

Multi-Objective Optimizations in Rotor Aerodynamics using Variable Fidelity Simulations

Gunther Wilke

German Aerospace Center (DLR) Braunschweig, Institute of Aerodynamics and Flow Technology
Lilienthalplatz 7, 38108 Braunschweig, Germany, gunther.wilke@dlr.de

Abstract

This work explains the methodological setup of a variable fidelity framework for the aerodynamic optimization of helicopter rotor blades and demonstrates its capabilities for a single and multi-objective test case. The optimization approach utilizes a Hierarchical Kriging surrogate model. First, a benchmark for single objective optimization containing up to three fidelities demonstrates that the variable fidelity technique significantly reduces the required computational resources, if high fidelity optimizations are necessary. For the hover optimization savings up to 45.0% are achieved and in forward flight optimizations even 83.1%. Extending the framework to multi-objective problems, the usefulness is tested for mid-fidelity optimizations. The resulting Pareto fronts of the single and variable fidelity optimizations are compared to a reference solution. The accuracy of variable fidelity Pareto fronts is generally higher measured by the distance to the reference Pareto front. In the end, three different rotor designs are picked from the Pareto front and the advantages and disadvantages of each rotor are discussed.

1 INTRODUCTION

The aerodynamic optimization of helicopter rotors is demanding due to the complex nature of their flow field. On the one hand, the rotating system causes the downwash as well as the tip vortex to interact with the following blades, while on the other hand dynamic stall and shocks can occur in forward flight. At the same time, the structural analysis cannot be neglected. The deformation of the blade has a considerable impact on the performance, and thus fluid-structure coupled simulations are necessary for optimization as mentioned by Imiela [14].

There are multiple fidelities available for computing the aerodynamics of a rotor blade, with some of them sketched in Fig. 1. Low fidelity models comprise the blade element theory with inflow corrections methods like simple wake models. Mid fidelity ranges from more advanced wake models with panel method to simplified CFD models, such as the inviscid formulation of the flow equations. Solving the Navier-Stokes equations which model boundary layer effects are considered and high fidelity methods. Detached or large eddy simulations are even more accurate, but are still not applicable for industrial problems. While the fidelity increases, the resources required grow enormously.

Recent optimizations are either speed up using the adjoint methodology for computational fluid dynamics as done by Dumont et al. [1] and Choi et al. [21], or through surrogate and low fidelity models connected with global optimizers. Visingardi et al. [4] use a blade element method coupled with a wake model to di-

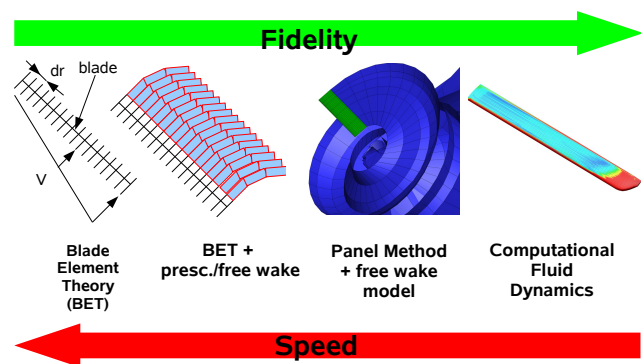


Figure 1: Fidelity range of Aerodynamic Models for Rotor blades

rectly optimize performance and acoustics of a variable speed rotor and deploy a genetic algorithm to locate viable configurations. Surrogate models in the form of artificial neural networks are adopted by Masaro et al. [3], who construct their surrogate model based on simplified aerodynamics. There are also high fidelity optimizations with computational fluid dynamics performed by Imiela [13], who exploits a Kriging model, or Johnson and Barakos [15] who train an artificial neural network.

The weighted function approach is often applied for multi-objective optimizations as for example done by Johnson and Barakos [15], who have a specific methodology on selecting the weights. The Nash game methodology is presented by Leon et al. [10] to obtain the Pareto front. The Nash game methodology gradually shifts from one goal function to the other.

However, there is also the methodology of searching directly for the Pareto front, as for example Visingardi et al. [4] for two mixed goal functions.

The research by Hollands et al. [18] reveals that for more advanced blade shapes the fidelity of simplified aerodynamics can be insufficient, and Wilke [23] presents a study on suitability of different aerodynamic modeling depth for planform optimizations. There is a concept followed by Collins [7], who takes advantage of different stages of fidelity in the optimization to achieve a speed up in the overall optimization process.

This paper focuses on the methodology of combining aerodynamic models of varying fidelity as an acceleration technique in surrogate based optimization for multiple objectives. Recent works by Han and Görtz [12] show a technique of data fusion they refer to as Hierarchical Kriging. A similar principle is investigated by Xiong et al. [24]. The idea is to build a surrogate model of the low fidelity aerodynamics, such as blade element theory coupled with prescribed wake models and calibrate this surrogate model with high fidelity CFD samples. The gain comes from using a lot of low fidelity samples to cover the global landscape of the optimization and only refine the model in the regions of interest with the high fidelity model.

The outline of this paper explains in detail the employed variable fidelity optimization framework from the general outline to the employed surrogates up the optimization techniques and goes over practical aerodynamic models for the planform optimization of rotor blades. The suitability of this methodology is demonstrated with two test cases. A single objective viscous high fidelity CFD optimization as well a multi-objective optimization for inviscid mid fidelity CFD using multiple planform parameters.

2 SURROGATE BASED OPTIMIZATION WITH VARIABLE FIDELITIES

This section explains the general outline of surrogate based optimization (SBO) and the additions to be made if variable fidelity methods are employed. The general aim of surrogate based optimization is to decrease the amount of required (high fidelity) simulations in an optimization process and substitute these with a surrogate model. The optimizer then seeks the point of interest in the surrogate model. Usually this point is recomputed with the simulation, and the result is fed back into the surrogate model to enhance it and re-perform the optimization. This is repeated until a stopping criterion is met. A good stopping criterion is the design confidence, which is explained in detail by Xiong et al. in [24]. The design confidence expresses the probability that the current best design cannot be improved by a pre-specified value anymore. For ex-

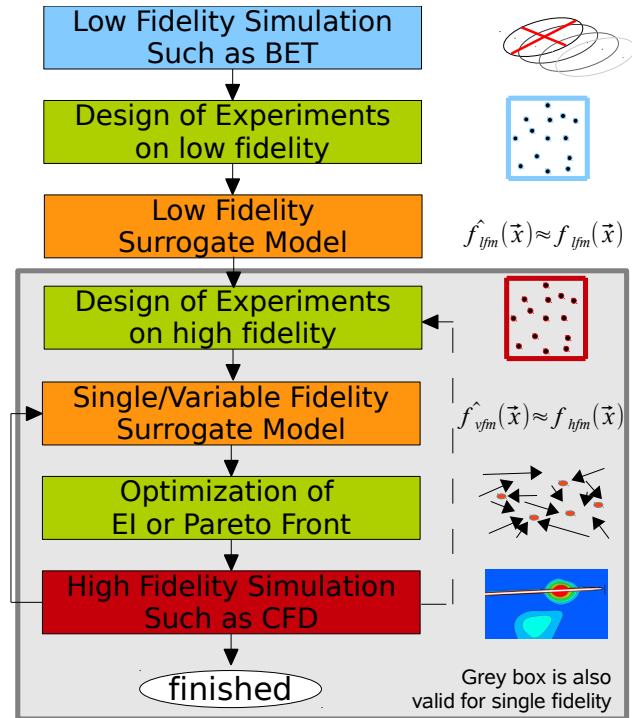


Figure 2: Variable Fidelity Optimization Process

ample if the a design confidence is above 90% twice in a row during the update cycle, the optimization is stopped.

The single fidelity surrogate based optimization is sketched in the gray box in Fig. 2. At the beginning of a surrogate based optimization, a design of experiments (DoE) is executed in which samples of the high fidelity simulation are spread out in the design space to gather initial information of the goal function. Afterwards the adaptive or update sampling is started. One well known approach given from the Efficient Global Optimizer (EGO) algorithm, as presented by Jones et al. [8] is to then seek for the expected improvement (EI) within the surrogate model, which yields a good trade-off between improving on the accuracy of the surrogate model as well as finding the best design. The point with the highest expected improvement is then sampled with the simulation.

The surrogate based optimization process is extended for variable fidelity optimizations by including a design of experiments for the low fidelity simulations, too. The low fidelity samples allow for the creation of a variable fidelity surrogate model in combination with the high fidelity samples. Looking at Fig. 2, the difference between single fidelity and variable fidelity optimizations is the top part above the gray box. The plan is to reduce the computational effort required to generate an accurate surrogate model. A demonstration is given in Fig. 3, where a 1D-slice through a 4D surrogate model is plotted. The three given lines are the visualizations of different surrogate models. The

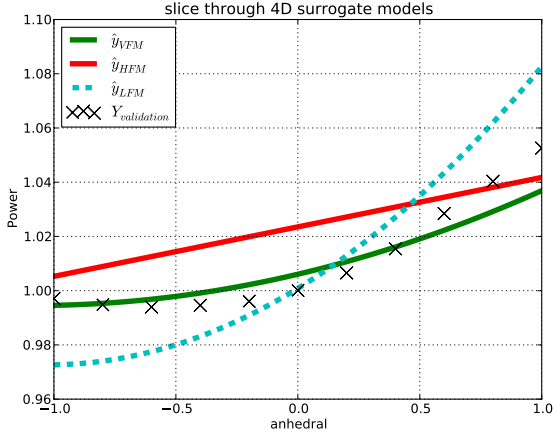


Figure 3: Demonstration of the improved accuracy of variable fidelity models.

dashed blue line corresponds to the surrogate model of the low fidelity model, here inviscid Euler computations, while the red line corresponds to the surrogate model of the high fidelity viscous Navier-Stokes solutions. The blue line was generated using 256 Euler samples, while the red line consists of only 8 Navier-Stokes samples. Combining both sample sets yields the green line, the variable fidelity surrogate model, which matches the validation samples with black Xs the best.

In the following, the technique behind the variable fidelity surrogate model is explained, as well as the optimization technique employed for this surrogate based optimization. Lastly the selected aerodynamic models are presented, before presenting the optimization results.

3 VARIABLE FIDELITY SURROGATES

The underlying variable fidelity method is referred to as Hierarchical Kriging as published by Han and Görtz [12]. It is similar to the technique presented by Xiong et al. [24] and is also related to Kriging with a drift. It differs from Co-Kriging as it does not directly correlate low- and high fidelity samples within the correlation matrix. To understand this method, the principal of universal Kriging is explained first, which is used for comparison between single- and variable fidelity surrogate based optimizations.

3.1 Regular Kriging

This section briefly covers the Kriging methodology. It shall help to understand the difference of single and variable fidelity Kriging. For a more detailed description of Kriging, the book by Forrester et al. [2] is recommended.

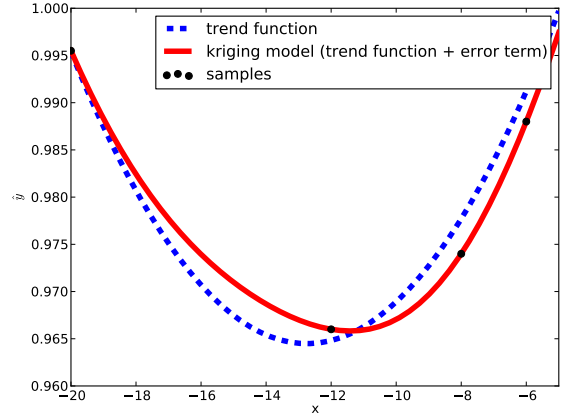


Figure 4: Trend Function and Kriging

ommended.

While ordinary Kriging assumes a constant trend referred to as the weighted mean of the sampled data responses \bar{y}_s , universal Kriging employs a trend function $f_{trend}(\vec{x})$ of higher polynomial degree based upon the design variables \vec{x} . Kriging then corrects the error between the trend function and the sample responses through the use of radial basis functions expressed by $\epsilon_{rbf}(\vec{x})$, and can be written as follows:

$$(1) \quad \hat{y}(\vec{x}) = f_{trend}(\vec{x}) + \epsilon_{rbf}(\vec{x})$$

A graphical interpretation is depicted in Fig. 4. The blue dashed curve is obtained from a polynomial trend function of degree two from the four samples in black. It regresses the data points as the system of equations is overdetermined for this type of polynomial function. The correction is represented by the red curve, which is the trend function $f_{trend}(\vec{x})$ plus the added radial basis function correction $\epsilon_{rbf}(\vec{x})$.

In Kriging the determination of the coefficients differs slightly from a regular regression model. While the ansatz function looks the same for a polynomial model of the degree of two:

$$(2) \quad f_{trend}(\vec{x}) = \beta_0 + \sum_k^d (\beta_k x_k + \sum_l^k (\beta_{(d+l)} x_k x_l)) = \vec{\beta} \cdot \vec{f}$$

the determination of these coefficients $\vec{\beta}$ is weighted with the inverse of the correlation matrix Ψ :

$$(3) \quad \vec{\beta} = (\mathbf{F}^T \Psi^{-1} \mathbf{F})^{-1} (\mathbf{F}^T \Psi^{-1} \vec{Y}_s)$$

with \mathbf{F} being the regression matrix composed of the regression vectors $\vec{f} = (1, x_1, \dots, x_d, x_1^2, \dots, x_d^2)$ for each data point \vec{x}_i . The correlation matrix is composed of the corrections of each sample with each other. The here applied basis function is the classical Kriging basis function, and an entry in the correlation matrix is

computed as follows:

$$(4) \quad RBF_{Kriging} = \exp\left(-\sum_k^d \theta_k |\vec{x}_i - \vec{x}_j|^{p_k}\right)$$

where θ_k and p_k are hyper parameters of the Kriging model, which allow it to be tuned for best prediction capabilities. To determine these hyper parameters, the concentrated maximum likelihood is maximized:

$$(5) \quad \ln(L) \approx -\frac{n}{2} \ln(\hat{\sigma}^2) - \frac{1}{2} \ln |\Psi|$$

with the model variance $\hat{\sigma}^2$ being:

$$(6) \quad \hat{\sigma}^2 = \frac{(\vec{Y}_s - \mathbf{F}\vec{\beta})^T \cdot (\Psi^{-1}(\vec{Y}_s - \mathbf{F}\vec{\beta}))}{n}$$

For a better condition of the correlation matrix as well as to deal with noisy functions, which can occur when flow solutions cannot be converged sufficiently, a noise constant λ can be added to the diagonal of the correlation matrix. This will regress the model slightly, yet not as much as a regression model. This procedure is valuable for deterministic computer experiments, where the error is of systematic nature and cannot be attributed to chance. A repetition of a sample will yield the same results again. This constant is then also tuned by adding it to the parameter list of the concentrated maximum likelihood optimization. A note on the tuning process is made in Section 4.2. The Kriging error function is then determined by:

$$(7) \quad \epsilon_{rbf}(\vec{x}) = \vec{\psi}(\vec{x}) \Psi^{-1}(\vec{y}_s - \mathbf{F}\vec{\beta}^T)$$

with $\vec{\psi}$ being the correlation vector build from all individual correlations of the predicted point with the points in the data set. The term $\Psi^{-1}(\vec{y}_s - \mathbf{F}\vec{\beta}^T)$ is stored as the weight vector \vec{W} for computational efficiency.

3.2 Hierarchical Kriging

The concept of Hierarchical Kriging is to exchange the polynomial trend function with a low fidelity surrogate model, which itself is based upon regular Kriging, or, in case of multiple fidelities used, upon another Hierarchical Kriging model. The proposal by Han and Görtz [12] is to multiply the low fidelity model \hat{y}_{lfm} with a model constant ρ which then yields the trend function:

$$(8) \quad f_{trend}(\vec{x}) = \rho \hat{y}_{lfm}(\vec{x})$$

while Xiong et al. [24] add an additional constant offset γ to the trend:

$$(9) \quad f_{trend}(\vec{x}) = \rho \hat{y}_{lfm}(\vec{x}) + \gamma$$

The author has his own variation and uses:

$$(10) \quad f_{trend}(\vec{x}) = \rho \hat{y}_{lfm}(\vec{x}) + \sum_k^d (\beta_k x_k)$$

which yields a mix of first order polynomial model and the low fidelity. The to be determined coefficients are handled similarly to universal Kriging, thus the extension to Hierarchical Kriging from an existing universal Kriging framework is made quickly. The regression vector is rewritten as $\vec{f} = (\hat{y}_{lfm}(\vec{x}), x_1, \dots, x_d)$. An advantage of Hierarchical Kriging, compared to Co-Kriging is that the problem is split into two smaller correlation matrices, instead of a big single one. The number of entries in correlation matrices for Hierarchical Kriging is $n_{lfm}^2 + n_{hfm}^2$, which grows slower than the size of a Co-Kriging correlation matrix $(n_{lfm} + n_{hfm})^2$. The statistical properties of expected improvement are then based upon the low fidelity trend and the high fidelity error estimation, which leads to an improved correction as high fidelity samples are further spread out than with Co-Kriging. Here the variable fidelity error estimation is also based upon the statistical error of the low fidelity error, which is zero at places where low fidelity simulations exists, thus the chance of placing a high fidelity sample there for correction is smaller.

4 OPTIMIZATION TECHNIQUES

4.1 Pareto Front Search

The optimization framework follows a similar idea as the EGO algorithm as presented in Jones et al. [8]. The surrogate model is sought for the point with the best expected improvement. A sample is then placed at this position and the response is fed back to the surrogate model. This technique is adapted with the major difference in the search strategy and the application of a variable fidelity surrogate.

The general outline of the multi-objective optimization technique is sketched in Fig. 5. First a design of experiments is spawned upon the **surrogate model**, which is used as the initial population of of an differential evolutionary algorithm. This algorithm is used to find the Pareto front. In the last step, a local search algorithm refines the front.

In detail, the design of experiments is based upon a central voronoi tessellated (CVT) hypercube as described in Ju et al. [17] with a Hammersley sequence [11] as the initial hypercube before the tessellation. This ensures a good space filling hypercube which allows to capture as much information as possible. The differential evolutionary algorithm is based upon the development by Das et al. [22] and is modified in such a way that individuals are replaced if they are superseded by their offspring in all goal functions. This yields the Pareto front after a sufficient amount of

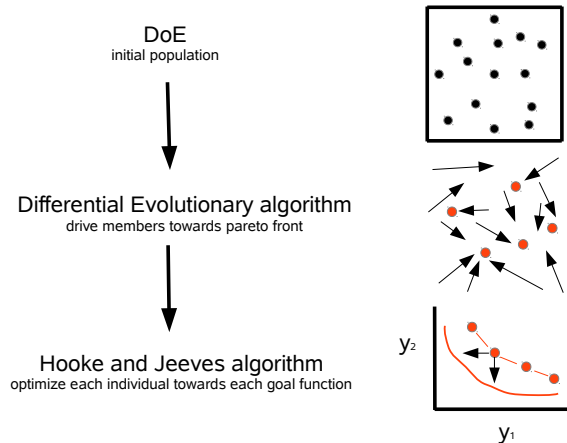


Figure 5: *The three stages of the optimization strategy*

generations. In the last step, each sample of the then given Pareto front is handed over as starting point for the simplex algorithm by Nelder and Mead [19] which is applied for each goal function. The history of points of the local searcher is kept, and later on the Pareto front is extracted from this history. The point with the highest statistical error from the Kriging model is sampled. This guarantees that samples do not bunch up in the same region and a certain diversity of the members on the Pareto front of the truth function is achieved.

4.2 Tuning of Surrogate Models

As mentioned in Section 3.1, the hyper parameters of the Kriging models can be tuned to increase the maximum likelihood function. The optimization process is similar to the one for finding the Pareto front, yet posed as multiple single objective problems. During the multi-objective optimization process, for each goal function an individual Kriging model is constructed and tuned.

5 AERODYNAMIC MODELS

A wide range of aerodynamic models for helicopter rotor blades is available. With increasing fidelity the computational effort of these methods grows rapidly. Simple aerodynamic models such as the blade element theory, or blade element theory enhanced with wake modeling can be computed in matter of minutes on current work stations, while panel methods with free wake models require hours to complete, and CFD methods can easily obtain run times of days or even weeks. Depending on the parametrization of the rotor blade, some methods may not be applicable. As the anhedral of the blade is strongly driven by the

blade-vortex interaction, a model neglecting this effect would be improper for (pure) anhedral optimizations of a blade. The selection of the different stages for the individual flight conditions are introduced in the following.

5.1 Low Fidelity Simulations

In this work, the aero-mechanic code HOST [6] computes the blade deformations and also takes over the role of trimming the rotor blade. For low fidelity hover simulations, the FISUW model based on an accelerated potential theory and developed by Basset et al. [20] improves the downwash computations of the blade element theory. The computational times are small and FISUW is robust, yet will not capture direct blade vortex interaction effects. For forward flight computations, the prescribed wake model METAR [5] corrects the induced velocities, which is not robust enough for hover computations, but grants good results in forward flight cases.

5.2 Mid Fidelity Simulations

For the mid-fidelity setup, Euler computations are performed using similarly coarse grids as Imiela [14], yet without resolving the boundary layer, which reduces the number of cells by a third, mesh size are printed in Table 1. Hover is assumed to be steady, while forward flight computations are modeled using a single blade approach. The flower solver is the DLR FLOWer code [16].

5.3 High Fidelity Simulations

The high fidelity aero-loads are also computed by FLOWer with which a complete helicopter [9] has already been simulated. In both cases, hover and forward flight, Navier-Stokes computations are performed, which sufficiently resolves the boundary layer as well as the tip vortex of the blade. In hover a periodic mesh is generated, which reduces the computational effort as only one blade is computed. In forward flight, the Chimera technique is applied to model the unsteady effects of all five blades. The Wilcox $k - \omega$ turbulence model is used. The grid setup is identical to the one by Imiela [14], with the mesh dimensions also listed in Table 1.

6 SINGLE OBJECTIVE OPTIMIZATION

This initial tests demonstrates the overall applicability of variable fidelity methods to the optimization process of helicopter rotor blades.

case	hover		forward flight	
	radial	chord	radial	chord
FISUW/METAR	31	-	25	-
Euler	25	61	32	96
	197,142 cells		380,016 cells	
Navier-Stokes	61	201	64	193
	1,376,256 cells		10,101,760 cells	

Table 1: Discretization of the blade and flow field for the individual solvers and flight conditions.

Parameter	baseline value	lower upper bound	
		hover/forward flight	
anhedral [$*c_{ref}$]	0.0	-1.0/-0.3	1.0/0.3
tip chord [$*c_{ref}$]	1.0	0.5	1.5
sweep [$*c_{ref}$]	0.0	-1.0/0.0	1.0
twist [$^{\circ}$]	-4.32	-20.0/-16.0	0.0/-4.0

Table 2: parameter ranges. reference chord length $c_{ref} = 0.14m$. hv=hover and ff=forward flight variables

6.1 Parameters and Flight Conditions

In order to have a reference to the multi-objective simulations, a single fidelity optimization is performed. This has been done for both flight conditions, hover (hv) and forward flight (ff). Both have been trimmed for a vertical force of $4400N$ at a tip Mach number of 0.646 . In forward flight the rotor had to counter a virtual fuselage drag of $530N$ at an advance ratio of $\mu = 0.38$. The considered parameters include anhedral, chord, sweep, and twist. Anhedral and sweep are specified through the offset at the blade tip, while chord refers to the tip chord length. All three parameters are varied parabolically, and start at 80.6% of the rotor radius. Twist on the other hand is varied linearly over the full radial span. The parametrization of the rotor is visualized in Fig. 6 and the ranges of the design parameters are listed in Table 2. The baseline values match the 7A description, a rectangular blade with a small blade twist. In forward flight the parameter boundaries are shifted towards a smaller design space, as for example a forward sweep or a strong an/- dihedral will cause severe aero-elastic problems leading to unachievable trim conditions. Hover in contrast allows a greater spectrum and thus the parameter space is enlarged in comparison.

The resulting goal function values are always scaled with results of the reference blade for each individual methodology and flight condition. Written out this means:

$$(11) \quad GF = \left(\frac{\text{Power current rotor}}{\text{Power 7A Rotor}} \right)$$




Figure 6: The parameter anhedral, chord, sweep and twist for the modification of a 7A rotor.

6.2 Optimization Setup

As the single objective optimization is part of the benchmark between single- and variable fidelity surrogate based optimizations, a set of various design of experiments is created. For the single fidelity (SF) optimizations three different design of experiments are created. The first sampling only includes one sample, which is the baseline rotor. The other two design of experiments are central voronoi tessellated (CVT) Hammersley sequences, one is created with 8 samples, the other one with 20.

For the variable fidelity (VF) optimizations, two dual-fidelity test are created and one triple fidelity case. The first dual fidelity cases also only consist of one sample, yet this sample is generated at the location of the optimum found in the mid-fidelity surrogate model. The mid-fidelity surrogate model is created by running a full factorial design of experiments with 256 inviscid Euler CFD computations. As the second design of experiments is for the dual-fidelity case, the central voronoi tessellated (CVT) Hammersley sequence with 8 samples from the single fidelity optimization is recycled.

The triple fidelity setup uses the surrogate models, which are the results of the optimizations in [23]. The design of experiments for the low fidelity model consists of a full factorial with 256 samples, while the mid fidelity Euler computations have 20 samples in the design of experiments and 40 in the update cycle. This yields a very accurate mid-variable-fidelity model around the mid fidelity optimum. The overall computational effort for the low- and mid fidelity samples is slightly less than two high fidelity computations in hover and forward flight for this case. The best rotor

case	low	mid	high	updates
	fidelity			
SF Point	-	-	1	15
SF 08 CVT	-	-	8	15
SF 20 CVT	-	-	20	15
VF Point	-	256	1	15
VF 08 CVT	-	256	8	15
VF Point 3F	256	20+40	1	15

Table 3: Number of samples used for each single objective test case.

from this mid-fidelity surrogate model is sampled as the initial sample for the high fidelity surrogate model.

Table 3 summarizes the sampling strategies and number of simulations performed for each optimization.

6.3 Results of Single- and Variable Fidelity Optimizations

Fig. 7 plots the convergence of the different optimization runs for the hover case. The dashed lines represent the single fidelity optimizations (SF), while the continuous lines are for the variable fidelity optimizations (VF). Iteration 0 is the first iteration where the surrogate model is sought for the point of the greatest EI. The iterations before the thick red line belong to the design of experiments and are thus random. Rotor designs that performed worse than the baseline rotor are not shown.

A greater design of experiments increases the chance of a sample being closer towards the optimum, which can be seen for the run of SF 20 CVT. Here a good point is spotted in the design of experiments early on, and thus the update cycle becomes more effective. In Table 4 the cost for reaching a hover performance better than 0.915 is listed, which is the highest best goal function value obtained from all runs. The cost is scaled with the average computational time for one high fidelity sample. For single fidelity, the difference between sampling a single point or 8 samples up front is similar in total cost, yet using 20 samples in the design of experiments already decreases the efficiency in terms of cpu time. However, a design of experiments of 20 samples can be parallalized if a sufficient number of processors is available and a theoretical wall clock time of 6 samples would be necessary if the load can be split up as such. For the variable fidelity optimization case, the number of required high fidelity samples is the lowest as expected. However, for dual fidelity case, the number of samples computed in the mid fidelity design of experiments is fairly costly. A shrinkage of design of experiments of the mid fidelity simulations or re-using the surrogate models of optimizations done in [23] might be beneficial. The triple fidelity sampling demonstrates this as less than fourth

case	low	mid	high	total cost
	fidelity			
SF Point	-	-	18	18.0
SF 08 CVT	-	-	18	18.0
SF 20 CVT	-	-	25	25.0
VF Point	-	8.0	7	15.0
VF 08 CVT	-	8.0	12	19.0
VF Point 3F	3.5e-3	1.9	8	9.9

Table 4: Cost for each hover optimization test case measured in theoretical number of high fidelity samples. The mean cost of a high fidelity sample is 160 cpu hours.

case	low	mid	high	total cost
	fidelity			
SF Point	-	-	16	16.0
SF 08 CVT	-	-	16	16.0
SF 20 CVT	-	-	29	29.0
VF Point	-	7.3	5	13.3
VF 08 CVT	-	7.3	10	17.3
VF Point 3F	1.6e-3	1.7	1	2.7

Table 5: Cost for each forward flight optimization test case measured in theoretical number of high fidelity samples. The mean cost of a high fidelity sample is 2600 cpu hours.

of the mid fidelity samples is necessary, while the cost for the lowest fidelity is negligible. As shown in [23], a mid-fidelity optimization with single-fidelity converges sufficiently with around 20 samples, 12 in the design of experiments and 8 in the update process. This would in theory be the cost of a less than one high fidelity sample. Under this assumption the application of the variable fidelity framework leads always to a reduction in total computational time in hover.

Looking at the examples in forward flight, the effectiveness of the variable fidelity technique increases. Fig. 8 shows that the triple fidelity setup already finds the best point as the first sample, which is attributed to the similarity of the mid- and high fidelity optimum as well as the good low fidelity trend, which was already observed in [23]. Similar to hover, the point sampling proves to be the most efficient technique for this optimization. Yet, for the 8 CVT run, the cost for variable fidelity is higher than the single fidelity, due to the expensive mid-fidelity sampling. Overall, the forward flight optimization requires less high fidelity update samples with the exception of the 20 CVT single fidelity run. Reason for this is that the large number of points in the initial stage already saturates the model too much that the best point has a fairly low expected improvement as the predicted error is too low. The costs for each individual optimization reaching a goal function smaller than 0.944 is given in Table 5.

Deducting from these optimizations the triple fidelity route proves to be viable and the most efficient way for

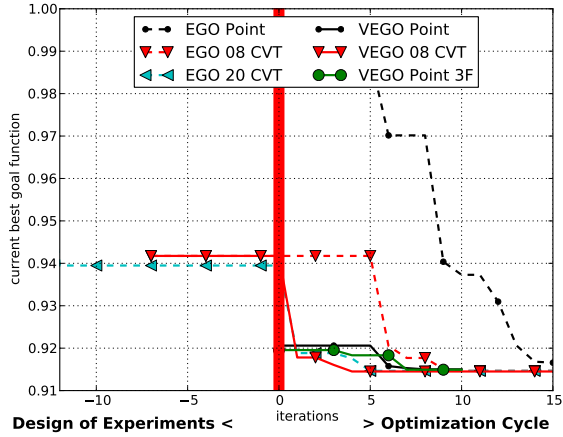


Figure 7: Comparison single (SF) and variable fidelity (VF) optimization in hover.

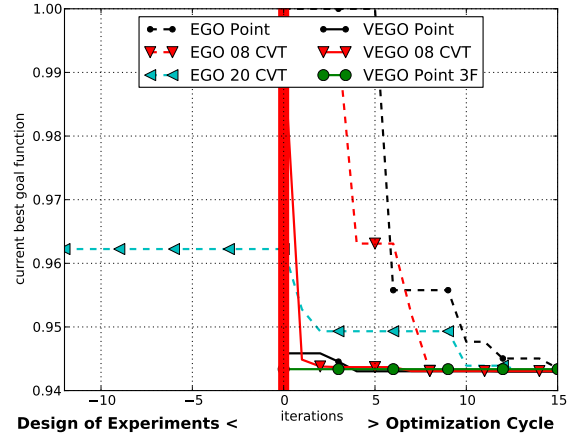


Figure 8: Comparison single (SF) and variable fidelity (VF) optimization in forward flight.

aerodynamic rotor blade optimizations. On the other hand, single point samplings in the initial design of experiments is a possibility to decrease the overall computational cost. However, if parallel computing is available, small random design of experiments minimize the number of update samples significantly, and depending on the resources available a larger design of experiments should be run.

7 MULTI OBJECTIVE OPTIMIZATION

This section covers the findings when applying the variable fidelity methodology towards multi-objective optimizations. The same planform parametrization is kept while the parameter boundaries for forward flight are chosen. A benchmark suite for mid-fidelity is performed, which allows the testing of different design of experiments setups. Later on, three different blade designs from the Pareto front are investigated.

7.1 Optimization Setup

Benchmarking single objective optimizations is easier in the sense that if the optimal configuration is known the current best objective function call can be compared against it and the strategy which reaches this point faster is the better. For multi-objective this is rather tricky as in the Pareto sense not one but multiple 'best' solutions exist. In order to determine the efficiency of single- and variable fidelity optimizations, the Pareto front is sought for a specific number of simulation calls after different initial design of experiments. The better optimization strategy should have a more refined Pareto front, which is compared to a reference solution. The reference solution is generated by a large sampling of the design space to create a

case	low fidelity		mid fidelity	
	DoE	updates	DoE	updates
SF 2 Points	-	-	2	40
SF 04 CVT	-	-	4	40
SF 12 CVT	-	-	12	40
SF 20 CVT	-	-	20	40
VF 2 Points	40	16	2	40
VF 4 CVT	40	16	4	40
VF 12 CVT	40	16	12	40
VF 20 CVT	40	16	20	40

Table 6: Number of samples used for each test case in the mid-fidelity multi-objective setup.

very accurate surrogate model. To minimize the computational effort, chord and sweep is set to 0.5 and 1.0, which are the optimal settings in both flight conditions for the mid-fidelity simulations. Thus, only the anhedral-twist design space is sampled uniformly with five samples in each direction leading to a full factorial cube of 25 samples. The reference Pareto front is then found by applying the optimization technique mentioned in Section 4.1.

The different sampling strategies are listed in Table 6 for the mid-fidelity optimization based on inviscid Euler CFD computations.

7.2 Results of Multi-Objective Optimizations

The Pareto fronts which are obtained after each optimization run are presented in Fig. 9 and Fig. 10 for single- and variable fidelity optimizations respectively. The scattered dots present the simulated samples, while the continues lines are the Pareto fronts which are found within the surrogate model after the final optimization run. The magenta line depicts the reference solution.

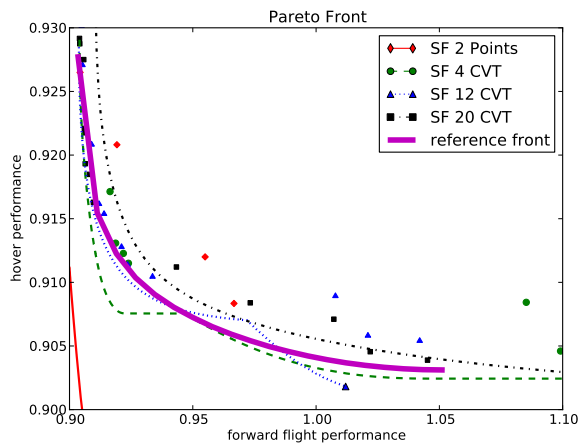


Figure 9: Single Fidelity Pareto front after 40 samples.

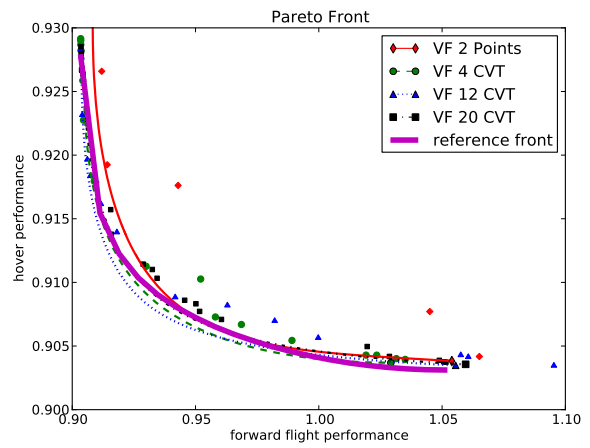


Figure 10: Variable Fidelity Pareto front after 40 samples.

case	forward flight		hover	
	design vector	performance	design vector	performance
reference	(0.17,0.5,1.0,-7.40)	0.903	(-0.3,0.5,1.0,-10.9)	0.903
SF 2 Points	(0.15,0.5,1.0,-7.20)	0.904	(0.01,0.5,0.0,-16.0)	0.905
SF 04 CVT	(0.13,0.5,1.0,-7.20)	0.904	(0.02,0.5,0.0,-14.1)	0.904
SF 12 CVT	(0.14,0.5,1.0,-7.31)	0.904	(-0.3,0.5,1.0,-11.9)	0.904
SF 20 CVT	(0.20,0.5,1.0,-7.71)	0.904	(-0.25,0.5,1.0,-11.2)	0.904
VF 2 Points	(0.16,0.5,1.0,-7.28)	0.904	(-0.3,0.5,1.0,-10.6)	0.904
VF 4 CVT	(0.16,0.5,1.0,-7.29)	0.903	(-0.25,0.5,1.0,-10.7)	0.904
VF 12 CVT	(0.19,0.5,1.0,-7.63)	0.904	(-0.3,0.5,1.0,-12.8)	0.903
VF 20 CVT	(0.16,0.5,1.0,-7.41)	0.904	(-0.3,0.5,1.0,-10.7)	0.904

Table 7: The Pareto optimal anchors of the simulated samples. The design vector consists of (anhedral, chord, sweep, twist).



Figure 11: Single Fidelity Pareto front.

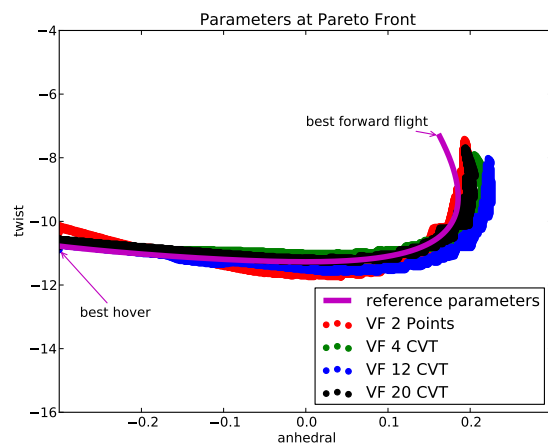


Figure 12: Variable Fidelity Pareto front.

Looking at the Pareto front of the single fidelity run in Fig. 9, it can be seen that for small design of experiments the samples are far away from the Pareto front, and the resulting surrogate model is still very inaccurate. In the case of the 2 Points sampling, the surrogate Pareto front under predicts the goal functions greatly, as the red line is far off to bottom right of the plot. The 4 CVT and 12 CVT optimizations gradually move towards the reference front. The 20 CVT sampling yields acceptable results, yet the hover goal function is over predicted. The samples are sparse in the region of better hover performance, the bottom right part on the plot leading to these inaccuracies. In comparison, all the surrogate Pareto fronts of the variable fidelity test cases are in close vicinity of the surrogate model with the simulated samples being closer to the Pareto front. However, the updates from the 2 Points sampling case of the variable fidelity optimization are also further away from the reference solution, and the same effect as for single fidelity is seen. A larger design of experiments improves the prediction accuracy of the surrogate model and as well will place more update samples on the Pareto front. In contrast to the single objective optimizations the saturation of the surrogate models is not seen. Reason for this is that the region of interest is greater than in a single objective optimization, thus the samples are further spread out and not centered around one optimum.

Another sanity check is the comparison of the anchor points of each Pareto front found. If these anchors match with the ones from reference front, it is clear that the multi-objective optimization also includes the optimal locations of individual objectives. In Table 7 the reference Pareto anchors along with the anchors of the different optimization runs are compared against each other. Except for the first two single fidelity optimizations, which indicate a different location of the hover optimum, all anchor points are found within an acceptable tolerance. The reason for the two single fidelity optimization predicting a different anchor point is that the hover optimum is shallow in the direction of sweep. On the one hand, this expresses the overall robustness of the process, while on the other hand the variable fidelity proves to be superior when using a reduced number of initial samples.

A different way of measuring the accuracy of the surrogates is to look at the spread of design variables on the Pareto front. Thus in Fig. 11 and Fig. 12 the variation of twist over anhedral is plotted, where it can be seen how one parameter is traded for the other. Hence the determination cannot be done analytically, the determination is done by using the numerical Pareto front search algorithm as mentioned in Section 4.1. This does not yield a smooth line, yet the trend of a hook like shape for the parameters on

the Pareto front is identified.

The parameters on the single fidelity front are more spread out in the anhedral-twist space. The more initial samples are used, the closer the agreement with parameters of the reference front. For the variable fidelity optimizations all samplings show the hook shape spread of the samples thus demonstrating that the accuracy of the resulting surrogate model is greater than that of the single fidelity model. Still, greater design of experiments lead to more points closer to the reference Pareto solution.

7.3 Rotors on the Pareto Front

With the Pareto front available the trade-offs of the different designs can be discussed. Three designs from the Pareto front are selected for a closer investigation. The two anchor points, namely the best forward flight and the best hover flight rotor, and a trade-off rotor are chosen, with the results taken from the optimization runs. The trade-off rotor is found in the middle of the Pareto front, where the change of hover performance is payed by almost the same change in forward flight performance. The design variables along with their respective performances are given in Table 8, and the rotors are visualized in Fig. 13, Fig. 14, and Fig. 15.

Looking at the difference in sectional torque between forward flight and trade-off rotor Fig. 16, the trade-off rotor is only little degraded in comparison with the hover rotor, where the difference to the forward flight rotor is plotted in Fig. 17. As the twist is almost similar for the best hover rotor to the twist of the trade-off rotor, the major losses are attributed to the strong anhedral. The great anhedral causes a loss of thrust on the outboard rotor stations and the collective pitch is increased in order to produce more thrust at the inner stations. At the same time, the pitching moment of the full rotor is grows as well as the roll moment decreased, which is reflected by change in cyclic control angles listed in Table 9. The rotational position of the tip is plotted in Fig. 18, which also leads to the indicator that more thrust needs be generated as the general angle of setting is greater. Even though the best hover rotor mostly saves power at the blade tip, the rise in drag to regain the thrust in the inner stations

case	anhedral	twist	forward flight performance	hover performance
Forward	0.170	-7.40	0.903	0.929
Trade-off	0.215	-9.81	0.910	0.915
Hover	-0.300	-10.9	0.947	0.903

Table 8: Parameters and performances for the individual rotors of the trade-off study.




Figure 13: Sketch of best forward flight rotor.




Figure 14: Sketch of trade-off rotor rotor.




Figure 15: Sketch of best hover rotor.

configuration	forward flight	trade-off	hover
collective pitch θ_0	17.0	18.4	20.7
cyclic cosine θ_c	1.2	1.1	-0.3
cyclic sine θ_s	-4.1	-4.2	-4.9

Table 9: Trim angles for the individual rotor configurations at forward flight.

is an unwanted side effect for forward flight.

Moving onto the hovering performances of the individual rotors, Fig. 19 depicts the gains and losses over the radial span. The forward flight rotor draws more power on the outside part than the hover optimized rotor. Due to the straight tip, the outboard station produces more lift and thus more induced drag at the tip. On the inboard part the rotor then requires less power as less lift has to be created in this region. The trade-off rotor is closer to the hover rotor, yet due the missing adjustment at the tip to alleviate the lift peak, the drag remains high in that region.

8 CONCLUSION

Based on findings in [23] the variable fidelity framework for helicopter rotors has been extended from low-mid fidelity optimization to low-mid-high fidelity optimizations as well as multi-objective problems. The initial single objective optimizations at high fidelity accuracy revealed that the triple fidelity approach is

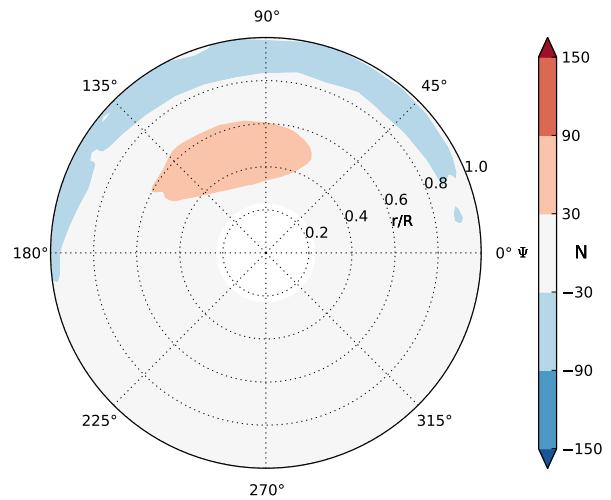


Figure 16: Torque difference of the trade-off rotor minus the best forward flight rotor in forward flight. Red zones indicate a greater power consumption, blue zones a power saving in comparison to the best forward flight rotor.

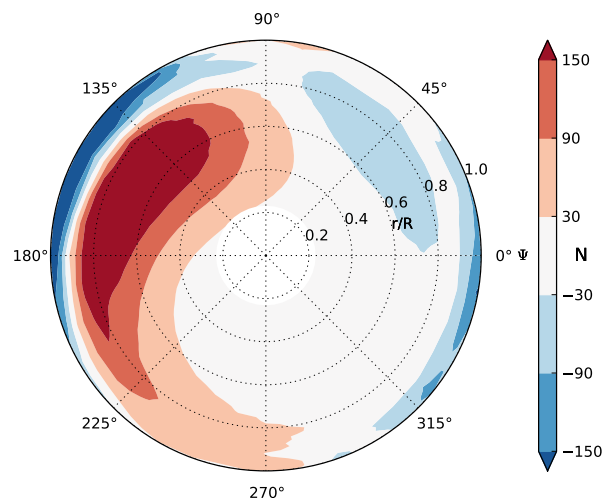


Figure 17: Torque difference of the best hover rotor minus the best forward flight rotor in forward flight. Red zones indicate a greater power consumption, blue zones a power saving in comparison to the best forward flight rotor.

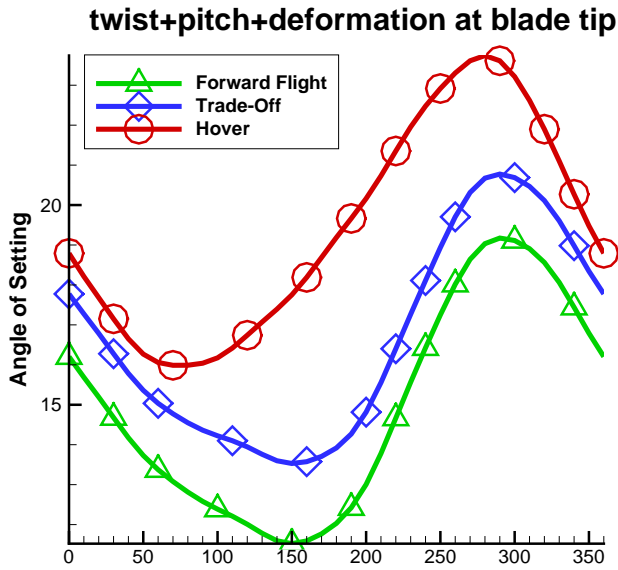


Figure 18: Blade tip angle of setting including pitch, twist and elastic deformation.

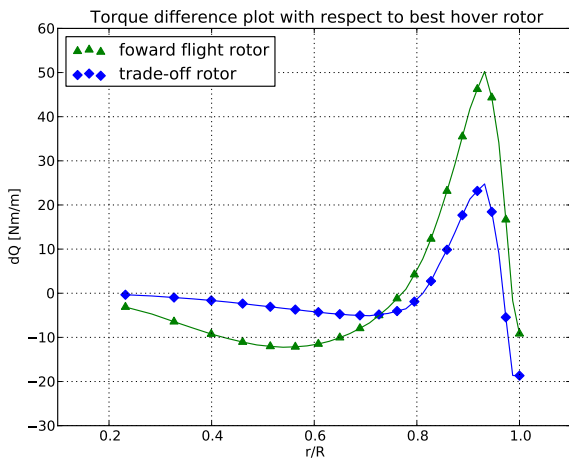


Figure 19: Torque difference of the best forward flight or trade-off rotor minus the best hover rotor. Positive areas require more power, negative areas mark an enhancement in comparison to the hover rotor.

well suited for rotor blade optimizations and saved up to 83.1% computational time in forward flight and 45.0% hover flight when comparing to the least costly single fidelity optimization.

The modification of the framework towards multi-objective capability proved to be very practical. For the single fidelity optimizations, only a few points could be placed close to the reference Pareto front and the according surrogate models required a great increase in samples to bring these sampled points closer to the reference front. For the variable fidelity surrogate model, all extracted Pareto fronts were in good agreement with the reference solution. Yet, for the smaller design of experiments, not too many update samples were placed towards the reference Pareto front. This effect improved with larger design of experiments. The overall gain of using the variable fidelity technique was clearly demonstrated.

Comparing three different rotors from the resulting Pareto front, the drawbacks and benefits of the different rotors are seen. With the sensitivities of the design parameters known from the surrogate models, a rotor suitable for a specific mission can be selected. In terms of aerodynamic performance a reduction in twist and a slight dihedral prove to be more advantageous in forward flight, while the opposite, high twist and anhedral, are helpful in hover flight, which is in agreement with the findings by Hollands et al. [18] for the an-/dihedral settings. A good trade-off has been identified, a rotor blade which featured a high twist with a slight dihedral. This blade has to struggle less with great elastic deformations in forward flight, while still maintaining a good hover performance.

Future work will include more parameters and constraints, while at the same time applying this framework to industrial scaled rotors.

9 Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ERF2013 proceedings or as individual offprints from the proceedings and for inclusion in a freely accessible web-based repository.

References

- [1] J. Peter A. Dumont, A. Le Pape and S. Huberson. Aerodynamic shape optimization of hover-

- ing rotors using a discrete adjoint of the reynolds-averaged navier-stokes equations. *Journal of American Helicopter Society*, 56:032002–1–11, 2011.
- [2] A. Søbester A. Forrester and A. Keane. *Engineering Design via Surrogate Modelling - A Practical Guide*. John Wiley & Sons Ltd., 2008.
- [3] A. D'Andrea A. Massaro and E. Benini. Multiobjective-multipoint rotor blade optimization in forward flight conditions using surrogate-assisted memetic algorithms. In *37th European Rotorcraft Forum*, 2011.
- [4] L. Federico A. Visingardi and M. Barbarino. Blade planform optimization for a dual speed rotor concept. In *38th European Rotorcraft Forum*, 2012.
- [5] G. Arnaud and P. Beaumier. Validation of r85/metar on the puma rae flight tests. In *18th European Rotorcraft Forum*, 1992.
- [6] K. Kampa W. von Grünhagen P.-M. Basset B. Benoit, A.-M. Dequin and B. Gimonet. Host, a general helicopter simulation tool for germany and france. In *American Helicopter Society 56th Annual Forum, Virginia Beach, Virginia, May 2-4, 2000*, 2000.
- [7] K. B. Collins. *A Multi-Fidelity Framework for Physics Based Rotor Blade Simulation and Optimization*. PhD thesis, Georgia Institute of Technology, 2008.
- [8] M. Schonlau D. R. Jones and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.
- [9] M. Dietz and O. Dieterich. Towards increased industrial application of rotor aeroelastic cfd. In *35th European Rotorcraft Forum 2009*, 2009.
- [10] J.-A. Desiderie D. Alfano-M. Costes E. R. Leon, A. Le Pape. Concurrent aerodynamic optimization of rotor blades using a nash game method. In *AHS 69th Annual Forum*, 2013.
- [11] J. M. Hammersley. Monte carlo methods for solving multivariate problems. *Annals of the New York Academy of Science*, 86:844–874, 1960.
- [12] Z.-H. Han and S. Görtz. A hierarchical kriging model for variable-fidelity surrogate modeling. *AIAA Journal*, 50-9:1885–1896, 2012.
- [13] M. Imiela. High-fidelity optimization framework for helicopter rotors. *Aerospace Science and Technology*, 23(1):2 – 16, 2012. 35th ERF: Progress in Rotorcraft Research.
- [14] M. Imiela. *Mehrpunktoptimierung eines Hub-schrauberrisors im Schweben- und Vorwärtsflug unter Berücksichtigung der Fluid-Struktur-Wechselwirkung*. PhD thesis, Institut für Aerodynamik und Strömungstechnik Braunschweig, 2012.
- [15] C. Johnson and G. N. Barakos. A framework for the optimization of a berp-like blade. In *38th European Rotorcraft Forum*, 2012.
- [16] Aumann P. Bartelheimer W. Bleecke H. Eisfeld B. Fassbender-J. Kuntz M. Lieser J. Monsen E. Heinrich R. Mauss M. Raddatz J. Reisch U. Roll B. Schwarz T. Kroll, N. Flower installation and user handbook, release 116. Technical report, German Aerospace Center (DLR), 2008.
- [17] Q. Du L. Ju and M. Gunzburger. Probabilistic methods for centroidal voronoi tessellations and their parallel implementations. *Parallel Computing*, 28:1477–1500, 2002.
- [18] E. Krämer M. Hollands, M. Kessler. Influence of an-/dihedral and of different blade-shapes on performance and aeroacoustics of an isolated rotor. In *38th European Rotorcraft Forum*, 2012.
- [19] J.A. Nelder and R. Mead. A simplex function for minimization. *Computer Journal*, 8-1:308–313, 1965.
- [20] J.V.R. Prasad P.-M. Basset, O. Heuze and M. Hamers. Finite state rotor induced flow model for interferences and ground effect. In *American Helicopter Society 57th Annual Forum*, 2001.
- [21] K. Lee G. Iaccarino S. Choi, M. Potsdam and J. J. Alonso. Helicopter rotor design using a time-spectral and adjoint-based method. In *12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference 10 - 12 September 2008, Victoria, British Columbia Canada*, 2008.
- [22] U. K. Chakraborty S. Das, A. Abraham and A. Konar. Differential evolution using a neighborhood-based mutation operator. *IEEE Transactions on Evolutionary Computation*, 13-3:526–, 2009.
- [23] G. Wilke. Variable fidelity optimization of required power of rotor blades: Investigation of aerodynamic models and their application. In *38th European Rotorcraft Forum*, 2012.
- [24] W. Chen Y. Xiong and K.-L. Tsui. A new variable-fidelity optimization framework based on model fusion and objective-oriented sequential sampling. *Journal of Mechanical Design*, 130(11):111401, 2008.