, the angle-of-attack of the airfoil is considered to have a symmetric beta distribution between 1.75 and 2.25 and shape parameters are equal to 2. The constrained optimisation problem under environmental uncertainty can be formulated as follows:

$$\begin{aligned}
& \min_x \mathcal{S}_{95}^{right}(C_d(x, u)) \\
& \text{s.t. } \mathcal{S}_{95}^{left}(C_l(x, u)) \geq 1.0 \\
& \quad t_{\%}(x) = 13.05 \\
& \quad \text{LER}(x) \geq 0.011377 \\
& \quad \text{TEA}(x) \geq 6.0^\circ
\end{aligned} \tag{6.3}$$

where $t_{\%}$ is the relative thickness, LER is the leading edge radius and TEA is the trailing edge angle of the airfoil.

To handle the constraints with a single-objective optimisation strategy, the constrained problem is scalarised with the penalisation approach. The single-objective penalised formulation can be expressed as:

$$\begin{aligned}
& \min_x \mathcal{S}_{95}^{right}(C_d(x, u)) + p_{C_l} \max(0, 1 - \mathcal{S}_{95}^{left}(C_l(x, u))) + \\
& + p_{\text{LER}} \max(0, 0.011377 - \text{LER}(x)) + p_{\text{TEA}} \max(0, 6 - \text{TEA}(x)),
\end{aligned} \tag{6.4}$$

where the penalisation factors $p_{C_l} = 1000$, $p_{\text{LER}} = 100000$ and $p_{\text{TEA}} = 100$ are chosen by taking the magnitude of the corresponding values into account. The equality constraint of the thickness

²This section discusses a collaboration work published in [Paper C: Morales and Korondi et al., 2020]. The results, tables and figures are identical to the published work. The discussion and interpretation of the results are modified according to the individual thoughts and work of the candidate.

is imposed by scaling the modified airfoil shape to the prescribed value. Only the aerodynamic force requirements (C_d and C_l) are affected by the uncertainty of the angle-of-attack. The geometrical properties of the airfoil are independent from the environmental uncertainty. Therefore, they are evaluated deterministically for each design configuration.

6.5.1 Constrained single-objective optimisation workflow for airfoil design under uncertainty

To reduce the computational burden of the optimisation under uncertainty, the proposed single-objective multi-fidelity surrogate-assisted optimisation workflow is employed. The proposed workflow was detailed in Section 4.4. In this study, the low-fidelity aerodynamic force predictions are calculated with the RANS solver of SU2 using a coarse mesh. Additionally, the drag coefficient predictions are calculated with the so-called far-field method. This drag prediction technique compensates for the truncation error of the numerical scheme and artificial dissipation due to the coarse mesh as was detailed in Section 6.3.3.

The computational chain for combining the high- and low-fidelity information is displayed in Figure 6.14. The high computational cost of RANS simulations using a fine grid motivates to populate the design space with cheaper low-fidelity information. Only a limited number of high-fidelity RANS simulations are allowed during the optimisation to respect the size of the computational budget. Despite the low number of high-fidelity samples, an accurate surrogate model of the performance landscape can be trained with the MFGPR technique. For both the lift and drag coefficient independent surrogates are trained.

The acquisition function presented in Section 4.3 is tailored for the penalised objective and complemented with constraint treatment as in Eq. (4.7). The expected improvement of the penalised objective is calculated by considering the prediction error of the drag prediction. The constraint is treated by multiplying the expected improvement of the objective with the probability of feasibility of the lift coefficient. The problem specific *Constrained*

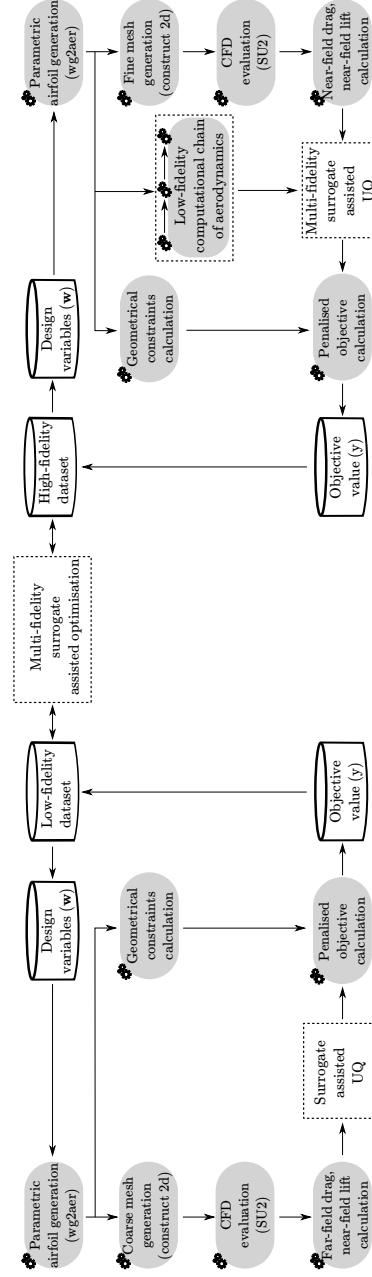


FIGURE 6.14: Computational chain of the aerodynamic forces with the probabilistic model.

Expected Improvement (cEI) is formulated as:

$$\begin{aligned} \text{cEI} &= \mathbb{E} \left[\max \left(0, F_{obj}^* - F_{obj}(\tilde{C}_d, \tilde{C}_l) \right) \right] \mathbb{P} \left[\tilde{C}_l \geq 1 \right] \\ &= \left(\left(F_{obj}^* - F_{obj}(\hat{C}_d, \hat{C}_l) \right) \Phi \left(\frac{F_{obj}^* - F_{obj}(\hat{C}_d, \hat{C}_l)}{\hat{\sigma}_{C_d}^2} \right) + \right. \\ &\quad \left. + \hat{\sigma}_{C_d}^2 \phi \left(\frac{F_{obj}^* - F_{obj}(\hat{C}_d, \hat{C}_l)}{\hat{\sigma}_{C_d}^2} \right) \right) \Phi \left(\frac{\hat{C}_l - 1}{\hat{\sigma}_{C_l}^2} \right), \end{aligned} \quad (6.5)$$

where $\hat{\sigma}_{C_d}^2$ and $\hat{\sigma}_{C_l}^2$ are the standard deviation of the drag and lift coefficient respectively. F_{obj} is the penalised objective given by Eq. (6.4). The best evaluated objective value is given by F_{obj}^* . The Φ and ϕ symbols denote the cumulative distribution function and probability density function of the standard normal distribution respectively. In every iteration, the multi-fidelity surrogate is updated with one additional sample. This is the design with a predicted performance that maximises the cEI function given in Eq. (6.5).

$$x_{new} = \underset{x}{\operatorname{argmax}} \text{cEI}(x) \quad (6.6)$$

The new design is evaluated with the high- or low-fidelity level on the basis of the *Scaled Expected Variance Reduction* (SEVR) measure Eq. (4.12) of the drag coefficient. For this study, we set the costs of the low- and high-fidelity as $c_{LF} = 1$ and $c_{HF} = 10$ respectively.

The calculation of the risk measures requires an additional step to train the local probabilistic models of the lift and drag coefficients. This surrogate-based uncertainty quantification approach overcomes the problem of querying computationally heavy aerodynamic solver by querying the cheap to evaluate local probabilistic model instead. The probabilistic model of the LF is a GPR trained with five LF samples. For predicting the high-fidelity risk measures, a MFGPR model is trained with five LF samples and three HF samples. These models are cheap to evaluate and a sufficiently large number of samples can be drawn to predict the risk measure. To keep the computational demand low, the minimum number of HF samples is used for training the local probabilistic

model. The number of LF samples is arbitrarily increased by 20% to ensure that the LF trend is properly captured. In Figure 6.15, the MFGPR local probabilistic models of the baseline lift and drag are displayed.






FIGURE 6.15: MFGPR models of the aerodynamic force coefficients of the baseline configuration (The LF and HF probabilistic models are drawn with light blue and grey lines respectively. The LF and HF samples are displayed with black and red crosses. The virtual samples are represented by blue dots.).

The convergence analyses of the virtual sample size are shown in Figure 6.16. Both the lift and drag coefficient shows negligible variation above 10^5 virtual samples. Therefore 10^5 virtual samples are generated for calculating the superpercentile risk measures.



(a) Lift coefficient



(b) Drag coefficient

FIGURE 6.16: Convergence of the risk measure value of the aerodynamic force coefficients. (The figures show the averaged value of the risk measure, where 10 independently generated virtual sample set were used at each magnitude level of the virtual sample size. The shaded area represents the approximated 95% confidence interval.)

The costs of the low- and high-fidelity evaluations are determined according to their true computational time. With the used computational infrastructure, a high-fidelity simulation takes approximately 16 times more time than a low-fidelity simulation. This reflects in an approximated cost of 5 and 53 for calculating

the superpercentile risk measure³. Consequently, the 1 to 10 cost-ratio was applied for this design optimisation study.

6.5.2 Results

To highlight the advantages of the proposed multi-fidelity surrogate-assisted optimisation, the constrained airfoil optimisation problem is first solved deterministically. For the deterministic study, the angle-of-attack has no perturbation. The local probabilistic model construction step and risk measure calculation are skipped.

In Table 6.9, the comparison of the obtained optimal designs is provided by solving the deterministic problem with multi-fidelity, single-fidelity surrogate-based and population-based techniques.

The single- and multi-fidelity surrogate-based optimisation were conducted with a computational budget of 300 units. This computational budget equals the total computational cost of 30 high-fidelity evaluations. To train the initial surrogates, equally 10 high-fidelity samples are generated for both the single- and multi-fidelity approach. Additionally, the multi-fidelity approach uses 100 LF samples for constructing the initial surrogate model. With these settings, the single-fidelity approach has a budget for acquiring 20 new high-fidelity samples. On the other hand, the initial surrogate of the multi-fidelity approach requires to burn 200 units of the total computational cost due to the additional 100 LF samples. This results in an optimisation budget which is sufficient only for 10 HF observations in the case of the multi-fidelity approach. Despite the smaller optimisation budget, the multi-fidelity approach finds a better design as the MFGPR can better explore the design landscape due to the high number of LF observations. The single-fidelity GPR model with 10 HF samples provides performance predictions with high model uncertainty. This forces the acquisition function to spend computational budget on exploring designs with sub-optimal aerodynamic performances.

³The LF probabilistic model is trained with 5 LF samples which cost 5 computational units. The HF probabilistic model is trained with 5 LF and 3 HF samples. The computational time required for a single HF evaluation is approximately 16 times slower. Therefore, to evaluate the required number of samples for the HF probabilistic model (i.e. local MFGPR), $5 + 3 \cdot 16 = 53$ computational units must be spent.

Table 6.9: Comparison of multi-fidelity, single-fidelity surrogate-based and population-based optimisation results of the constrained single-objective airfoil design problem. (The cost of a single high- and low-fidelity simulations are 10 and 1 respectively.) [Paper C: Morales and Korondi et al.,2020]

	C_d [dc]	C_l	TEA[°]	LER	Samples	Cost
baseline (MH114)	117.8	1.09	6.60	0.014324	[0, 1]	10
best design MFGPR	112.0	1.00	6.03	0.013228	[120, 18]	300
best design GPR	117.6	1.08	7.76	0.011973	[0, 30]	300
best design CMA-ES ($h = 10.7$)	117.1	1.02	6.17	0.018025	[1800, 0]	1800
best design CMA-ES ($h = 8$)	115.6	1.02	6.34	0.020900	[1800, 0]	2700

The limited number of allowed high-fidelity evaluations prohibits the single-fidelity approach to converge to an optimal design.

Table 6.10: Comparison of the prediction error of MFGPR and GPR surrogate models of the constrained single objective airfoil design problem. (Prediction error is defined as the arithmetic mean value of the relative error of the high-fidelity predictions during the course of optimisation.)

	\hat{C}_d	\hat{C}_l	\hat{F}_{obj}	$F_{obj}(\hat{C}_d, \hat{C}_l, \text{TEA, LER})$	HF iters.
MFGPR	2.04 %	0.71 %	34.51 %	5.43 %	8
GPR	3.11 %	5.65 %	17.14 %	11.53 %	20

A quantitative comparison of the single- and multi-fidelity surrogate models is presented in Table 6.10. The mean prediction errors of the surrogate models are calculated using the HF observations evaluated during the optimisation. MFGPR predicts the lift and the drag coefficients significantly better. The poor performance of the single-fidelity surrogate-based optimisation approach can be explained also by the fact that the lift coefficient is predicted by the GPR with a relatively high error. This might result in evaluating unfeasible designs with eventually high objective value. The results also reveal that training a surrogate directly using the objective values results in a weak surrogate model. The prediction error of the objective is smaller when the objective is calculated using the independently predicted aerodynamic forces.

When the uncertainty of the angle-of-attack is taken into account, the deterministic optimum might show significant constraint violations as displayed in Figure 6.17a. The superpercentile S_{95}^{left} of the lift distribution of the deterministic optimum is 0.976 and illustrated with the blue vertical line. This violates the imposed constraint level drawn with the red vertical line at value 1.00. This is why the design optimisation must be performed by taking the uncertainty into account. To ensure the reliable performance of the airfoil under various angle-of-attack values, the optimisation problem in Eq. (6.3) is solved with a penalisation approach. To solve the problem under uncertainty the MFGPR is employed also to accurately model the local probabilistic space. This enables a

fast calculation of the risk measure which allows us to perform the optimisation under uncertainty. The advantage of performing the design optimisation under uncertainty becomes observable in Figure 6.17a. The probabilistic optimum respects the imposed constraints on the lift coefficient even in the case of uncertain environmental conditions. To gain a more reliable airfoil, the drag coefficient must be increased due to the high correlation between the lift and drag coefficients. This can be seen in Figure 6.17b.

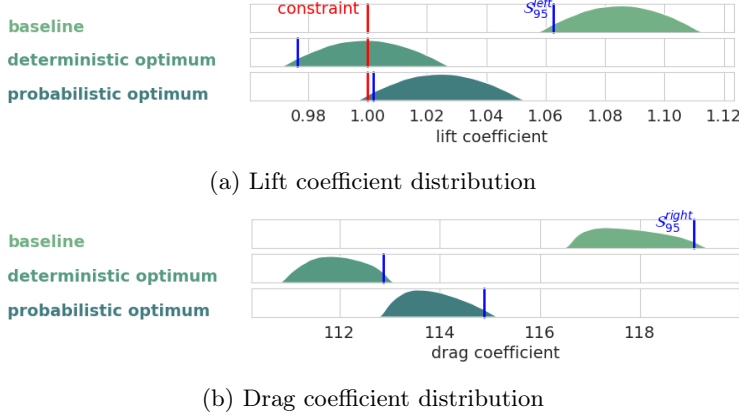



FIGURE 6.17: Prediction of the distributions for the baseline and optimal designs.

The design optimisation under uncertainty is performed with an increased computational budget. To evaluate the superpercentile risk measure of the lift and drag, the employed probabilistic model requires the evaluation of 3 HF samples. Therefore, the total computational budget is tripled. Nevertheless, the size of the budget is sufficient only for 90 HF evaluations, which is still considered a limited budget. For the probabilistic optimisation, the DoE size is set to 50 risk measure evaluations at the low-fidelity level and 5 risk measure evaluations at the high-fidelity level. The obtained probabilistic optimum is compared against the deterministic optimum and the baseline in Table 6.11. To reach the probabilistic optimum, 17 HF and 57 LF risk measure evaluations are performed which means a total of 51 HF and 370 LF solver evaluations. The geometries of the baseline, deterministic and


probabilistic optimum are displayed in Figure 6.18. Both the deterministic and probabilistic optimum have a decreased camber line curvature. The imposed lift constraint stands well below the lift generated by the MH114 airfoil. This allows the optimisation strategy to reduce the drag by reducing the lift coefficient to the target value. The higher reliability of the probabilistic optimum is achieved by finding an airfoil design with a stronger S-shaped lower side. This increases the lift and drag coefficients with the right amount such that the constraint is satisfied. The investigation of the pressure coefficient distributions in Figure 6.19 reveals that the probabilistic optimum has a higher expansion rate, the maximum is reached at approximately 10% of the chord. The deterministic optimum reaches the maximum with a smoother expansion rate at approximately 30% of the chord.

Table 6.11: Comparison of the requirements considering environmental uncertainty.


	baseline (MH114)	deterministic optimum	probabilistic optimum
Objective	119.07	136.55	114.89
$\mathcal{S}_{95}^{right}(\tilde{C}_d)$ [dc]	119.07	112.87	114.89
$\mathcal{S}_{95}^{left}(\tilde{C}_l)$	1.063	0.976	1.002
TEA	6.60	6.03	6.38
LER	0.014324	0.013228	0.016937
Samples [LF,HF]	[5, 3]	[5, 3]	[370, 51]
Cost	35	35	880



(a) Baseline design




(b) Best design of deterministic problem




(c) Best design of probabilistic problem


FIGURE 6.18: Comparison of optimal airfoil geometries of the constrained single-objective problem



(a) Pressure distribution of baseline



(b) Pressure distribution of the best design of deterministic problem



(c) Pressure distribution of the best design of probabilistic problem

6.5.3 Conclusion of the constrained single-objective airfoil design optimisation

This section discussed the constrained single-objective optimisation of the MH114 airfoil. Due to the computationally demanding aerodynamic force analysis of an airfoil design with RANS simulations, the optimisation under uncertainty was performed by using low-fidelity aerodynamic force predictions of RANS simulations with a coarser mesh. The drag predictions of the low-fidelity RANS simulations were calculated with the so-called far-field formula. This drag prediction technique compensates the spurious drag stemming from the employed numerical scheme and numerical dissipation; hence, it provides an accurate drag prediction despite the coarse spatial discretisation. By using low- and high-fidelity information, the optimisation problem under uncertainty was solved efficiently with the proposed optimisation strategy.

“The purpose of the experiment is not to verify a proposed theory but to replace a computation from an unquestioned theory by direct measurements. Thus, wind tunnels are, for example, used at present, at least in part, as computing devices of the so-called analogy type . . . to integrate the nonlinear partial differential equations of fluid dynamics.”

— John von Neumann

7.1 Introduction to propeller blade optimisation

A propeller (or rotor) is an aircraft component which transforms the engine angular kinetic energy into forward motion. The angular rotation of propeller blades generates the so-called thrust. The thrust force moves the aircraft forward in the air. To rotate the propeller blade, the engine has to work against the drag force. Consequently, the efficiency of a propeller can be evaluated with the ratio of the generated thrust and the required power. The efficiency of the propeller (i.e. propulsive efficiency) is governed by the surface geometry of the propeller blades. Therefore, the shape optimisation of a propeller aims to find blade geometries which maximise the propulsive efficiency. To increase the efficiency of the propeller, the propeller can be placed inside a duct (or shroud) which typically increases the mass flow. The theo-

retical calculations credit the increased mass flow to a reduced slipstream contraction [Weir, 1987; Black and Rohrbach, 1968]. However, for higher Mach numbers, the slipstream contraction decreases anyway, and the drag induced by the duct increases. This mitigates the advantages of ducted propellers for high-speed aircraft [El-Sayed, 2008]. Remaining in the low-speed regime allows benefiting the most from the increased efficiency of a ducted propeller propulsion unit. Therefore, ducted propellers can be applied to small scale aircraft which operate at lower speeds. The increased propulsion efficiency makes ducted propellers promising candidates for electrical aircraft where the ratio of thrust and electricity consumption must be highly optimised. Due to uncertain environmental conditions, the aerodynamic shape optimisation of the propeller must be performed considering uncertainties. The following section explains how the uncertainty is treated in our propeller optimisation study.

7.2 Uncertainty quantification with local probabilistic model


The design process of a ducted propeller aims to estimate the performance of the propulsion system in various conditions. During the operation, the loading of the blades can vary depending on the environmental conditions. For the rotor blade optimisation studies, the free-stream pressure, temperature and velocity are assumed to be probabilistic. The parameters of the probabilistic distribution is given in Table 7.1.

Similarly to the analytical function and airfoil shape optimisation studies, the probabilistic space is modelled locally at each


Table 7.1: The uncertain environmental parameters of the propeller blade optimisation problems

Environmental parameter	Probability distribution	Lower bound	Upper bound	Shape parameters
pressure [Pa]	Beta	96325	106325	$\alpha = \beta = 2$
temperature [K]	Beta	283.15	293.15	$\alpha = \beta = 2$
velocity [m/s]	Beta	31.53	36.53	$\alpha = \beta = 2$

blade design configuration. The expensive 3D *Reynolds-averaged Navier-Stokes* (RANS) simulations make it challenging to build an accurate probabilistic model using only *High-Fidelity* (HF) evaluations. Therefore, at high-fidelity level, the local probabilistic space is modelled by a *Multi-Fidelity Gaussian Process Regression* (MFGPR) using 30 *Low-Fidelity* (LF) and 3 HF samples. The number of HF samples is chosen to be as low as possible. When constant regression functions are chosen for both the mean trend and the low-fidelity coupling term, the MFGPR model operates with 2 regression coefficients (β and ρ) at the top-level which requires a minimum of 3 samples to train the model. The number of low-fidelity evaluations is determined based on the thumb rule of *Gaussian Process Regression* (GPR) methods. Typically, GPR methods scale linearly with the number of dimensions (n_d). Therefore, the suggested number of samples is 30 (considering that the number of samples should be at least ten times the dimension of the problem by the rule of thumb). The convergence analysis of the number of low-fidelity samples is displayed in Figure 7.1 which suggests that even a lower number of low-fidelity samples would provide an accurate local probabilistic model. The local probabilistic models can be utilised to generate a larger set of virtual samples which can be used to predict the chosen risk measure accurately. The convergence studies in Figure 7.2 show that the 95th superpercentile of the thrust and power coefficients can be sufficiently accurately calculated with 10^5 samples using the trained probabilistic models. Surrogate-assisted uncertainty quantification highly depends on the quality of the surrogate. Due to the expensive computational cost of the high-fidelity level, the quality cannot be assessed directly with test data. However, the presented convergence test provides a certain level of assurance that our probabilistic models are well predicting the desired risk measure.




(a) Convergence of the risk measure of the thrust in function of low-fidelity samples




(b) Convergence of the risk measure of the power in function of low-fidelity samples.

FIGURE 7.1: Convergence analysis of the number of low-fidelity samples of the local probabilistic model. (The figures show the averaged value of the risk measure, where 10 independently generated LF sample sets were used at each level of the low-fidelity sample size. The shaded area represents the approximated 95% confidence interval.)



(a) Convergence of the risk measure of the thrust in function of virtual samples



(b) Convergence of the risk measure of the power in function of virtual samples

FIGURE 7.2: Convergence analysis of the number of virtual samples of the local probabilistic model (The figures show the averaged value of the risk measure, where 10 independently generated virtual sample sets were used at each magnitude level of the virtual sample size. The shaded area represents the approximated 95% confidence interval.)

7.3 Aerodynamic performance prediction of a propeller

The prediction of the aerodynamic performance of a propeller is a challenging task. The theoretical foundation of the propeller theory was laid down in [Betz, 1919]. The idealised models of the propeller and the surrounding flow let the research community construct a powerful propeller analysis tool called *Blade Element Momentum Theory* (BEMT). This propeller performance predictor can accurately analyse the generated aerodynamic state of the propeller at an economic cost. The details of the technique are discussed in Section 7.3.2. The aerodynamic performance of a propeller can be calculated also by solving the RANS partial differential equations. This equation has to be solved numerically which is typically an expensive calculation due to the high number of discretisation cells required. The necessary steps to conduct the aerodynamic analysis of a propeller with the RANS solver of SU2 is detailed in Section 7.3.3.

7.3.1 Parametric geometry definition of propeller blades

In this work, the geometry of the propeller blade is defined by the chord and twist distributions of the propeller blade. The radial blade section distribution is considered to be constant and the MH114 airfoil is used uniformly. The radial blade sections are sized and twisted around the quarter-chord point of the local airfoil section. The design variables at the root and tip are illustrated in Figure 7.3. To provide parametric chord and twist distributions for the design optimisation strategy, the superposition of design shape modes is calculated. Five shape modes for the chord and five shape modes for the twist distributions are considered. The shape modes are selected in a way that they can be connected to physical values. Namely, the five modes are related to the chord and twist value at five selected locations of the blade span. The chosen modes are listed in Table 7.2 and the design variables are given in Table 7.3. The first mode determines the value at the tip and the second mode adds on the top of the constant mode to determine the value at the root. Therefore, the factor of the second mode is the deviation of the root and tip values. The

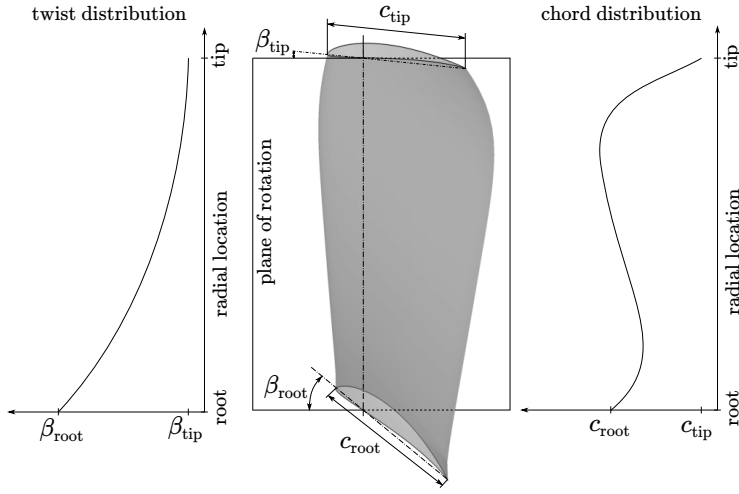


FIGURE 7.3: Design variables: twist and chord distributions

 Table 7.2: Shape modes of chord and twist distribution functions (r_R is the relative radial location)

Mode	Function type	Physical interpretation
mode 1	Constant	value at tip
mode 2	Linear	deviation of the root and tip value
mode 3	Quadratic	value at the middle ($r_R = 0.5$)
mode 4	Hicks-Henne	value at $r_R = 0.8$ relative coordinate
mode 5	Hicks-Henne	value at $r_R = 0.3$ relative coordinate

other three modes can be hierarchically superposed on the two first modes without changing the values at the root and tip. The selected shape modes are visualised in Figure 7.4.

Table 7.3: Design variables of the twist and chord distributions (r_R is the relative radial location)

Design var.	Corresponding function	Physical interpretation
x_1	$x_1 \cdot (\text{mode 1})$	chord length at tip
x_2	$(x_2 - x_1) \cdot (\text{mode 2})$	chord length at root
x_3	$x_3 \cdot (\text{mode 3})$	factor of quadratic chord mode
x_4	$x_4 \cdot (\text{mode 4})$	local chord mode at $r_R = 0.8$
x_5	$x_5 \cdot (\text{mode 5})$	local chord mode at $r_R = 0.3$
x_6	$x_6 \cdot (\text{mode 1})$	twist angle at tip
x_7	$(x_7 - x_6) \cdot (\text{mode 2})$	twist angle at root
x_8	$x_8 \cdot (\text{mode 3})$	factor of quadratic twist mode
x_9	$x_9 \cdot (\text{mode 4})$	local twist mode at $r_R = 0.8$
x_{10}	$x_{10} \cdot (\text{mode 5})$	local twist mode at $r_R = 0.3$




FIGURE 7.4: Shape modes of chord and twist distribution functions (mode 1: blue line, mode 2: orange line, mode 3: green line, mode 4: red line, mode 5: purple line)

7.3.2 Low-fidelity predictor: Blade element momentum theory




FIGURE 7.5: Blade Element velocities and forces

BEMT combines the *Blade Element Theory* (BET) and *Actuator Disk Theory* (ADT) into an iterative solver [Glauert, 1983; Veldhuis, 2005; Hansen, 2015]. In both BET and ADT, the propeller blade is discretised with a given number of annuli. The streamtubes are considered radially independent and the interaction between the adjacent streamtubes are neglected [Whale et al., 2000]. The effect of the actual blade elements is averaged over time. Therefore, the technique is applicable to any number of blades. Each annulus is characterised by their local velocities and forces. At each radial station the velocity state is given by Eq. (7.1):

$$V_x = V_\infty(1 + a), \quad (7.1)$$

$$V_\theta = \omega r(1 - b), \quad (7.2)$$

$$V = \sqrt{V_x^2 + V_\theta^2}, \quad (7.3)$$

where V_∞ is the free-stream velocity, V_θ is the angular velocity and V is the local velocity seen by the blade. r is the radius of the annulus and ω is the angular velocity of the propeller. a and b denote the induced axial and angular inflow factor respectively. The velocity vectors and resulting forces are depicted in Figure 7.5.

By knowing the induced velocities a and b , BET can determine the thrust T and power P of each blade element with Eqs. (7.4), (7.5):

$$\delta T = \frac{1}{2} \rho V^2 c (C_l \cos(\varphi) - C_d \sin(\varphi)) B dr, \quad (7.4)$$

$$\delta P = \frac{1}{2} \rho V^2 c (C_d \cos(\varphi) + C_l \sin(\varphi)) r \omega B dr, \quad (7.5)$$

where ρ is the fluid density, c is the chord length and B is the number of blades. C_l and C_d are the 2D lift and drag coefficients of the blade element section. The lift $C_l(\alpha)$ and drag $C_d(\alpha)$ are functions of the angle-of-attack α . Following the angle orientations in Figure 7.5, the angle-of-attack can be calculated by the following equations:

$$\varphi = \tan^{-1} \frac{V_x}{V_\theta}, \quad (7.6)$$

$$\alpha = \theta - \varphi, \quad (7.7)$$

where φ is the relative flow angle seen by the blade and θ is the geometrical twist of the blade element.

The induced velocities, however, are not known and their direct calculation would be a tedious work. Therefore the thrust and power are alternatively calculated according to the ADT :

$$\delta T = \rho 4\pi r V_\infty^2 a (1 + a) dr, \quad (7.8)$$

$$\delta P = \rho 4\pi r^3 V_\infty b (1 + a) \omega^2 dr, \quad (7.9)$$

The Eqs. (7.4), (7.5) and (7.8), (7.9) are equated respectively in BEMT and the a and b induced velocity factors are calculated by iteratively minimising the deviation between the two theories. By considering that $V = \frac{V_x}{\sin \varphi} = \frac{V_\infty(1+a)}{\sin \varphi}$ and the blade solidity is $\sigma_r = \frac{Bc}{2\pi r}$, the problem to be solved iteratively can be reduced to Eqs. (7.10), (7.11):

$$\frac{a}{1+a} = \frac{\sigma_r}{4 \sin^2(\varphi)} (C_l \cos(\varphi) - C_d \sin(\varphi)), \quad (7.10)$$

$$\frac{b}{1-b} = \frac{\sigma_r}{4 \sin(\varphi) \cos(\varphi)} (C_d \cos(\varphi) + C_l \sin(\varphi)). \quad (7.11)$$

The idealised models of BEMT allows to calculate a fast prediction of the propeller performance. In this work, the BEMT

predictions are done with the XROTOR [Drela and Youngren, 2013] software which has an advanced implementation of the theory.

The computational chain for calculating the propeller aerodynamic performance with XROTOR is illustrated in Figure 7.6. XROTOR requires the distribution of twist and chord length as input. Furthermore, the cross-sectional aerodynamic properties of the blade must be also given in the XROTOR configuration file. The aerodynamic properties of the local cross-section geometries are calculated with XFOIL if the aerodynamic parameters are not available.

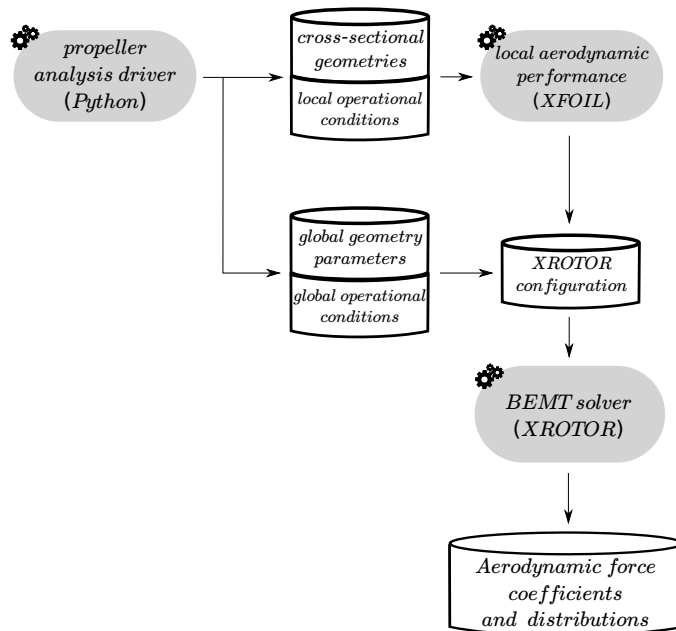


FIGURE 7.6: Computational chain for calculating the propeller aerodynamic performance with XROTOR.

7.3.3 High-fidelity predictor: Reynolds Averaged Navier-Stokes solver

In this study, the 3D compressible RANS equation is solved with SU2. To perform the aerodynamic analysis of the propeller the computational chain illustrated in Figure 7.7 is followed. The

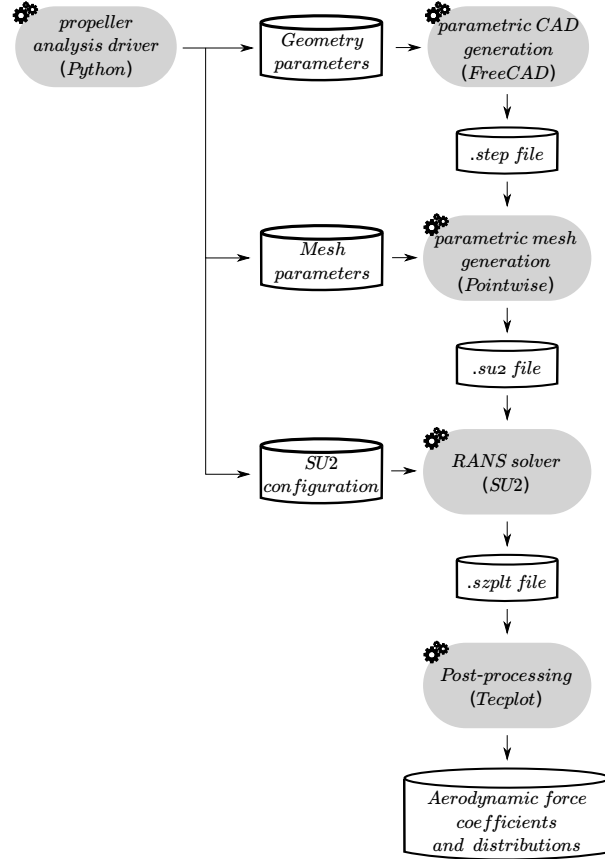


FIGURE 7.7: Computational chain for solving RANS equations numerically.

parametrisation of the propeller is implemented in a *python* script. With the geometrical parameters, *FreeCAD* [FreeCAD, 2018] gen-

erates the *Computer-Aided Design* (CAD) model of the propeller. The CAD model alongside with the surrounding domain is meshed with *Pointwise* [Pointwise, 2017] which exports the grid in *.su2* format. The compressible solver of SU2 solves the RANS equations. The obtained flow field is then post-processed with *Tecplot* [Tecplot, 2019] which calculates the aerodynamic coefficients and distributions.

The thrust and power distributions of the baseline geometry calculated with XROTOR and SU2 are compared in Figure 7.8.




FIGURE 7.8: Comparison of thrust and power distributions of the baseline design calculated with XROTOR and SU2.

7.3.3.1 Ducted propeller geometry

The ducted propeller geometry is illustrated in Figure 7.9. The configuration considers a double ring mounting which avoids the presence of energy losses due to the flow over the tip. The parametrisation of the propeller blade geometry is detailed in Section 7.3.1.




FIGURE 7.9: Propeller blade geometry with double ring mounting. (The rotating parts are highlighted with turquoise colour.)

7.3.3.2 Parametric mesh generation

The mesh generation process is still among the most challenging and time-consuming pre-processing step of *Computational Fluid Dynamics* (CFD). To reduce the geometrical complexity and computational burden of the CFD simulation, only one blade and its surrounding wedge domain are discretised. This is a common practice for propeller simulations as the geometry is axisymmetric and the generated wake has a cyclic nature [Klein, 2017; Gaggero et al., 2017].

In this study, the rotor consists of 5 blades, therefore, the flow conditions are analysed in a 72° wedge as shown in Figure 7.10. The dimensions of the wedge domain are $10R$, $15R$ and $35R$ in radial, upstream and downstream directions respectively, where $R = 0.28829$ is the radial position of the blade tip.




FIGURE 7.10: Outer wedge faces.






FIGURE 7.11: Near propeller faces.

Due to some limitations of the SU2 solver the periodic faces cannot share a common edge. The cells connected to one periodic


face will be used to generate halo cells at the other periodic face. This requires having more than one cell between the periodic faces. Therefore, the wedge is closed by a dummy surface generated at $r_{dummy} = 0.003$ in the upstream region and by a dummy surface extending the hub into the downstream region. The modelled surfaces near the propeller are presented in Figure 7.11.



(a) Structured boundary layer at tip (2D top view)



(b) Structured boundary layer at root (3D view zoom on root)




(c) Structured boundary layer at tip (3D view zoom on tip)


FIGURE 7.12: Structured boundary layer mesh.

The investigated domain is meshed with a hybrid structured


and unstructured grid. The advantage of a structured grid becomes prominent when the flow physics properties are not changing in all dimensions with the same magnitude [Tu et al., 2018]. Then it is enough to generate a dense discretisation only in the direction of the dominant flow physics change while the directions with lesser flow physics changes can be sparsely discretised. Therefore, a fully structured boundary layer grid is used to capture the near-field flow around the blade. To properly capture the flow physics in the boundary layer a large number of small cells are clustered near the blade surface. The clustering is required only in the direction perpendicular to the blade surface due to the nature of the boundary layer profile. The structured boundary layer mesh is shown in Figure 7.12. In other regions of the domain, an unstructured grid has been generated. The main motivation to facilitate the mesh generation of different blade configurations. The side view of the full domain is presented in Figure 7.13a. A refinement box is placed around the ducted propeller in which the maximal cell size is reduced to 0.005 as shown in Figure 7.13b. The mesh is further refined in the regions close to the blade at the LE and TE of the duct which is displayed in Figure 7.13c. The fine-tuning of the mesh parameters is performed with numerous trial and error iterations considering deformation measures of the cells, the dimensionless wall distance number of propeller surfaces and the aerodynamic force sensitivity of the mesh.



(a) Mesh sizing of the outer domain



(b) Mesh sizing of the refinement box



(c) Mesh sizing near the duct and blade

FIGURE 7.13: Mesh sizing

7.3.3.3 Boundary conditions

To each boundary face, a boundary condition is assigned according to Table 7.4. The free-stream fluid properties are imposed at the inlet, outlet and top faces. The two sides of the wedge domain are

Table 7.4: Boundary condition definition of periodic propeller mesh.

Face	SU2 Boundary Condition
inlet	far-field
outlet	far-field
top far-field	far-field
side+	periodic
side-	periodic
dummy upstream	Euler
dummy downstream	Euler
hub static	Euler
duct static	Euler
hub propeller band	Euler
duct propeller band	Euler
blade	zero heat-flux

connected by periodic boundary enforcement. The blade surface is treated as a no-slip wall with zero heat flux. The other faces are set as Euler boundaries.

7.4 Multi-objective optimisation under uncertainty with XROTOR and SU2

In this section, a multi-objective optimisation of the propeller blade design is conducted considering uncertain environmental conditions and a constant angular velocity. The assumed probabilistic environmental conditions are given in Table 7.1. The constant angular velocity is set to $\omega = 476.475$ [rad/s]. Propeller blade designs are sought with maximal thrust and minimal power consumption. To assure the reliable operation of the propeller over a range of uncertain environmental conditions, the 95th superpercentiles of the thrust and power coefficients are optimised. The problem is formulated as follows:

$$\min_x \mathcal{S}_{95}^{right}(-T_c(x, \tilde{u})) \quad (7.12a)$$

$$\min_x \mathcal{S}_{95}^{right}(P_c(x, \tilde{u})) \quad (7.12b)$$

where the negative thrust is minimised instead of maximised; hence, the right-tail superpercentile. The thrust coefficient (T_c) and power coefficients (P_c) are:

$$T_c = \frac{T}{\frac{1}{2}\rho V_\infty^2 \pi R^2}, \quad (7.13)$$

$$P_c = \frac{P}{\frac{1}{2}\rho V_\infty^3 \pi R^2}, \quad (7.14)$$

where ρ is the air density. V_∞ is the free-stream velocity and R is the radius of the propeller.

7.4.1 Computational chain for multi-objective propeller blade design optimisation under uncertainty

The computational chain of the multi-objective propeller optimisation is presented in Figure 7.14.

The acquisition function of the multi-objective propeller optimisation problem is based on Eq. (4.14). The *Hypervolume Improvement of the Lower Confidence Bound* (HVI_{LCB}) of the thrust and power coefficients are predicted for each candidate design. The design which provides the maximum HVI_{LCB} . The HVI_{LCB} measure of the problem is:

$$\begin{aligned} HVI_{LCB} = \\ HV \left(\mathcal{P} \cup \left\{ -\hat{T}_c^{S_{95}^{right}} - 2\hat{\sigma}_{T_c}, \hat{P}_c^{S_{95}^{right}} - 2\hat{\sigma}_{P_c} \right\}_{\text{new}} \right) - HV(\mathcal{P}), \end{aligned} \quad (7.15)$$

where \hat{T}_c and \hat{P}_c denote the predicted risk measures of the thrust and power respectively. $\hat{\sigma}_{T_c}$ and $\hat{\sigma}_{P_c}$ denote the Gaussian process variance of the predicted risk measures.

The strategy of fidelity level selection is performed based on the *Scaled Hypervariance Reduction* (SHVR) of the thrust and power. The costs of the low- and high-fidelity simulations are $c_{LF} = 1$ and $c_{HF} = 200$ respectively. The costs of the fidelity levels are set after conducting some trial and error experiments targeting an approximately equal budget allocation between the fidelity levels.

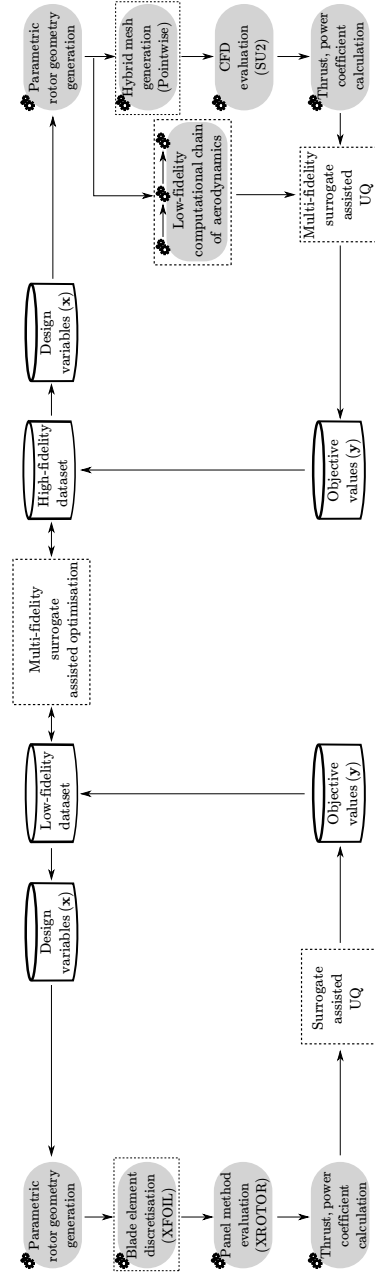


FIGURE 7.14: Computational chain of the propeller aerodynamics with reliability measure calculation for multi-objective optimisation problem

Table 7.5: Computational budget spent on multi-objective propeller design optimisation. (Total computational budget is 21240 units which are equal to the cost of 120 LF and 22 HF reliability samples.)

	LF DoE	LF total	HF DoE	HF total
Solver	XROTOR	XROTOR	SU2	SU2
Evaluation cost	1	1	200	200
Reliability samples	120	294	6	18
Total samples	3780	10440	18	54
Budget spent	3780	10440	3600	10800
Budget percentage	17.8%	49.2%	16.9 %	50.8%

7.4.2 Results of multi-objective propeller optimisation under uncertainty

The budget allocation of the multi-objective optimisation problem is given in Table 7.5. The *Design of Experiments* (DoE) is composed of 120 LF and 6 HF risk measure calculations which takes a total of 3780 LF and 18 HF aerodynamic analyses¹. The total budget is set to 21240 units which are equal to the cost of 120 LF and 22 HF risk measure calculations of the underlying problem. The total cost of DoE is 34.70% of the total budget and the computational resources were allocated approximately equally between the high- and low-fidelity solvers.

The history of the multi-objective optimisation is given in Figure 7.15. After evaluating the DoE, the optimisation strategy performs low-fidelity calculations to identify promising design candidates. Most of the HF evaluations are performed in the fourth quarter of the optimisation run when the design landscape is already explored by the low-fidelity observations.

The low- and high-fidelity Pareto fronts are shown in Figure 7.16. Out of the 18 HF reliability evaluations, 10 of them are on the obtained Pareto front. We can see that the low- and

¹The LF probabilistic model is trained with 30 LF samples resulting in 3600 LF evaluations in total for 120 LF reliability calculations. Additionally, 180 LF samples are evaluated for training 6 local probabilistic models of the HF level (local MFGPR models). A single HF local probabilistic model is trained with 30 LF samples and 3 HF samples.




FIGURE 7.15: History of the low- and high-fidelity thrust and power coefficients, and the computational budget in the multi-objective problem. (The positive thrust coefficients are given here: $T_c^{left} = -(-T_c^{right})$.)




FIGURE 7.16: Pareto front of multi-objective propeller optimisation problem

high-fidelity fronts are meeting in the low-thrust-low-power side of the objective space. The deviation of the fronts is more significant in the high-thrust-high-power side (the Pareto fronts are displayed for simultaneously minimising the negative thrust and power in Figure 7.16).

The thrust and power coefficient distributions of five selected Pareto optimal designs are displayed in Figures 7.17 and 7.18. The designs with high power and thrust coefficients are more sensitive to uncertainties (as indicated by the larger spread of the distribution range). This is due to the fact that higher thrust can be reached by bigger chord length or higher twist values. Such blade properties are increasing the probability of flow separation in uncertain environmental conditions. Therefore, it is not surprising that the Pareto optimal design no. 1 and 4 shows higher response uncertainty than design no. 6 and 8. The Pareto optimal design no. 10 has very low thrust and power; however, with much robust performance.




FIGURE 7.17: Thrust coefficient distributions of selected Pareto optimal designs.




FIGURE 7.18: Power coefficient distributions of selected Pareto optimal designs.

The list of design variables and superpercentile values of the Pareto optimal designs are given in Table 7.6.

By plotting the sectional properties of the Pareto optimal designs, a further explanation can be provided why the LF and HF Pareto fronts have a significant deviation in the high-thrust-high-power region. In Figure 7.19k and 7.19p, the flow separation resulted in the drop of thrust and power coefficients close to the blade tip. The flow separation close to the tip is not predicted by XROTOR for this blade; hence, the XROTOR and SU2 predict significantly different aerodynamic coefficients. The Pareto optimal design no.4 has smaller chord length close to the blade tip, the flow separation is less severe. Therefore, the predicted values of the two aerodynamic solvers are also closer to each other.

7.4.3 Conclusion of multi-objective propeller blade optimisation study

A multi-objective propeller blade design problem under environmental uncertainty was tackled by the proposed optimisation algorithm in this section. The proposed multi-objective multi-fidelity surrogate-assisted optimisation technique made it possible to perform the optimisation under uncertainty despite the high computational cost of the RANS simulations. Multi-fidelity surrogate models were trained with the data from high-fidelity RANS simulations and low-fidelity XROTOR calculations. The multi-fidelity optimisation strategy helped to allocate the computational resources efficiently: the exploration of the propeller designs was performed with XROTOR and only the most promising blade designs were evaluated with SU2.

Table 7.6: Design variables and objective values of the Pareto optimal designs.

no.	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	$-T_c^{S^{right}}_{95}$	$P_c^{S^{right}}_{95}$
1	0.057	0.070	0.010	-0.008	0.010	60.0	38.4	10.0	8.2	-10.0	-2.713	9.237
2	0.046	0.070	0.010	-0.010	0.010	60.0	39.4	-3.2	10.0	4.0	-2.635	8.869
3	0.055	0.061	0.004	0.004	-0.004	45.0	30.0	4.0	4.0	8.0	-2.610	7.035
4	0.070	0.040	0.010	0.003	0.010	60.0	30.1	-10.0	9.6	6.3	-2.358	6.354
5	0.049	0.067	-0.004	-0.004	-0.008	57.0	30.0	-4.0	-4.0	8.0	-1.936	4.574
6	0.070	0.043	0.002	0.004	0.010	37.4	23.2	3.1	-6.1	8.9	-1.509	3.443
7	0.040	0.040	0.010	-0.008	-0.008	57.8	40.7	-5.8	-2.0	-6.5	-1.405	3.337
8	0.060	0.050	0.010	0.010	-0.010	60.0	20.0	-10.0	0.0	0.0	-1.145	2.488
9	0.057	0.044	0.010	0.003	-0.010	30.0	15.0	-10.0	9.0	10.0	-0.814	1.789
10	0.070	0.048	-0.010	-0.010	-0.010	43.4	15.0	-10.0	-5.4	6.0	-0.310	0.753

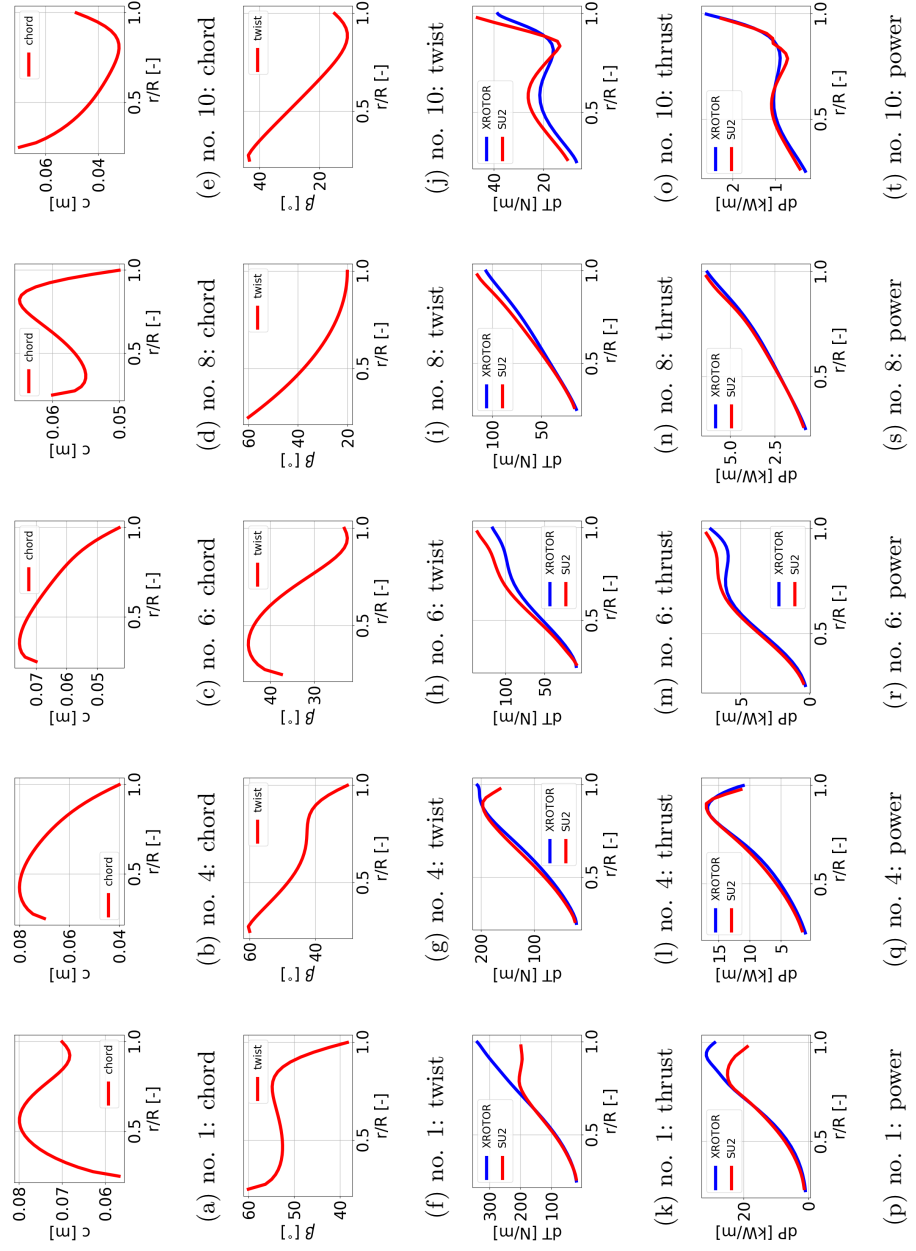


FIGURE 7.19: Comparison of the sectional properties of the Pareto optimal designs.

7.5 Constrained single-objective optimisation under uncertainty with XROTOR and SU2

In this section, a constrained single-objective optimisation of the rotor blade is carried out. The optimisation is performed considering environmental uncertainties and a constant angular velocity as in the multi-objective case study. The right-tail 95th superpercentile of the power coefficient (P_c) is minimised such that the left-tail 95th superpercentile of the thrust coefficient (T_c) is required to be greater than one as given in Eq. (7.16).

$$\begin{aligned} \min_x \quad & \mathcal{S}_{95}^{right}(P_c(x, u)) \\ \text{s.t. :} \quad & \mathcal{S}_{95}^{left}(T_c(x, u)) \geq 1.0 \end{aligned} \quad (7.16)$$

To satisfy the imposed constraint on the thrust, the constrained problem is transformed into the following minimisation problem of a penalised objective:

$$\min_x \quad \mathcal{S}_{95}^{right}(P_c(x, u)) + p_{T_c} \max(0, 1 - \mathcal{S}_{95}^{left}(T_c(x, u))), \quad (7.17)$$

where $p_{T_c} = 100$ is the penalisation factor. The single-objective optimisation problem of the penalised objective in Eq. (7.17) is solved with the single-objective surrogate-based optimisation proposed in Algorithm 1.

7.5.1 Computational chain for constrained single-objective propeller blade design optimisation under uncertainty

The problem-specific computational chain of the constrained single-objective optimisation problem is presented in Figure 7.20. The computational chain resembles the one presented for the multi-objective case; however, an additional step is added to calculate the penalised objective function. It is important to note that the penalised objective is calculated with the reliability measures of the thrust and power coefficients.

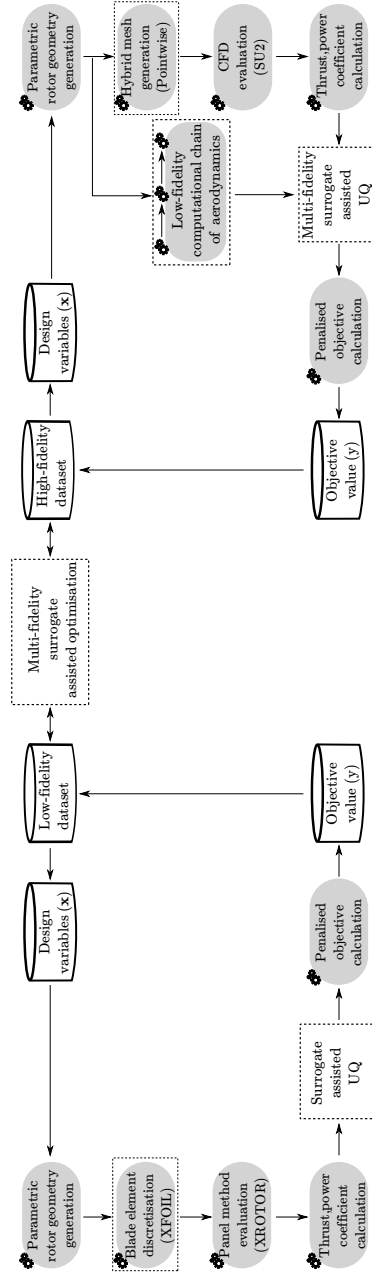


FIGURE 7.20: Computational chain of the propeller aerodynamics with reliability measure calculation for constrained single-objective optimisation problem.

The acquisition function employed in this optimisation study is formulated according to Eq.(4.7) and given by the following equation:

$$\begin{aligned} \text{cEI} &= \mathbb{E} \left[\max \left(0, F_{obj}^* - F_{obj}(\tilde{P}_c^{S_{95}^{right}}, \tilde{T}_c^{S_{95}^{left}}) \right) \right] \mathbb{P} \left[\tilde{T}_c^{S_{95}^{left}} \geq 1 \right], \\ &= \left(\left(F_{obj}^* - F_{obj}(\hat{P}_c, \hat{T}_c) \right) \Phi \left(\frac{F_{obj}^* - F_{obj}(\hat{P}_c, \hat{T}_c)}{\hat{\sigma}_{P_c}^2} \right) + \right. \\ &\quad \left. + \hat{\sigma}_{P_c}^2 \phi \left(\frac{F_{obj}^* - F_{obj}(\hat{P}_c, \hat{T}_c)}{\hat{\sigma}_{P_c}^2} \right) \right) \Phi \left(\frac{\hat{T}_c - 1}{\hat{\sigma}_{T_c}^2} \right), \end{aligned} \quad (7.18)$$

where $F_{obj}(\hat{P}_c, \hat{T}_c)$ is the penalised objective of the superpercentiles of the power and thrust. The \hat{T}_c and \hat{P}_c denote the predicted risk measures of the thrust and power respectively. $\hat{\sigma}_{T_c}$ and $\hat{\sigma}_{P_c}$ denote the Gaussian process variance of the predicted risk measures.

The strategy of fidelity level selection is performed based on the variance reduction of the $P_c^{S_{95}^{right}}$:

$$l = \underset{LF, HF}{\operatorname{argmax}} \text{SEVR}_l, \quad (7.19)$$

where the SEVR is defined as:

$$\text{SEVR}_{LF} = \frac{\rho^2(x_{new}) \hat{\sigma}_{P_c, LF}^2(x_{new})}{c_{LF}}, \quad (7.20)$$

$$\text{SEVR}_{HF} = \frac{\rho^2(x_{new}) \hat{\sigma}_{P_c, LF}^2(x_{new}) + \hat{\sigma}_{P_c, \delta_{HF}}^2(x_{new})}{c_{HF}}, \quad (7.21)$$

where $c_{LF} = 1$ and $c_{HF} = 200$ are the costs of the low- and high-fidelity simulations respectively. The costs of the fidelity levels are kept unchanged compared to the multi-objective study in Section 7.4.

7.5.2 Results of constrained single-objective propeller optimisation under uncertainty

The history of the optimisation process is displayed in Figure 7.21. The initial multi-fidelity DoE samples provided sufficient information to build an accurate initial MFGPR model. This helped the




FIGURE 7.21: History of the low- and high-fidelity objectives and the computational budget in the constrained single-objective problem.

7.5. Constrained single-objective optimisation under uncertainty with
XROTOR and SU2

Table 7.7: Computational budget spent on constrained single-objective propeller design optimisation. (Total computational budget is 26910 units which are equal to the cost of 120 LF and 37 HF samples.)

	LF DoE	LF total	HF DoE	HF total
Solver	XROTOR	XROTOR	SU2	SU2
Evaluation cost	1	1	200	200
Reliability samples	120	408	6	23
Total samples	3780	12930	18	69
Budget spent	3780	12930	3600	13800
Budget percentage	14.0%	48.0%	13.4 %	51.3%


identification of promising regions of the design landscape. Then the optimisation evaluated samples from the promising regions to find the minimum of our penalised objective. In the first part of the optimisation (first 150 iterations), the strategy chose to evaluate most of the samples with the low-fidelity solver. The computationally more expensive high-fidelity simulations were performed mostly in the second half of the optimisation run. We can also notice that the high-fidelity evaluations were performed only for highly promising candidate designs. Designs with high constraint violation or high power consumption were evaluated with the LF solver in the first half of the optimisation.

The total computational budget is set to 26910 units which are equal to the sum of the budget required for DoE and 31 additional HF reliability evaluations. In Table 7.7, the number of evaluations performed with the high- and low-fidelity solver are summarised. In total, 99.3% of the budget is spent allocating approximately equal portions of the budget for high- and low-fidelity evaluations.


The multi-fidelity surrogate-based optimisation approach finds a significantly better design than the available baseline design. A comparison of the performance indicators of the baseline and the optimal designs is provided in Table 7.8. The optimal design slightly violates the constraint imposed on the thrust coefficient. This is due to the penalisation approach. The level of constraint violation acceptance can be controlled by the penalisation factor p_{T_c} in Eq. (7.17).

Table 7.8: Comparison of the baseline and the optimal designs of the constrained single-objective problem


	Objective	$T_c^{S_{95}^{left}}$	$P_c^{S_{95}^{right}}$
baseline	2.49	1.14	2.486
probabilistic opt.	2.04	0.999	2.034




(a) Comparison of the chord distributions along blade span



(b) Comparison of the twist distributions along blade span



(c) Comparison of the thrust distributions along blade span



(d) Comparison of the power distributions along blade span

FIGURE 7.22: Comparison of the baseline and the probabilistic optimum

The sectional geometrical values are compared in Figures 7.22a and 7.22b. The chord distribution of the optimal propeller blade design is close to the lower boundary limit of the design space. The obtained optimal twist distribution resembles an S-shape with a wide and plateau-like first part and a steeply decreasing second part with rapid flattening at the tip of the blade. The sectional

thrust and power distributions of the baseline and the optimal design are depicted in Figures 7.22c and 7.22d. The smaller sectional chord length values result in smaller thrust and power coefficient values. However, the higher sectional twist angles in the second half of the blade are compensated with the decreased thrust coefficient values due to the shorter sectional chord length values. This results in a significant power coefficient reduction 18.18% at the expense of reducing the thrust coefficient only by 12.36%.

The predicted aerodynamic coefficient distributions indicate that the optimal design has a significantly smaller power coefficient compared to the baseline while the probability of violating the thrust constraint is kept at a predefined acceptable reliability level (Figure 7.23). The optimal design also performs more robustly stemming from the fact that left-tail of the thrust coefficient distribution is constrained and the right-tail of the power distribution is minimised simultaneously.




FIGURE 7.23: Comparison of the baseline and the probabilistic optimum of the constrained single-objective propeller blade optimisation

7.5.3 Conclusion of the constrained single-objective propeller blade optimisation study

This section discussed a constrained single-objective optimisation case study of a propeller blade under uncertain environmental conditions. The proposed single-objective multi-fidelity surrogate-based design optimisation approach was able to tackle the constraint optimisation problem with limited computational resources. The obtained propeller blade design satisfies the reliability constraint imposed on the thrust and significantly reduced the required power consumption of the propeller.

“The future cannot be predicted, but futures can be invented.”

— Dennis Gabor

8.1 Conclusions

This dissertation presents a novel multi-fidelity surrogate-assisted design optimisation approach targeted for computationally expensive aerospace applications under uncertainty.

In aerospace engineering, the relative computational burden of virtual aerodynamic analyses is rapidly increasing despite the increasing availability of computational resources. This hinders the applicability of direct uncertainty quantification techniques and the exploration of robust and reliable designs.

The proposed optimisation framework overcomes the challenges of probabilistic design optimisation of computationally expensive problems and is capable of finding designs with optimal statistical performance for both single- and multi-objective problems, as well as constrained problems. Our approach performs the design optimisation with a limited computational budget thanks to the integrated multi-fidelity surrogates for design exploration and uncertainty quantification. The design optimisation is realised following the principles of Bayesian optimisation. The acquisition function balances exploration and exploitation of the design space and allocates the available budget efficiently considering the cost and accuracy of the fidelity levels.

To validate the proposed optimisation framework, available multi-fidelity test functions were tailored for benchmarking problems under uncertainty. The benchmarks showed that it is profitable to use multi-fidelity surrogates when the computational budget is too limited to allow for the construction of an accurate surrogate with high-fidelity simulations but is large enough to generate a great number of low-fidelity data. The type of error present in the low-fidelity simulation code is non-negligible. Surrogate training on data resulting from instable simulations results in potentially biased predictions which eventually drive the optimisation to sub-optimal designs. The applicability of a multi-fidelity optimisation framework highly depends on the good correlation between the low-fidelity models and their high-fidelity counterparts.

To investigate the applicability of the proposed optimisation framework to aerospace applications, we performed optimisation studies of a propeller blade airfoil and a 3D propeller blade. The studies revealed the advantage of multi-fidelity surrogate-assisted optimisation. The otherwise impractical optimisation under uncertainty with high-fidelity RANS simulations is rendered to be feasible by using low-fidelity information, like potential flow solver predictions or RANS simulations with coarse mesh.

We considered various optimisation problems: multi-objective airfoil optimisation under geometric uncertainties, single-objective airfoil optimisation under uncertain environmental conditions, multi-objective and constrained single-objective propeller blade optimisation under uncertain environmental conditions.

The multi-objective airfoil optimisation study considering geometric uncertainties showed that an accurate Pareto front approximation can be found even with a limited computational budget combining information from XFOIL and the RANS solver of SU2. The optimal airfoil designs found by the proposed approach display significant dominance over the baseline solution in the objective space of the reliability measures of the lift and drag. By optimising the opposite tails of the conflicting lift and drag distributions, the optimal solutions are likely to be more robust as well.

The study of a constrained single-objective airfoil optimisation under uncertain environmental conditions showed that the computational load of the probabilistic optimisation can be significantly reduced by using low-fidelity aerodynamic force predictions

of RANS simulations with a coarser mesh. This is particularly true, when a compensation technique, such as the far-field drag prediction formula, exists to provide an accurate drag prediction despite the coarse spatial discretisation.

We found that a multi-objective propeller blade optimisation under uncertain environmental conditions can be efficiently carried out with the proposed multi-fidelity approach. Our strategy allocated the available computational resources efficiently such that the exploration of the propeller designs is performed with XROTOR and only the most promising blade designs are evaluated with the RANS solver of SU2.

The discussion on a constrained single-objective propeller blade optimisation considering environmental uncertainties found that the proposed multi-fidelity approach can find a propeller blade design with significantly reduced power consumption while respecting the imposed reliability constraint on the thrust.

8.2 Possible developments

The proposed multi-fidelity surrogate-assisted optimisation approach was developed with a modular structure which facilitates any extensions and further improvements of the approach.

The extensions of GPR can be also applied for MFGPR such as including the gradient information in the surrogate training process (Gradient-Enhanced GPR). To solve high-dimensional problems, this framework could be complemented with dimensionality reduction techniques, such as *principal component analysis*. Nevertheless, these extensions require caution and additional considerations due to the integration of low-fidelity information. For instance, the gradient information obtained from low-fidelity data can be severely inaccurate.

In this dissertation an auto-regressive implementation of *Multi-Fidelity Gaussian Process Regression* (MFGPR) was presented. It would be interesting to investigate the performance of the optimisation framework with different MFGPR formulation. For instance, a formulation based on *Gaussian Markov Random Fields* is discussed in Appendix A.1. This formulation offers a computationally cheaper training process benefiting from the sparsity of the inverted covariance matrix.

Alternatively, the application of *deep Gaussian processes* [Korondi et al., 2019a] would be also interesting in a multi-fidelity design optimisation context, since deep learning techniques could model nonlinear transformation between the fidelity levels.

“Adventure is worthwhile in itself.”

— Amelia Earhart

A.1 An alternative Gaussian process regression formulation

This section discusses the work published in [Korondi et al., 2019b]. An alternative multi-fidelity surrogate model is given here exploiting the mathematical properties of *Markov random fields*. Gaussian regression models, such as *Gaussian Process Regression* (GPR) and *Multi-Fidelity Gaussian Process Regression* (MFGPR), scale with the number of observations. The basic complexity of Gaussian processes is $\mathcal{O}(N^3)$ where N is the number of data points, stemming from inverting an $N \times N$ matrix. Therefore the *Gaussian Markov Random Field* (GMRF) models are introduced which assume that a random variable associated with a local region depends primarily on its neighbours. Gaussian regression models approximate a sampled function as a random process with Gaussian probability density function.

A.1.1 Gaussian Markov random field

A *random process* can be characterised as a *Markov process* when the conditional probability distribution of future states of a random process depends only on the present state (i.e. the *Markov property* is satisfied). In a multi-dimensional space, a random process is

commonly referred as *random field* and a Markov process as *Markov random field*.

A Gaussian process (i.e. random field) is a random process with Gaussian probability distribution. When a Gaussian process also satisfies the Markov property, it is a GMRF. Following [Rue and Held, 2005], when the realisations of a Gaussian process $z(x) \sim \mathcal{N}(\mu, \Sigma)$ depends only on the neighbouring realisations the Markov property is satisfied. Therefore, $z(x)$ is a GMRF when:

$$P(z_i|z_{-i}) = P(z_i|\{z_j|j \in N_i\}), \quad (\text{A.1})$$

where $\{z_j|j \in N_i\}$ is the set of neighbouring realisations, $z_i = z(s_i)$ and $z_{-i} = (z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_n)$. As GMRF is still a Gaussian process it can be specified by its mean and covariance matrix. Since the mean influences the global properties (trend) of the distribution, the additionally satisfied Markov property is fully contained in the covariance matrix. By investigating the inverse of the covariance matrix, it reveals that Σ^{-1} has non-zero elements when the corresponding realisations are neighbours. Therefore, the *precision matrix* $Q = \Sigma^{-1}$ is defined. This allows to specify a GMRF by its mean and precision [Rue and Held, 2005]:

$$\mathbb{E}(z_i|z_{-i}) = \mu_i - \frac{1}{Q_{ii}} \sum_{j:j \in N_i} Q_{ij}(x_j - \mu_j), \quad (\text{A.2})$$

$$\text{Prec}(z_i|z_{-i}) = Q_{ii}, \quad (\text{A.3})$$

$$\text{Corr}(z_i, z_j|z_{-ij}) = \frac{-Q_{ij}}{\sqrt{Q_{ii}Q_{jj}}} \quad (\text{A.4})$$

The joint PDF of a GMRF denoted by $z(x) \sim \mathcal{N}(\mu, Q^{-1})$ can be formulated as:

$$\text{PDF}(z) = (2\pi)^{-n/2} |Q|^{1/2} \exp\left(-\frac{1}{2}(z - \mu)^T Q (z - \mu)\right). \quad (\text{A.5})$$

where Q is a positive definite matrix.

Consider a GMRF in a form of:

$$z = \begin{pmatrix} z_A \\ z_B \end{pmatrix}, \quad (\text{A.6})$$

with mean:

$$\mu_{\text{joint}} = \begin{bmatrix} \mu_A \\ \mu_B \end{bmatrix}, \quad (\text{A.7})$$

and precision:

$$Q_{\text{joint}} = \begin{bmatrix} Q_{AA} & Q_{AB} \\ Q_{BA} & Q_{BB} \end{bmatrix}. \quad (\text{A.8})$$

Then this GMRF can be transformed into a conditional distribution. The conditional distribution is also a GMRF and given by its mean and precision:

$$\mathbb{E}(z_A|z_B) = \mu_A - Q_{AA}^{-1}Q_{AB}(z_B - \mu_B) \quad (\text{A.9})$$

$$\text{Prec}(z_A|z_B) = Q_{AA} \quad (\text{A.10})$$

At this point the question is how to construct a sparse precision matrix such that the corresponding GMRF can capture the characteristics of the functions which we would like to surrogate.

It has been shown that a Gaussian process with Matérn covariance is a *stationary* solution to the following *Stochastic Partial Differential Equation* (SPDE) [Whittle, 1954, 1963]:

$$(\chi^2 - \Delta)^{\alpha/2} z(x) = W(x), \quad (\text{A.11})$$

where $W(s)$ is a white noise, $\Delta = \sum_i \frac{\partial^2}{\partial s_i^2}$ is the Laplacian operator and $\alpha = \nu + d/2$. The parameter ν controls the smoothness and the parameter χ controls the range. From [Rozanov, 1977], when α is an integer the solution satisfies the Markov property, hence, it can be modelled as a GMRF.

The solution can be approximated as a finite basis expansion [Lindgren et al., 2011]:

$$z(x) = \sum_k \varphi_k(x) w_k, \quad (\text{A.12})$$

with a suitable distribution for the weights w_k . Then, a stochastic weak solution to the SPDE is given by:

$$\langle \psi_j, (\chi^2 - \Delta)^{\alpha/2} z(x) \rangle = \langle \psi_j, W \rangle \quad \forall j. \quad (\text{A.13})$$

where $\langle f, g \rangle = \int f(x)g(x)dx$ is the inner product over the region of interest. The test functions are chosen according the *least squares* finite element approximation $\psi_k = (\chi^2 - \Delta)^{1/2} \varphi_k$ when $\alpha = 1$. For $\alpha > 2$, the *Galerkin* solution is employed $\psi_k = \varphi_k$. However, when $\alpha \geq 3$ the Eq. (A.11) is modified by setting the $\alpha = 2$ on

the left-hand side and the right-hand side is replaced with a field generated by $\alpha - 2$. This recursive Galerkin formulation then terminates in $\alpha = 1$ or $\alpha = 2$.

With the explicit link between Gaussian processes and GMRF [Lindgren et al., 2011], the precision matrix for weights w for $\alpha = 1, 2, \dots$ is:

$$Q_1 = K, \quad (\text{A.14})$$

$$Q_2 = KM^{-1}K, \quad (\text{A.15})$$

$$Q_\alpha = KM^{-1}Q_{\alpha-2}M^{-1}K, \quad (\text{A.16})$$

where M and K are required to be sparse. They are defined as:

$$M_{ij} = \langle \varphi_i, \varphi_j \rangle, \quad (\text{A.17})$$

$$S_{ij} = \langle \nabla \varphi_i, \nabla \varphi_j \rangle, \quad (\text{A.18})$$

$$K_{ij} = \chi^2 M_{ij} + S_{ij}, \quad (\text{A.19})$$

Additionally, to keep the sparsity of the precision matrix, the M matrix is replaced with a diagonal matrix \tilde{M} having $\tilde{M}_{ii} = \langle \varphi_i, 1 \rangle$ entries. This introduces an additional approximation error but assures the sparsity of the precision matrix.

A.1.2 Gaussian Markov Random Field Regression

The unknown function which we would like to surrogate is assumed to be a linear transformation of a GMRF w and some additional noise:

$$z(x) = \varphi^T(x)w + \mathcal{E}(x), \quad (\text{A.20})$$

with

$$w \sim \mathcal{N}(\mu_w, Q_w^{-1}), \quad (\text{A.21})$$

$$\mathcal{E}(x) \sim \mathcal{N}(0, \sigma_\epsilon^2), \quad (\text{A.22})$$

where $\phi^T w$ is a spatial basis expansion with k basis functions, for instance, linear finite element basis. The noise term $\mathcal{E}(x)$ is normally distributed with constant variance σ_ϵ^2 . In this model, w is a latent field as we cannot observe it, but only its linear transformation with some noise.

Given the observation vector $z^{(n)} = (z_1, \dots, z_n)^T$ containing the realisations of a $z(x) = \varphi^T(x)w + \mathcal{E}(x)$ at $x^{(n)}$ observation

locations and the noise vector $\mathcal{E}^{(n)} = (\sigma_\epsilon \mathcal{E}_1, \dots, \sigma_\epsilon \mathcal{E}_n)^T$ with $\mathcal{E}_{i=1, \dots, n}$ independent and identically distributed standard normal realisations, the following hierarchical model can be written:

$$(z^{(n)} | w) \sim \mathcal{N}(\Phi w, Q_\epsilon^{-1}), \quad (\text{A.23})$$

$$w \sim \mathcal{N}(\mu_w, Q_w^{-1}), \quad (\text{A.24})$$

where $Q_\epsilon^{-1} = \sigma_\epsilon^2 I$ is the $n \times n$ covariance matrix of observations. Φ is a $k \times n$ matrix with the values of the basis functions.

The joint distribution for the observations and the latent field w forms a GMRF [Simpson et al., 2012] with mean:

$$\mathbb{E} \begin{pmatrix} w \\ z^{(n)} \end{pmatrix} = \begin{pmatrix} \mu_w \\ \Phi \mu_w \end{pmatrix}, \quad (\text{A.25})$$

and precision:

$$\text{Prec} \begin{pmatrix} w \\ z^{(n)} \end{pmatrix} = \begin{bmatrix} Q_w + \Phi^T Q_\epsilon \Phi & -\Phi^T Q_\epsilon \\ -Q_\epsilon \Phi & Q_\epsilon \end{bmatrix}. \quad (\text{A.26})$$

The conditional distribution for $(w | z^{(n)})$ can be described by Eqs. (A.9) and (A.10):

$$\mu_{w|z^{(n)}} = \mu_w + Q_{w|z^{(n)}}^{-1} \Phi^T Q_\epsilon (z^{(n)} - \Phi \mu_w), \quad (\text{A.27})$$

$$Q_{w|z^{(n)}} = Q_w + \Phi^T Q_\epsilon \Phi, \quad (\text{A.28})$$

where $\mu_{w|z^{(n)}} = \mathbb{E}(w | z^{(n)})$ and $Q_{w|z^{(n)}} = \text{Prec}(w | z^{(n)})$ are the mean and the precision of the conditional distribution. Then, the variance can be computed as $s_{w|z^{(n)}}^2 = \text{diag}(Q_{w|z^{(n)}}^{-1})$. The elements of $m_{w|z^{(n)}}$ are the basis function coefficients and covariate effect estimates in the GMRF regression predictor:

$$\hat{m}_z(x) = \varphi(x) \mu_{w|z^{(n)}} \quad (\text{A.29})$$

with squared error:

$$\hat{s}_z^2(s) = \text{diag}(\varphi(x) Q_{w|z^{(n)}}^{-1} \varphi^T(x)). \quad (\text{A.30})$$

The hyper-parameters of the precision matrix ($\theta = \{\nu, \chi\}$) can be found with the MLE approach. The likelihood for w given the

parameters θ is:

$$\mathcal{L}(w|\theta) = \frac{1}{(2\pi)^{\frac{k+p}{2}} \sqrt{|Q_w|}} \exp\left(-\frac{1}{2}(w - \mu_w)^T Q_w (w - \mu_w)\right) \quad (\text{A.31})$$

so that the log-likelihood is:

$$\begin{aligned} \log \mathcal{L}(w|\theta) = & -\frac{k+p}{2} \log(2\pi) + \\ & + \frac{1}{2} \log |Q_w| - \frac{1}{2} (w - \mu_w)^T Q_w (w - \mu_w). \end{aligned} \quad (\text{A.32})$$

For known $w = \hat{w}$, the likelihood for $z^{(n)}$ given the parameters θ is:

$$\mathcal{L}(z^{(n)}|\theta) = \frac{\mathcal{L}(\theta|z^{(n)})}{\mathcal{L}(\theta)} = \frac{\mathcal{L}(w|\theta) \mathcal{L}(z^{(n)}|\theta, w)}{\mathcal{L}(w|\theta, z^{(n)})} \Bigg|_{w=\hat{w}} \quad (\text{A.33})$$

so that the log-likelihood is:

$$\begin{aligned} \log \mathcal{L}(z^{(n)}|\theta) = & \log \mathcal{L}(\hat{w}|\theta) + \log \mathcal{L}(z^{(n)}|\theta, \hat{w}) - \log \mathcal{L}(\hat{w}|\theta, z^{(n)}) = \\ & -\frac{k+p}{2} \log(2\pi) + \frac{1}{2} \log |Q_w| - \frac{1}{2} (\hat{w} - \mu_w)^T Q_w (\hat{w} - \mu_w) \\ & -\frac{n}{2} \log(2\pi) + \frac{1}{2} \log |Q_\epsilon| - \frac{1}{2} (z^{(n)} - \Phi \hat{w})^T Q_\epsilon (z^{(n)} - \Phi \hat{w}) \\ & + \frac{k+p}{2} \log(2\pi) - \frac{1}{2} \log |Q_{w|z^{(n)}}| + \\ & \frac{1}{2} (\hat{w} - \mu_{w|z^{(n)}})^T Q_{w|z^{(n)}} (\hat{w} - \mu_{w|z^{(n)}}). \end{aligned} \quad (\text{A.34})$$

The likelihood for $z^{(n)}$ given the parameters θ can be evaluated with $\hat{w} = \mu_{w|z^{(n)}}$, so that:

$$\begin{aligned} \log \mathcal{L}(z^{(n)}|\theta) = & -\frac{n}{2} \log(2\pi) + \frac{1}{2} \log |Q_w| + \frac{1}{2} \log |Q_\epsilon| - \frac{1}{2} \log |Q_{w|z^{(n)}}| \\ & - \frac{1}{2} (\mu_{w|z^{(n)}} - \mu_w)^T Q_w (\mu_{w|z^{(n)}} - \mu_w) \\ & - \frac{1}{2} (z^{(n)} - \Phi \mu_{w|z^{(n)}})^T Q_\epsilon (z^{(n)} - \Phi \mu_{w|z^{(n)}}). \end{aligned} \quad (\text{A.35})$$

A.1.3 Multi-fidelity Gaussian Markov random field regression

The GMRF regression can be extended to train a surrogate with multi-fidelity data. Similarly to the recursive MFGPR, there are L levels of response $z_l(x)$ sorted by increasing order of fidelity (i.e. $l = 1$ is the lowest and $l = L$ is the highest fidelity level). An auto-regressive model using GMRF can be formulated starting from the 2nd level:

$$\begin{cases} z_l(x) = \varphi^T(x)w_l + \mathcal{E}_l(x), \\ w_l = \rho_{l-1}^T w_{l-1} + \delta_l, \\ w_{l-1} \perp \delta_l, \end{cases} \quad (\text{A.36})$$

where δ_l is a GMRF with mean μ_{w_l} and precision matrix Q_{w_l} . ρ_{l-1} is the scale factor relating the consecutive fidelity levels.

Given $z_l^{(n_l)}$ observed values of $z_l(x)$, the following hierarchical model can be written:

$$\left(z_l^{(n_l)} \middle| w_l \right) - \rho_{l-1}^T \odot \varphi^T(z_l^{(n_l)})w_{l-1} \sim \mathcal{N}(\Phi_l \delta_l, Q_{\epsilon_l}^{-1}), \quad (\text{A.37})$$

$$\delta_l \sim \mathcal{N}(\mu_{w_l}, Q_{w_l}^{-1}), \quad (\text{A.38})$$

where $Q_{\epsilon_l}^{-1} = \sigma_{\epsilon_l}^2 I$ is the $n_l \times n_l$ covariance matrix of observations and \odot is the *Hadamard product* (i.e. element-wise product). The joint distribution for the observations and the latent variables w_l is given by:

$$\left(\left(z_l^{(n_l)} \middle| w_l \right) - \rho_{l-1}^T \odot \varphi^T(z_l^{(n_l)})w_{l-1} \right) \sim \quad (\text{A.39})$$

$$\mathcal{N} \left(\begin{pmatrix} \mu_{w_l} \\ \Phi_l \mu_{w_l} \end{pmatrix}, \begin{bmatrix} Q_{w_l} + \Phi_l^T Q_{\epsilon_l} \Phi_l & -\Phi_l^T Q_{\epsilon_l} \\ -Q_{\epsilon_l} \Phi_l & Q_{\epsilon_l} \end{bmatrix}^{-1} \right). \quad (\text{A.40})$$

The conditional distribution for $(w_l | z_l^{(n_l)})$ can be described with:

$$\begin{aligned}
m_{w_l | z_l^{(n_l)}} &= \rho_{l-1}^T m_{w_{l-1} | z_{l-1}^{(n_{l-1})}} + \mu_{w_l | z_l^{(n_l)}}, \\
\mu_{w_l | z_l^{(n_l)}} &= \mu_{w_l} + \\
&+ Q_{w_l | z_l^{(n_l)}}^{-1} \Phi_l^T Q_{\epsilon_l} \left(z_l^{(n_l)} - \rho_{l-1}^T \odot \varphi^T(z_l^{(n_l)}) m_{w_{l-1} | z_{l-1}^{(n_{l-1})}} - \Phi_l \mu_{w_l} \right), \\
Q_{w_l | z_l^{(n_l)}} &= Q_{w_l} + \Phi_l^T Q_{\epsilon_l} \Phi_l, \\
\Sigma_{w_l | z_l^{(n_l)}} &= \rho_{l-1}^T \Sigma_{w_{l-1} | z_{l-1}^{(n_{l-1})}} \rho_{l-1} + Q_{w_l | z_l^{(n_l)}}^{-1}.
\end{aligned} \tag{A.41}$$

The elements of $m_{w_l | z_l^{(n_l)}}$ are the basis function coefficients and covariate effect estimates in the multi-fidelity GMRF regression predictor at l level:

$$\hat{m}_{z_l}(x) = \varphi(x) m_{w_l | z_l^{(n_l)}}, \tag{A.42}$$

with squared error:

$$\hat{s}_{z_l}^2(x) = \text{diag} \left(\varphi(x) \Sigma_{w_l | z_l^{(n_l)}} \varphi^T(x) \right). \tag{A.43}$$

The hyper-parameters of the precision matrix can be found with the MLE approach. The likelihood for $z_l^{(n_l)}$ given the parameters θ_l is:

$$\begin{aligned}
\log \mathcal{L} \left(z_l^{(n_l)} \middle| \theta_l \right) &= \\
&- \frac{n}{2} \log(2\pi) + \frac{1}{2} \log |Q_{w_l}| + \frac{1}{2} \log |Q_{\epsilon_l}| - \frac{1}{2} \log |Q_{w_l | z_l^{(n_l)}}| \\
&- \frac{1}{2} \left(\mu_{w_l | z_l^{(n_l)}} - \mu_{w_l} \right)^T Q_{w_l} \left(\mu_{w_l | z_l^{(n_l)}} - \mu_{w_l} \right) \\
&- \frac{1}{2} \left(z_l^{(n_l)} - \rho_{l-1}^T \odot \varphi^T(z_l^{(n_l)}) m_{w_{l-1} | z_{l-1}^{(n_{l-1})}} - \Phi_l^T m_{X_l | z_l^{(n_l)}} \right)^T Q_{\epsilon_l} \\
&\quad \left(z_l^{(n_l)} - \rho_{l-1}^T \odot \varphi^T(z_l^{(n_l)}) m_{w_{l-1} | z_{l-1}^{(n_{l-1})}} - \Phi_l^T m_{X_l | z_l^{(n_l)}} \right).
\end{aligned} \tag{A.44}$$

A.1.4 Comparison of multi-fidelity Gaussian process regression and multi-fidelity Gaussian Markov random field regression

A simple one-dimensional problem is investigated in this section. The test function for multi-fidelity surrogates is presented in [Forrester et al., 2007]. The high- and low-fidelity functions are the following:

$$f_{high} = (6x - 2)^2 \sin(12x - 4), \quad (\text{A.45})$$

$$f_{low} = \frac{1}{2}f_{high} + 10(x - 0.5) - 5. \quad (\text{A.46})$$

In this case four observations are available at the high-fidelity level $X_{high} = \{0, 0.4, 0.6, 1\}$ and eleven at the low-fidelity level $X_{low} = \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$. The surrogate built-on variable fidelity data is depicted in Figure A.1.




FIGURE A.1: Multi-fidelity GMRF regression.

Function	f_{low}	f_{low}	f_{high}	f_{high}
Surrogate	GPR	GMRF	MFGPR	MF-GMRF
Error	0.0389	0.0459	0.0852	0.1255

Table A.1: Comparison of MFGPR and multi-fidelity GMRF regression. (The mean absolute error is calculated with 100 instances.)

The result clearly shows that the single-fidelity learning technique GPR is not able to capture correctly the function landscape due to the limited number of observation points. The multi-fidelity learning technique is able to fuse the information from the low-fidelity function into the high-fidelity approximation and thus provides an adequate approximation of the true function. The multi-fidelity learning technique with GMRF is not able to properly learn the function landscape at the domain boundaries because *Neumann boundary conditions* with zero value are assumed. This is also confirmed by Figure A.2. This results in a slightly higher approximation error compared to MFGPR as it can be seen in Table A.1.

Nevertheless, MF-GMRF regression provides a viable alternative to MFGPR.




FIGURE A.2: Comparison of the prediction error distribution of MFGPR and MF-GMRF.

A.2 Space systems resilience optimisation under epistemic uncertainty

This section summarises a collaboration work published in [Paper F: Filippi et al.,2020].

The design optimisation of a space system requires a different perspective. Our understanding of the space is significantly limited compared to our knowledge of the atmosphere of the Earth. Additionally, the high complexity and costly physical experimenting of space systems increase the uncertainty of a space mission. Therefore, the uncertainty must be treated with more conservative approaches considering the presence of significant epistemic uncertainty. The entire design process of a space mission can last for several years. Traditionally, space systems are designed with well-established safety margins and redundancies. Unfortunately, these traditional approaches may overestimate or underestimate the effect of uncertainties.

In [Paper F: Filippi et al.,2020], a *resilience-based design optimisation* technique is proposed for designing space systems under epistemic uncertainty. The resilience-based approach is introduced through the optimisation of a five-component *CubeSat*. The five components are: *Attitude and Orbit Control*, *Telecommunication*, *On Board Data Handling*, *Power* and *Payload*. The *Quantities of Interest* are the total mass (M_{TOT}) of the satellite and the total amount of transmitted data to the ground station (V). The CubeSat is designed considering 12 design variables and 20 epistemic uncertain variables. Additionally, the amount of transmitted data is calculated with a stochastic dynamic model of the space system considering time dependence and 3 operational states: *system failure*, *partially functional* and *fully functional system*.

The optimisation is formulated as a *constrained min-max* problem. The optimal design parameters are sought which minimise the total mass of the satellite in a worst-case scenario (i.e. maximal mass value among all possible total mass values due to the uncertainty of a fixed design configuration) such that the amount of transmitted data is not smaller than a threshold value.

The resilience-based design optimisation of the satellite is performed with the so-called *Evidence-Based Robust Optimisation* technique exploiting the properties of *Evidence Network Models* and the principles of *Dempster-Shafer Theory*. The *cumulative belief* (i.e. lower bound of the CDF) of the total mass of the obtained resilience-based design is depicted with blue colour in Figure A.3. This resilience-based solution ensures that total mass of the space system will be below 13.44 when the transmitted data volume is at least 600 GB. The design solution obtained with the margin-based approach predicts that the total mass will be 13.17 as represented by the vertical red line in Figure A.3. However, the margin-based approach fails to properly quantify the uncertainty. By evaluating the uncertainty of the design obtained with the margin approach, it is revealed that the evidence supports only a low belief ($Bel=0.025$) that the satellite will have a mass of 13.17. Correspondingly, the total mass of the satellite designed with the margin approach can reach as high a value as 14.33 which is significantly heavier than the worst-case of the resilience-based optimum.




FIGURE A.3: Comparison of the margin-based and the resilience-based optimum when the transmitted data volume is at least 600 GB as published in [Paper F: Filippi et al.,2020].

Bibliography

- Adler, R. J. and Taylor, J. E. (2009). *Random fields and geometry*. Springer Science & Business Media.
- Alexandrov, N., Lewis, R., Gumbert, C., Green, L., and Newman, P. (2000). Optimization with variable-fidelity models applied to wing design. In *38th aerospace sciences meeting and exhibit*, page 841.
- Ariyarit, A. and Kanazaki, M. (2017). Multi-fidelity multi-objective efficient global optimization applied to airfoil design problems. *Applied Sciences*, 7(12):1318.
- Arizmendi, B., Bellosta, T., del Val, A. I., Gori, G., Prazeres, M. O., and Reis, J. (2019). On real-time management of on-board ice protection systems by means of machine learning. In *AIAA Aviation 2019 Forum*, page 3464.
- Bak, C., Anderson, P., Madsen, H., Gaunaa, M., Fuglsang, P., and Bove, S. (2008). Design and verification of airfoils resistant to surface contamination and turbulence intensity. In *26th AIAA applied aerodynamics conference*, page 7050.
- Bandler, J. W. (1969). Optimization methods for computer-aided design. *IEEE Transactions on Microwave Theory and Techniques*, 17(8):533–552.
- Barrett, R. and Ning, A. (2016). Comparison of airfoil precomputational analysis methods for optimization of wind turbine blades. *IEEE Transactions on Sustainable Energy*, 7(3):1081–1088.
- Betz, A. (1919). Airscrews with minimum energy loss. *Report, Kaiser Wilhelm Institute for Flow Research*.

- Beyer, H.-G. and Sendhoff, B. (2007). Robust optimization—a comprehensive survey. *Computer methods in applied mechanics and engineering*, 196(33-34):3190–3218.
- Black, D. and Rohrbach, C. (1968). Shrouded propellers—a comprehensive performance study. In *5th Annual Meeting and Technical Display*, page 994.
- Bonfiglio, L., Perdikaris, P., Brizzolara, S., and Karniadakis, G. (2018). Multi-fidelity optimization of super-cavitating hydrofoils. *Computer Methods in Applied Mechanics and Engineering*, 332:63–85.
- Bouhlef, M. A., Hwang, J. T., Bartoli, N., Lafage, R., Morlier, J., and Martins, J. R. (2019). A python surrogate modeling framework with derivatives. *Advances in Engineering Software*, 135:102662.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799.
- Chaudhuri, A., Marques, A. N., and Willcox, K. E. (2019). mfegra: Multifidelity efficient global reliability analysis. *arXiv preprint arXiv:1910.02497*.
- Choi, S.-K., Grandhi, R. V., Canfield, R. A., and Pettit, C. L. (2004). Polynomial chaos expansion with latin hypercube sampling for estimating response variability. *AIAA journal*, 42(6):1191–1198.
- Construct2D (2018). v2.0.4 [software]. <https://sourceforge.net/projects/construct2d/>.
- Deb, K. (2014). Multi-objective optimization. In *Search methodologies*, pages 403–449. Springer.
- Deb, K., Agrawal, S., Pratap, A., and Meyarivan, T. (2000). A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *International conference on parallel problem solving from nature*, pages 849–858. Springer.
- Deb, K., Thiele, L., Laumanns, M., and Zitzler, E. (2005). Scalable test problems for evolutionary multiobjective optimization. In *Evolutionary multiobjective optimization*, pages 105–145. Springer.

- Drela, M. (1989). Xfoil: An analysis and design system for low reynolds number airfoils. In *Low Reynolds number aerodynamics*, pages 1–12. Springer.
- Drela, M. (2001). Xfoil 6.9 user primer, xfoil_doc.txt [software]. Accessed: 2020-06-01.
- Drela, M. and Youngren, H. (2013). Xrotor [software]. URL: <http://web.mit.edu/drela/Public/web/xrotor>.
- Economou, T. D., Palacios, F., Copeland, S. R., Lukaczyk, T. W., and Alonso, J. J. (2016). Su2: An open-source suite for multi-physics simulation and design. *Aiaa Journal*, 54(3):828–846.
- El-Sayed, A. F. (2008). *Aircraft propulsion and gas turbine engines*. CRC press.
- Emmerich, M. T., Giannakoglou, K. C., and Naujoks, B. (2006). Single-and multiobjective evolutionary optimization assisted by gaussian random field metamodels. *IEEE Transactions on Evolutionary Computation*, 10(4):421–439.
- Fernández-Godino, G. M., Park, C., Kim, N. H., and Haftka, R. T. (2019). Issues in deciding whether to use multifidelity surrogates. *AIAA Journal*, 57(5):2039–2054.
- Filippi, G., Vasile, M., Krpelik, D., Korondi, P. Z., Marchi, M., and Poloni, C. (2019). Space systems resilience optimisation under epistemic uncertainty. *Acta Astronautica*, 165:195–210.
- Forrester, A., Sobester, A., and Keane, A. (2008). *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons.
- Forrester, A. I., Sobester, A., and Keane, A. J. (2007). Multifidelity optimization via surrogate modelling. *Proceedings of the royal society a: mathematical, physical and engineering sciences*, 463(2088):3251–3269.
- Frazier, P. I. (2018). A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*.
- FreeCAD (2018). v0.18.4 [software].

- Gaggero, S., Villa, D., Tani, G., Viviani, M., and Bertetta, D. (2017). Design of ducted propeller nozzles through a ranse-based optimization approach. *Ocean Engineering*, 145:444–463.
- Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., and Cunningham, J. P. (2014). Bayesian optimization with inequality constraints. In *ICML*, volume 2014, pages 937–945.
- Gittins, J. C. and Jones, D. M. (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66(3):561–565.
- Glauert, H. (1983). *The elements of aerofoil and airscrew theory*. Cambridge University Press.
- Habib, A., Singh, H. K., and Ray, T. (2019). A multiple surrogate assisted multi/many-objective multi-fidelity evolutionary algorithm. *Information Sciences*, 502:537–557.
- Han, Z.-H., Görtz, S., and Zimmermann, R. (2013). Improving variable-fidelity surrogate modeling via gradient-enhanced kriging and a generalized hybrid bridge function. *Aerospace Science and technology*, 25(1):177–189.
- Hansen, M. O. (2015). *Aerodynamics of wind turbines*. Routledge.
- Hansen, N. (2016). The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*.
- Helton, J. C., Johnson, J. D., Oberkampf, W. L., and Sallaberry, C. J. (2010). Representation of analysis results involving aleatory and epistemic uncertainty. *International Journal of General Systems*, 39(6):605–646.
- Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A. (2006). Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32(5):369–382.
- Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492.

- Kandasamy, K., Dasarathy, G., Oliva, J. B., Schneider, J., and Póczos, B. (2016). Gaussian process bandit optimisation with multi-fidelity evaluations. In *Advances in Neural Information Processing Systems*, pages 992–1000.
- Keane, A. J. (2006). Statistical improvement criteria for use in multiobjective design optimization. *AIAA journal*, 44(4):879–891.
- Kennedy, M. C. and O’Hagan, A. (2000). Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13.
- Klein, P. C. (2017). Parametric modeling and optimization of advanced propellers for next-generation aircraft. *TU Delft*.
- Knowles, J. (2006). Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66.
- Korondi, P. Z. (2016). Advanced constraint handling in crash-worthiness design optimization with hybrid cellular automata. Master’s thesis, Technical University of Munich (TUM).
- Korondi, P. Z., Elisa, M., and Christian, S. (2019a). Deepg-plaen: Deep gaussian process learning for aerospace engineering. <http://utopiae.eu/2-2/utopiae-training/outreach/ltw-ii-outreach-programme/> [Project in progress].
- Korondi, P. Z., Marchi, M., Parussini, L., and Poloni, C. (2020a). Multi-fidelity design optimisation strategy under uncertainty with limited computational budget. *Optimization and Engineering*.
- Korondi, P. Z., Marchi, M., Parussini, L., Quagliarella, D., and Poloni, C. (2020b). Multi-objective design optimisation of an airfoil with geometrical uncertainties leveraging multi-fidelity gaussian process regression. In *UQOP:International Conference on Uncertainty Quantification & Optimisation*.
- Korondi, P. Z., Marchi, M., and Poloni, C. (2020c). *Optimization Under Uncertainty with Applications to Aerospace Engineering: Response Surface Methodolgy (Ch.12)*. Springer.

- Korondi, P. Z., Parussini, L., Marchi, M., and Poloni, C. (2019b). Reliability-based design optimisation of a ducted propeller through multi-fidelity learning. In *UNCECOMP: 3rd International Conference on Uncertainty Quantification in Computational Sciences and Engineering Proceedings*.
- Kraft, D. and Schnepfer, K. (1989). Slsqp—a nonlinear programming method with quadratic programming subproblems. *DLR, Oberpfaffenhofen*, 545.
- Lam, R., Poloczek, M., Frazier, P., and Willcox, K. E. (2018). Advances in bayesian optimization with applications in aerospace engineering. In *2018 AIAA Non-Deterministic Approaches Conference*, page 1656.
- Le Gratiet, L. (2013). *Multi-fidelity Gaussian process regression for computer experiments*. PhD thesis, Paris Diderot University.
- Lemieux, C. (2009). *Monte carlo and quasi-monte carlo sampling*. Springer Science & Business Media.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.
- Maleki, A. and Do, T. (2000). Review of probability theory. *CS*, 229(2):1.
- Martin, M. (2002). On-line support vector machine regression. In *European Conference on Machine Learning*, pages 282–294. Springer.
- McKay, M. D., Beckman, R. J., and Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61.
- Meliani, M., Bartoli, N., Lefebvre, T., Bouhlef, M.-A., Martins, J., and Morlier, J. (2019). Multi-fidelity efficient global optimization: Methodology and application to airfoil shape design. In *AIAA Aviation 2019 Forum*, page 3236.

- Menter, F. R., Kuntz, M., and Langtry, R. (2003). Ten years of industrial experience with the sst turbulence model. *Turbulence, heat and mass transfer*, 4(1):625–632.
- Montrone, T. and Rigoni, E. (2018). Mego algorithm: General description. Technical report, ESTECO Technical Report 2018-002, Trieste, Italy, 2018.
- Morales, E., Korondi, P. Z., Quagliarella, D., Tognaccini, R., Marchi, M., Parussini, L., and Poloni, C. (2020). Multi-fidelity surrogate assisted design optimisation of an airfoil under uncertainty using far-field drag approximation. In *UQOP: International Conference on Uncertainty Quantification & Optimisation*.
- Morales, E., Quagliarella, D., and Tognaccini, R. (2019). Airfoil optimization using far-field analysis of the drag force. In *AIAA Scitech 2019 Forum*, page 0972.
- Myers, D. E. (1982). Matrix formulation of co-kriging. *Journal of the International Association for Mathematical Geology*, 14(3):249–257.
- Papalambros, P. Y. (2002). The optimization paradigm in engineering design: promises and challenges. *Computer-Aided Design*, 34(12):939–951.
- Perdikaris, P., Raissi, M., Damianou, A., Lawrence, N., and Karniadakis, G. E. (2017). Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2198):20160751.
- Pointwise (2017). User manual v18, [software].
- Poloczek, M., Wang, J., and Frazier, P. (2017). Multi-information source optimization. In *Advances in Neural Information Processing Systems*, pages 4288–4298.
- Poloni, C. (1995). Hybrid ga for multi objective aerodynamic shape optimisation. *Genetic Algorithms in Engineering and Computer Science*.

- Poloni, C., Giurgevich, A., Onesti, L., and Pediroda, V. (2000). Hybridization of a multi-objective genetic algorithm, a neural network and a classical optimizer for a complex design problem in fluid dynamics. *Computer Methods in Applied Mechanics and Engineering*, 186(2-4):403–420.
- Poloni, C. and Pediroda, V. (1997). Ga coupled with computationally expensive simulations: tools to improve efficiency. *Genetic Algorithms and Evolution Strategies in Engineering and Computer Science. Recent Advances and Industrial Applications*, pages 267–288.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press.
- Quagliarella, D. (2016). wg2aer (weights-to aerofoils) [software]. Software developed at the Italian Aerospace Research Centre (CIRA).
- Quagliarella, D. and Diez, M. (2020). An open-source aerodynamic framework for benchmarking multi-fidelity methods. In *AIAA Aviation 2020 Forum*, page 3179.
- Quagliarella, D., Petrone, G., and Iaccarino, G. (2014). Optimization under uncertainty using the generalized inverse distribution function. In *Modeling, Simulation and Optimization for Science and Technology*, pages 171–190. Springer.
- Quagliarella, D., Serani, A., Diez, M., Pisaroni, M., Leyland, P., Montagliani, L., Iemma, U., Gaul, N. J., Shin, J., Wunsch, D., et al. (2019). Benchmarking uncertainty quantification methods using the NACA 2412 airfoil with geometrical and operational uncertainties. In *AIAA Aviation 2019 Forum*, page 3555.
- Quagliarella, D., Tirado, E. M., and Bornaccioni, A. (2020). Risk measures applied to robust aerodynamic shape design optimization. In *Flexible Engineering Toward Green Aircraft*, pages 153–168. Springer.
- Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer.

- Rockafellar, R. T. and Royset, J. O. (2015). Engineering decisions under risk averseness. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 1(2):04015003.
- Rockafellar, R. T. and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of banking & finance*, 26(7):1443–1471.
- Rozanov, J. A. (1977). Markov random fields and stochastic partial differential equations. *Mathematics of the USSR-Sbornik*, 32(4):515.
- Rubinstein, R. Y. and Kroese, D. P. (2016). *Simulation and the Monte Carlo method*, volume 10. John Wiley & Sons.
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. CRC press.
- Sabater, C., Le Maître, O., Congedo, P., and Görtz, S. (2020). A bayesian approach for quantile optimization problems with high-dimensional uncertainty sources. *H2020-MSCA-ITN-2016 UTOPIAE*.
- Sacher, M., Le Maître, O., Duvinneau, R., Hauville, F., Durand, M., and Lothodé, C. (2020). A non-nested infilling strategy for multi-fidelity based efficient global optimization. *International Journal for Uncertainty Quantification*.
- Sareen, A., Sapre, C. A., and Selig, M. S. (2014). Effects of leading edge erosion on wind turbine blade performance. *Wind Energy*, 17(10):1531–1542.
- Schobi, R., Sudret, B., and Wiart, J. (2015). Polynomial-chaos-based kriging. *International Journal for Uncertainty Quantification*, 5(2).
- Schuëller, G. I. and Jensen, H. A. (2008). Computational methods in optimization considering uncertainties—an overview. *Computer Methods in Applied Mechanics and Engineering*, 198(1):2–13.

- Serafino, A., Obert, B., and Cinnella, P. (2020). Multi-fidelity gradient-based strategy for robust optimization in computational fluid dynamics. *Algorithms*, 13(10):248.
- Simpson, D., Lindgren, F., and Rue, H. (2012). Think continuous: Markovian gaussian models in spatial statistics. *Spatial Statistics*, 1:16–29.
- Singer, S. and Nelder, J. (2009). Nelder-mead algorithm. *Scholarpedia*, 4(7):2928.
- Song, W. and Keane, A. (2004). A study of shape parameterisation methods for airfoil optimisation. In *10th AIAA/ISSMO multidisciplinary analysis and optimization conference*, page 4482.
- Stein, M. L. (2012). *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media.
- Swersky, K., Snoek, J., and Adams, R. P. (2013). Multi-task bayesian optimization. In *Advances in neural information processing systems*, pages 2004–2012.
- Takeno, S., Fukuoka, H., Tsukada, Y., Koyama, T., Shiga, M., Takeuchi, I., and Karasuyama, M. (2019). Multi-fidelity bayesian optimization with max-value entropy search. *arXiv preprint arXiv:1901.08275*.
- Tecplot (2019). 360 ex r1 [software].
- Tu, J., Yeoh, G. H., and Liu, C. (2018). *Computational fluid dynamics: a practical approach*. Butterworth-Heinemann.
- Turco, A. (2011). Hybrid-description. Technical report, ESTECO Technical Report 2011-003, Trieste, Italy, 2011.
- Vaithyanathanasamy, R. (2017). Cfd analysis of 2d and 3d airfoils using open source solver su2. *University of Twente, Internship report*.
- Vasile, M. and Zuiani, F. (2011). Multi-agent collaborative search: an agent-based memetic multi-objective optimization algorithm applied to space trajectory design. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 225(11):1211–1227.

- Veldhuis, L. L. M. (2005). *Propeller wing aerodynamic interference*. PhD thesis, Delft University of Technology.
- Von Mises, R. (2014). *Mathematical theory of probability and statistics*. Academic Press.
- Wang, H., Jin, Y., and Doherty, J. (2017). A generic test suite for evolutionary multifidelity optimization. *IEEE Transactions on Evolutionary Computation*, 22(6):836–850.
- Weir, R. J. (1987). *Ducted propeller design and analysis*. Sandia National Laboratories.
- Whale, J., Anderson, C., Bareiss, R., and Wagner, S. (2000). An experimental and numerical study of the vortex structure in the wake of a wind turbine. *Journal of Wind Engineering and Industrial Aerodynamics*, 84(1):1–21.
- Whittle, P. (1954). On stationary processes in the plane. *Biometrika*, pages 434–449.
- Whittle, P. (1963). Stochastic-processes in several dimensions. *Bulletin of the International Statistical Institute*, 40(2):974–994.
- Wolpert, D. H. and Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1):67–82.
- Xiao, M., Zhang, G., Breikopf, P., Villon, P., and Zhang, W. (2018). Extended co-kriging interpolation method based on multi-fidelity data. *Applied Mathematics and Computation*, 323:120–131.
- Zhang, Y., Han, Z.-H., and Zhang, K.-S. (2018). Variable-fidelity expected improvement method for efficient global optimization of expensive functions. *Structural and Multidisciplinary Optimization*, 58(4):1431–1451.
- Zhang, Y., Hoang, T. N., Low, B. K. H., and Kankanhalli, M. (2017). Information-based multi-fidelity bayesian optimization. In *NIPS Workshop on Bayesian Optimization*.

- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., and Da Fonseca, V. G. (2003). Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on evolutionary computation*, 7(2):117–132.