Approved for Public Release.

MULTIFIDELITY ROTOR OPTIMIZATION USING A MODERN COMPUTATION AEROMECHANICS TOOLCHAIN

Jeffrey D. Sinsay, U.S. Army Combat Capabilities Development Command Aviation & Missile Center, USA
Ananth Sridharan, Science & Technology Corp., Moffett Field, CA, USA

ABSTRACT

A multifidelity surrogate optimization procedure is demonstrated for aerodynamic design of a rotor. The rotorcraft comprehensive analysis CAMRADII is used as a lower fidelity, fast running solution. Higher fidelity solutions are calculated using the CFD solver HELIOS. Best practices in meshing and solution settings which balance overall accuracy with computational efficiency for the higher fidelity method are explored. The higher fidelity solution takes approximately 6 orders of magnitude more computational resources to complete. Through successive epochs of optimization and surrogate updating it is shown that the process can reach similar optimal values as optimization on a more densely sampled multifidelity surrogate, addressing the curse of dimensionality concern that often arises in optimization and lower the overall computational cost.

NOTATION

Symbols

c	chord design variable
R	rotor radius, m
r/R	non-dimensional rotor radius
x	vector of design variables
x_i	ith vector of design variables
x_{lb}	lower bounds on design variables
x_{ub}	upper bounds on design variables
x^*	optimal design
y	higher fidelity analysis objective value
$\bar{y}$	lower fidelity analysis objective value
$\hat{y}$	multifidelity surrogate objective value
$\hat{y}_{\text{hi}}$	higher fidelity surrogate objective value
$\hat{y}_{\text{lo}}$	lower fidelity surrogate objective value
θ	twist design variable
ρ	scaling factor between $\hat{y}_{\text{hi}}$ and $\hat{y}_{\text{lo}}$
$\sigma(\hat{y})$	variance of multifidelity surrogate

Acronyms

AvMC	Aviation & Missile Center
DE	Digital Engineering
CA	Comprehensive Analysis
CPU	Central Processing Unit
CFD	Computational Fluid Dynamics
CSD	Computational Structural Dynamics
DDES	Delayed Detached Eddy Simulation
DEVCOM	U.S. Army Development Command
DLR	German Aerospace Center
DSRC	Defense Supercomputing Resource Center
GPU	Graphical Processing Unit
GPR	Gaussian Process Regression
HART–II	2nd HHC Aeroacoustic Rotor Test
HHC	Higher Harmonic Control
InDeWo	International Design Workshop on Rotor Blade Optimization
JAXA	Japanese Aerospace Exploration Agency
ONERA	The French Aerospace Lab
RANS	Reynold's Averaged Navier-Stokes
YAML	Yet Another Mark-up Language

INTRODUCTION

The U.S. Army Combat Capabilities Development Command Aviation & Missile Center (DEVCOM AvMC) continues to maintain an interest in enhancing its ability to perform Digital Engineering (DE). A recent Army Directive (Ref. 1) on DE further emphasized the importance of moving from a design-build-test approach to a model-analyze-build approach in system development and acquisition. Advances in computational power and software algorithms in recent decades have directly lead to an improved ability to accurately predict performance, loads and flight dynamics of rotorcraft using physics-based modeling and simulation (Refs. 2–4). In the area of rotorcraft aeromechanics analyses, recent results (e.g. (Refs. 5, 6)), have demonstrated that simulation models can predict test results with accuracy sufficient to enable the move from a build-test to model-analyze paradigm.

While the prediction accuracy of computation analysis can approach physical testing, the computational cost of obtaining these results is still a major hindrance to realizing the vision of DE. Recent work on GPU based rotorcraft computational fluid dynamics (CFD) (Refs. 7, 8) offers hope of a significant speed up in computational times compared to legacy CPU based approaches. However, when used in designing a virtual prototype, as envisioned in DE, this computational speed-up may be insufficient to offset the increased cost in total computational resources needed to support the inherent iterations required of a design process. All design/optimization methods require some form of iterative information gathering where a design space is sampled at multiple points, information gathered and used to update the design and then sampled again based on that update, repeating until some terminating criteria is met. These methods, whether they be ad hoc human in-the-loop approaches, formal gradient based optimization searches, or gradient free search algorithms, all typically require orders of magnitude more simulations than is necessary for the simple fixed design analysis problem.

A variety of multifidelity approaches to rotor blade optimization have been reported on previously. Leusink *et al.* (Ref. 9) used a lower-fidelity comprehensive analysis (CA) to search the full design space. This was then followed by sampling a subset of the design vector space near the initially identified optimal design using higher fidelity analysis to built a Gaussian process regression surrogate upon which a 2nd round of optimization of was preformed. Such an approach demonstrated an ability to identify an optimal design using fewer higher fidelity functional evaluations than would have been required for an approach that started with building a higher fidelity surrogate directly on the full design vector space. Collins and Sankar, (Ref. 10) and Wilke (Ref. 11) have demonstrated the use of multifidelity surrogates for rotor optimization, reducing the number of higher fidelity simulation evaluations required. These multifidelity approaches compare favorably to those which operate on surrogates built solely from higher fidelity simulations, or attempt to optimize on the higher fidelity simulation model itself. In multifidelity surrogate modeling, results from a lower fidelity, computationally cheaper, model are corrected with a small sampling of the higher fidelity analysis. The primary concern in multifidelity surrogate modeling then becomes establishing a heuristic for choosing how and when to sample the higher fidelity analysis to balance between accuracy and cost.

In this paper, use of the rotorcraft aeromechanics CA tool, CAMRADII (Ref. 12), as a lower fidelity analysis combined with the CREATE™-AV HELIOS (Ref. 13) CFD framework for higher fidelity analysis in a multifidelity approach is explored. Such an approach is beneficial because lower fidelity CA models has been shown to capture many of the relevant aeromechanical phenomena of rotors, while enjoying execution times of around 0.25 CPU-hr per analysis point. This compares very favorably to the 10,000s of CPU-hr typically required to calculate an analysis point in CFD. DE virtual prototyping becomes a tractable problem if an analysis approach can be developed where the computational cost is similar to comprehensive analysis, but with accuracy approaching that of CFD.

Table 1: 2nd InDeWo Rotor Blade Optimization Design Conditions

	Hover	Cruise
Pressure Altitude	30 m	
Air Temperature	17.3 °C	
Lift	5,582 N	
Drag Area	0 sq-m	0.2 sq-m
Roll & Pitching Moment	0 Nm	
Horizontal Speed	0 m/s	66 m/s
Rotor Speed	1042 RPM	
Thrust Weighted Chord	0.121 m	
Solidity (TW)	0.077	

As a test case for exploring the multifidelity optimization approach proposed here, optimization is performed on a common hover case proposed as part of the International Design Workshop (InDeWo) on Rotor Blade Optimization (Ref. 14). The case is based on optimizing the HART–II rotor for a hover and high speed forward flight condition, and extends work performed by DLR, JAXA and ONERA in a tri-lateral research activity (Refs. 15, 16). Their approaches are variations of using blade element theory with fast prescribed wake or dynamic inflow model for lower fidelity predictions that are then combined with RANS CFD based higher fidelity predictions using various surrogate modeling techniques. Focusing on a common optimization problem allows for easier comparison of results and methodologies between the participants, helping to indentify strengths and weaknesses of a variety of approaches. Details of the multiobjective optimization problem are given in Table 1. Hong *et al.* (Ref. 17) reported related work for the hovering condition, including airfoil shape optimization using a neural net airfoil surrogate coupled with CAMRADII. Unlike others in the workshop, Hong's research tried to incorporate higher fidelity information by adjusting free wake modeling parameters in CAMRADII using wake geometry information extracted from 3D CFD. The optimum found using the tuned CA is then assessed in CFD to establish expected performance.

The present paper explores the use of the CAMRADII free wake as the lower fidelity analysis, similar to Hong, but coupled with higher fidelity HELIOS CFD analysis using a multifidelity surrogate model similar to the approaches of DLR, ONERA and JAXA. The current approach successively samples the higher fidelity analysis in a search for an optimum design proceeds. Thanks to the availability of U.S. DoD supercomputing resources, it is possible to sample higher fidelity cases with finer mesh resolution than done by other InDeWo participants.

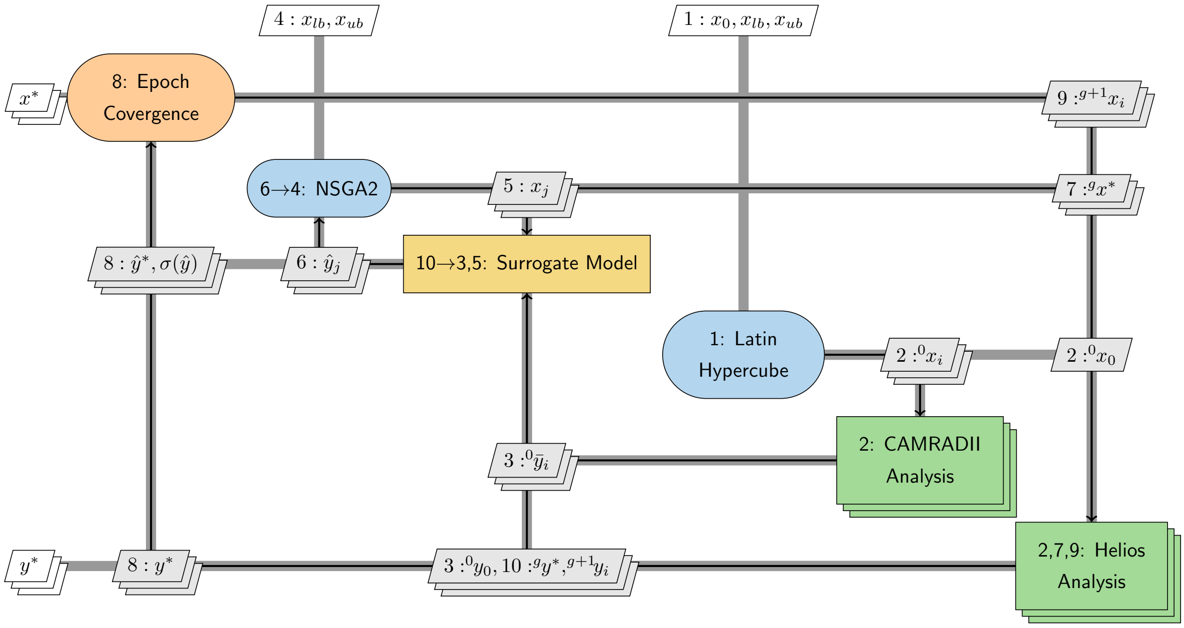

Fig. 1: Overall multifidelity optimization process

MULTIFIDELITY TOOLCHAIN OVERVIEW

The overall design process is outlined in Figure 1. The process starts with sampling the design space using the lower fidelity CAMRADII model on 100s of design points using a Latin hypercube sampling strategy. From this initial sample population, Gaussian process regression (GPR) models are constructed for each of the objective and constrain functions of interest using the CAMRADII results, along withHELIOS results for the initial design to initialize the multifidelity surrogate model. An initial optimization epoch is then performed on the resulting surrogate using the NSGA2 optimizer (Ref. 18). The optimization is allowed to continue for up to 40 generations, using a population size that is at least 5x the number of design variables. At this point in the process, additional higher fidelity information is introduced by sampling HELIOS using design points from the Pareto frontier anchor points. Anchor points are those design points which are individually most optimal in a particular objective function, the total number of anchor points therefore corresponds to the total number of objectives in the multiobjective problem. If the multifidelity surrogate model is not yet converged, additional points, which do not lie on the Pareto frontier, are chosen in the design space and sampled with the higher fidelity analysis as well. These additional points are needed to prevent the algorithm from only searching in the local neighborhood of the initial surrogate optimal values. This diversity in the search accounts for the potential that the initial multifidelity surrogate is not of sufficient accuracy to correctly identify all potential optima in the higher fidelity design space. Optimizing on such a surrogate may falsely converge to a globally sub-optimal region, unless additional information is available to improve the multifidelity surrogate in areas away from that region. The results from HELIOS are then used to update the surrogates for each objective and constraint function.

These updated surrogates are the basis for optimization in the proceeding epoch, which again consists of performing an NSGA2 search on the update surrogate, followed by high fidelity sampling based on the Pareto frontier anchor points. The process proceeds through successive epochs until the following convergence criteria are met:

1. the change in location of the Pareto anchor points, as measured by the L2-norm of the distance between the successive anchor points is less than a specified tolerance, and

2. the maximal change in a set of surrogate values for a set of well spaced control points, chosen using an initial Latin hypercube sampling, between successive epochs is less than a specified tolerance (indicating global convergence of the surrogate model).

Additionally, a limit on computational resources and time may terminate the algorithm and require human-in-the-loop determination if additional computations are required, or if the solution found is of sufficient quality to meet the engineering need.

The intent of this approach is to judiciously limit the number of higher fidelity solutions required, recognizing that for the purpose of design, accuracy of the surrogate is only of primary importance near the ultimate optimal design point. It does not strictly matter if a point away from the optimal design point is 1%, 10% or 100% worse than the optimal design. Practical rotor design problems can often be nonconvex; some exploration using high fidelity away from the initial optimal points is desirable to avoid premature settling of the optimization algorithm into a local optima. Furthermore, the use of a free wake CA model is explicitly chosen as the lower fidelity model because it has been shown to reasonably capture trends and in some instances approach the same level of accuracy as higher fidelity CFD models with much lower computational cost.

The algorithm described above, notably has very heterogeneous computational requirements for its various elements. The multifidelity optimization problem has been implemented on the U.S. Defense Supercomputing Resource Center (DSRC) using our previously developed Tool for Optimization of Rotorcraft Concepts (TORC) (Ref. 19). TORC is a lightweight Python framework which allows users to rapidly assemble analysis workflows using either Python native (e.g. HELIOS) or Python wrapped legacy analysis tools (e.g. CAMRADII). Users specify the workflow in a computer agnostic manner via a YAML input file, and TORC manages the file system, analysis execution, and data collection processes to complete the specified workflow. For the present work, the HELIOS wrapper has been updated to include Python utility functions which automate the generation a rotor blade strand meshes from an arbitrary input of the distribution of chord, twist, sweep, anhedral, and airfoil cross-sections using the BLADEGEN tool (Ref. 20), along with generation of the associated MELODI input files need for force calculations and CFD/CSD coupling.

Helios Modeling

The HPCMP CREATE™program and DEVCOM AvMC have collaboratively developed the high-fidelity simulation framework HELIOS for rotorcraft analysis. The HELIOS framework provides the ability to select from a number of different modules to perform overset multi-mesh CFD analysis. The ability to select analysis modules, and tailor the size and refinement of the temporal and spatial solution mesh, presents opportunity for making important marginal reductions in computational cost per solution point. These marginal savings can result in large improvements in performance for an optimization problem where many 10s of CFD solutions may be required to reach optimality. In the off-body region, ORCHARD (Ref. 7) is utilized to solve the Reynolds-averaged Navier Stokes (RANS) equations on a fifth-order spatially accurate Cartesian mesh. The RANS equations were solved in the near-body region using the MSTRAND solver (Ref. 20). This solver is ideal for the current application because, when coupled with the BLADEGEN surface mesh generation tool, the entire mesh generation process is automated, starting from a parametric description of the geometry.

BLADEGEN generates a structured surface mesh along the span of the blade, terminated on each end with unstructured end planes or rounded caps. As seen in Figure 2, the surface mesh has clustering of points chordwise towards the leading and trailing edges, as well as spanwise with clustering towards the tip and root. The exact location and minimum mesh spacing is specified as part of the input, along with airfoil 2D cross

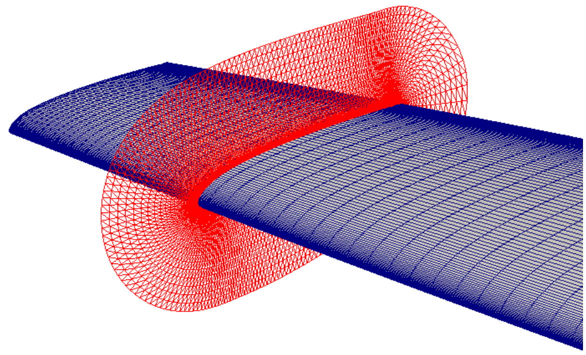

Fig. 2: BLADEGEN generated surface mesh and cross section of near-body mStrand mesh

section point information and spanwise variation of chord, twist, sweep and anhedral. In this study the airfoil has a blunt trailing edge, with 5 points placed along the trailing edge, resulting in an O-like strand mesh topology around the airfoil cross-section. The number of points in both the spanwise and chordwise wrap-around directions were varied, as shown in Table 2, to help establish settings for HELIOS to use when optimizing. Compared with previous work by other participants in InDeWo (Ref. 15), the near-body meshes used in the current work are at least twice as fine.

For the purpose of optimization, capturing the correct trends and location of optimality with respect to the design parameter values is the primary concern, and not the absolute magnitude of the objective functions themselves at a given combination of design parameter values. Once an optimal design has been identified, post-optimization runs in higher fidelity using alternative solution settings and more computational resource may be acceptable to reduce the uncertainty in the predicted magnitude of objective functions for the previously identified optimal design parameters.

Changes in azimuthal time step and near-body and off-body mesh size were explored on the baseline HART–II rotor case at a nominal hover collective of $10°$. A summary of the various modeling permutations with resulting computational cost is given in Table 2, while the predicted thrust and power is plotted in Figure 3 along with results from CAMRADII free wake and wind tunnel test data at identical atmospheric conditions (Ref. 21). Convergence histories for the 6 cases are shown in Figure 4, all cases were run as time accurate, unsteady in HELIOS. The cases which used a time step corresponding to an azimuthal advance of $0.25°$ per step show good convergence of figure of merit after 12 to 14 rotor revolutions, while the coarser $0.5°$ per step cases need upwards of 20 rotor revolutions to reach a periodic state. In particular, case 6 with a coarser time step and coarsest off-body mesh size requires not only longer to settle to a steady mean figure of merit value, but also shows a significantly higher oscillatory variation at frequencies higher than 1/rev.

Case 1 is slightly coarser in both near-body and off-body meshes than the finest meshes run previously for HART–II

Table 2: Summary of mesh and solution parameters examined for HART–II 10° collective condition

Case	1	2	3	4	5	6
Number of points around airfoil	494	202	494	202	202	202
Number of points spanwise along blade	177	121	177	121	121	121
Tip spanwise spacing (%R)	0.0086	0.05	0.0086	0.05	0.05	0.05
Near-body grid size per blade ($\times 10^6$)	8.2	2.0	8.2	2.0	2.0	2.0
Finest off-body cell size (%c)	6.5	6.5	6.5	6.5	6.5	12.9
Off-body grid size ($\times 10^6$)	123.1	123.1	123.1	123.1	184.4	133.3
Vertical extent finest refinement region	1.25R	1.25R	1.25R	1.25R	2R	2R
Turbulence Model Form	DDES	DDES	RANS	RANS	RANS	RANS
$\Delta\psi$ (deg)	0.25	0.25	0.25	0.25	0.5	0.5
CPU hours per rev	9,613	4,430	12,734	4,135	4,332	1,409

validation (Ref. 22). It is the finest mesh run for the present study, since, for optimization usage, the solution cost of an individual case needs to be less than the cost of a individual validation case. As mentioned above, some loss in absolute predictive accuracy is acceptable when searching for the optimal design shape. The fact that case 1 agrees with free wake to within about 3 counts of figure of merit, suggests that at least at the nominal HART–II design point, the difference in accuracy between free wake and CFD is relatively small. Coarsening the blade surface mesh in case 2 results in a 7 count reduction in figure of merit, but at least a 2x reduction in computational cost. Comparing cases 1 & 2 with cases 3 & 4 suggests that there is no significant difference in the computational cost of DDES compared to RANS. Comparing case 6 in Figure 3 with the other cases shows that coarsening the finest off-body grid cell to 10% chord results in a significant change in thrust. This is likely due to artificial dissipation of the tip vortex, which can be seen in the visualization of Figure 5c compared to the two cases which retain distinct tip vortices.

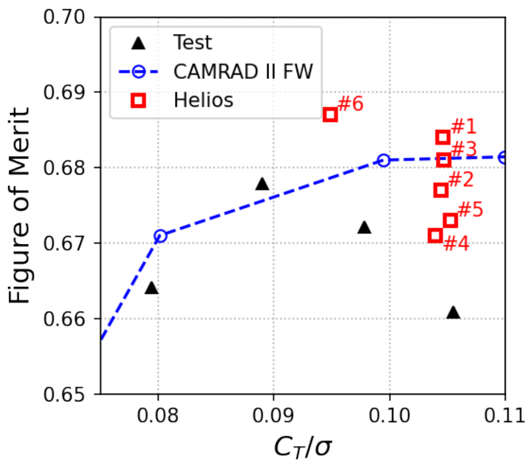

Fig. 3: Figure of merit and blade loading for 6 different Helios solution and mesh settings compared with CAMRAD II free wake and wind tunnel test results.

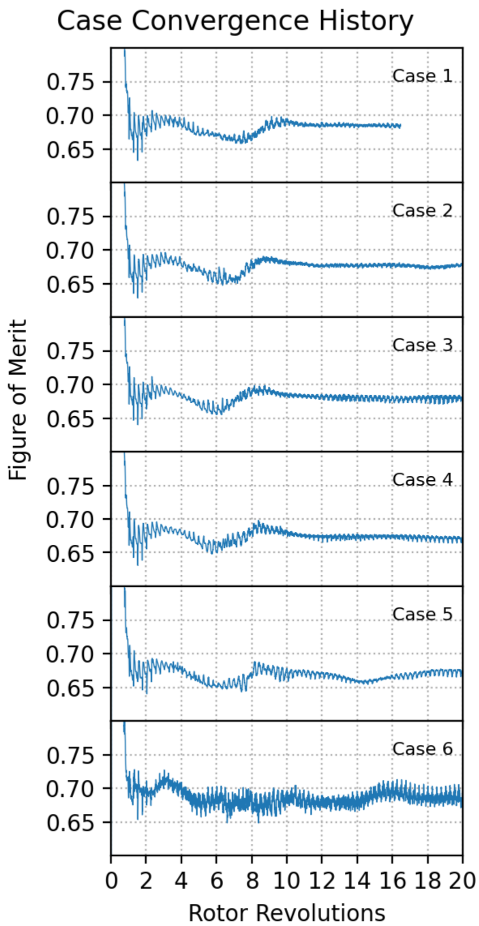

Fig. 4: Convergence history of baseline HART–II rotor Helios cases.

Based on these results, the settings of case 2 are used for the optimization process, achieving a balance between computational cost and solution predictive accuracy. Further improvements in computational cost are desirable. With these solution settings and the 12-14 rotor revolutions required to achieve a steady figure of merit solution, each solution is expected to require 60,000 CPU hours per solution. The emerging HELIOS capability (Ref. 13) of a complete suite of GPU based solvers offers potential to reduce this cost by roughly a factor of 10. Another significant factor in the cost of each HELIOS solution is the need to compute upwards of 15 revolutions of the rotor to achieve trim and establish the rotor wake and associated global flow field when starting from a quiescent initial condition. In particular, the ability to restart slightly modified geometries from a previous case of similar (but not identical geometry), may offer a way to reduce the number of revolutions needed to reach a steady state by removing the initial transients as seen in Figure 4. Additionally running for a coarser initial time step to develop the general flow field, followed by a few revolutions of a finer time step has been shown to be helpful with other HELIOS solvers.

Surrogate Modeling

Gaussian process regression (GPR) surrogate models were chosen for the present design optimization study. GPR models build an estimate of functional values based on known values at a set of training data points and an assumed probability distribution from a family of Gaussian kernel functions at locations away from the training data. The parameters that define the kernel functions are selected such that the resulting surrogate is a maximum likelihood predictor of the observed functional values. An advantage of GPR is that it provides not only estimates for functional values, but also an estimate of uncertainty in that value. This uncertainty measure can be searched for location of maximum uncertainty in the design space to inform infill sampling.

For multifidelity optimization a number of modifications to the basic GPR approach have proven popular in literature. Variations on Kriging modeling have shown great popularity, where the Kriging approach extends GPR to include individual variance parameters for each dimension of the regression space. Wilke (Ref. 23) has made extensive use of Hierarchical Kriging methods for surrogate rotor optimization. In Hierarchical Kriging lower fidelity samples are used to build a Kriging model for the general trend in the design space. This Kriging model then serves as the trend function in a second Kriging model which takes a limited set of higher fidelity analysis samples to generate another response surface which represents the high fidelity function estimator. Such an approach is assumed useful when the lower fidelity analysis generally captures the trends in the design space, and it is desirable to just use higher fidelity analysis to sample and correct in regions of interest.

Co-Kriging approaches have been used by other InDeWo participants (Refs. 16, 17). In the co-Kriging approach, low fidelity and high fidelity samples are combined into a single Kriging model using a method similar to that outlined by Forrester *et al.* (Ref. 24). In this method, the higher fidelity surrogate model is assumed to be a function of a scaling factor applied to the lower fidelity model plus a second Gaussian process model that captures any residual difference between the scaled lower fidelity surrogate and the higher fidelity sample points.

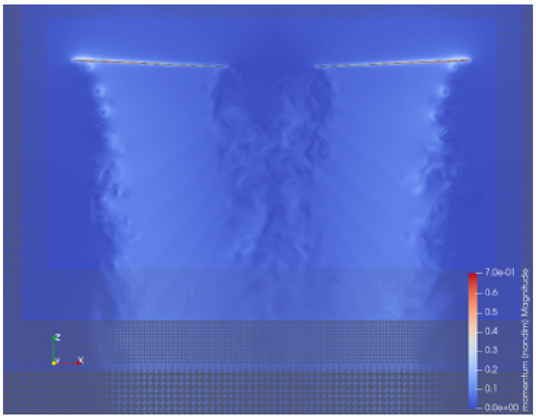

(a) Case 1

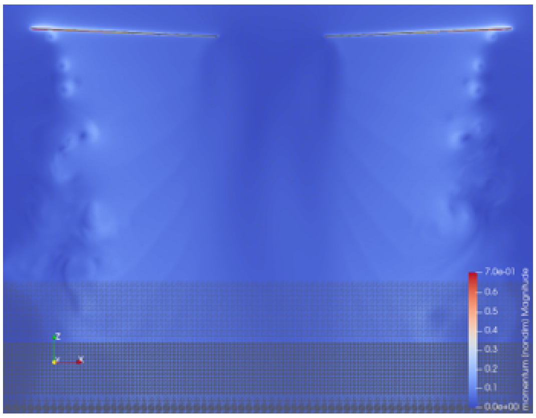

(b) Case 5

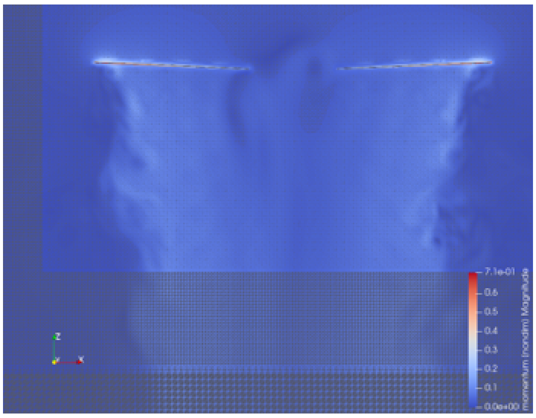

(c) Case 6

Fig. 5: Qualitative comparison of the vorticity magnitude for three HART–II solution and mesh settings, not the loss of well defined tip vortices in Case 6.

$$\hat{y}_{\text{hi}} = \rho \hat{y}_{\text{lo}} + \hat{y}_{err}, \quad (1)$$

where $\hat{y}_{err}$ is a Gaussian process describing the bias error between the high fidelity and low fidelity results. It should be noted that this approach includes an explicit assumption that GPR of the low fidelity ($\hat{y}_{lo}$), is statistically independent of the error basis.

Walizer *et al.* (Ref. 25), recently, successfully applied a co-Kriging method for prediction of HELIOS hover performance data using a lower fidelity comprehensive analysis and a lim-

Table 3: Summary of Design Variables for 1st and 2nd International Design Workshop on Rotor Blade Optimization

Design Variable	Initial Value	Upper Bound	Lower Bound
c_1	1.0	1.5	1.0
c_2	1.0	1.5	0.5
θ_1	-8.0	0.0	-10.0
θ_2	0.0	5.0	-5.0

ited set of higher fidelity points. Their approach leveraged an existing Python framework, EMUKIT (Ref. 26), which provides easy access to multifidelity modeling using GPRs as implemented in the related Python package GPY (Ref. 27). The linear multifidelity model of Eq. 1, as implemented in EMUKIT, has been incorporated in TORC, and used in the present optimization studies.

CHORD OPTIMIZATION IN HOVER

The InDeWo 2023 meeting looked initially at optimization of an isolated rotor in hover (Ref. 17), using just the hover conditions listed in Table 1. The baseline design is based on the HART–II rotor and optimization was performed on four design variables as shown in Figure 6 and summarized in Table 3.

Two variables: θ_1, θ_2, describe a non-linear twist law, where θ_1 is the linear twist rate, which is maintained to r/R=0.75, and θ_2 is a $\Delta\theta$ at the tip (r/R=1.0). The actual twist law from r/R=0.75 to 1.0 is defined by a cubic spline whose slope at r/R=0.75 matches the value of θ_1 and the value of twist at r/R=1.0 corresponds to $\theta(r/R = 1.0) = 0.25\theta_1 + \theta_2$, where the overall twist law has adjusted to the usual rotorcraft convention of $\theta(r/R = 0.75) = 0.0$. Two additional variables, c_1, c_2 are used to describe the chord variation along the span. c_1 is the normalized chord at r/R=0.75 and c_2 the normalized chord at r/R=1.0. The actual chord law is a constant value of 1.0 from root (r/R=0.22) to r/R=0.5.

A cubic spline then smoothly varies from the chord at r/R=0.5 to the value of c_1 at r/R=0.75 and finally a second cubic spline varies smoothly from c_1 to c_2 at the tip, where the inboard portion of the 2nd spline is constrained to match the chord slope at r/R=0.75. This chord law is then multiplied by a reference chord c_{ref} such that the thrust weighted solidity of the rotor is held constant at 0.121. The constant on thrust weighted solidity is automatically handled in the blade geometry generation process, and is not an active constraint in the rotor optimization problem itself.

The initial optimization search was performed using the NSGA2 genetic algorithm. Because of the relatively low computational cost of CAMRADII, and the ability to parallelize the search process in each generation of NSGA2, the initial search was performed directly on the CA without the use of a surrogate. This optimization was run with a population of 128 members for 30 generations. The resulting optimal twist and chord law results are shown in Figure 7 and compared

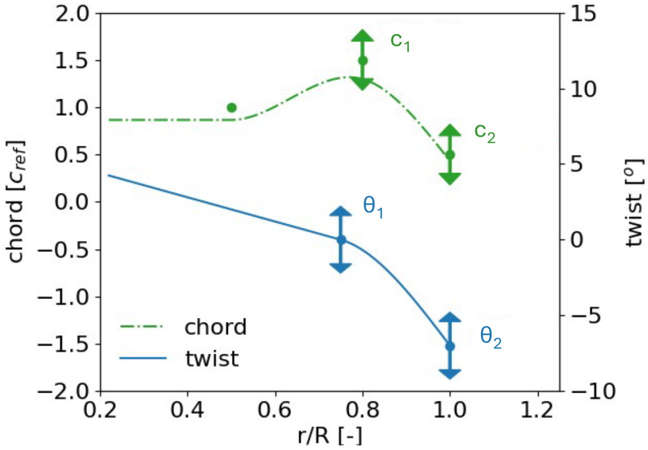

Fig. 6: Four design variables used in the initial rotor hover optimization problem, note that the chord law does not go exactly through the design points because of the rescaling to meet the thrust weighted solidity constraint.

with higher fidelity results from DLR which used their Hierarchical Kriging method to include RANS CFD results. The two approaches yield very similar results for optimal twist. Additionally, the small break in the twist at r/R=0.75 in the CAMRADII results is related to how the twist input was discretized in CAMRADII, a twist input value at r/R=0.75 was not directly available, so the exact details of the twist break are

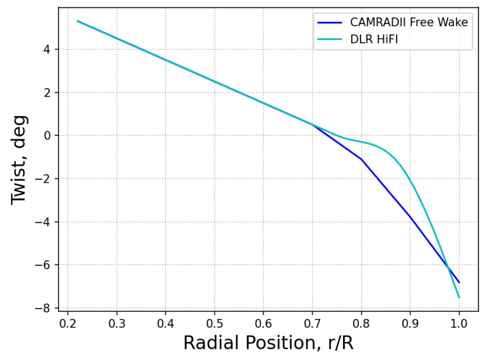

(a) Optimum twist law

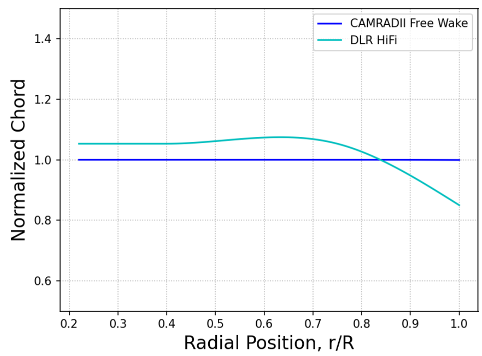

(b) Optimum chord law

Fig. 7: Comparison of optimal twist and chord laws from CAMRADII optimization with higher fidelity CFD surrogate results from InDeWo participant DLR (Ref. 14).

not properly captured. The lack of large differences in twist as seen in Figure 7a, suggest that the fidelity of free wake is sufficient to capture the impact of the twist design variables in the optimization problem. In terms of surrogate modeling therefore, we can expect CA to be a good lower fidelity surrogate in twist, and the portion of the Co-Kriging model that capture twist changes will likely be accurate without additional higher fidelity sampling.

Figure 7b, shows that there is a very large difference in the optimal chord law. In particular, the DLR results suggest that taper at the tip is beneficial to the overall hover figure of merit. In general, larger variations in optimal chord where seen among InDeWo participants. The use of a multifidelity surrogate to better capture the chord variations is desirable.

A multifidelity, single objective optimization was performed using the framework described above. The objective, again, was to maximize hover figure of merit, by adjusting just the two chord variables, subject to the constrain on rotor thrust weighted solidity. Twist is held constant at the baseline HART–II value of -8 deg linear twist along the blade radius. A Co-Kriging surrogate is built initially using a Latin hypercube sampling of the CA and one case of high fidelity run at the baseline HART–II design (constant chord). On this surrogate NSGA2 was run with a population of 10 members for 40 generations. Helios was then executed on the resulting optimal design and compared with the expected result from the surrogate model, completing one epoch of the design process outlined in Figure 1, steps 1-8. After each epoch of NSGA2 optimization, the surrogate is updated as shown in Figure 1, steps 9- & 10, using additional information from high fidelity analysis of the optimal design point and an additional point randomly selected in the design space. Figure 8 shows comparisons of subsequent epoch multifidelity GPR estimates of figure of merit variation over the design space of the two chord design variables.

As a point of comparision, Figure 8a is the resulting multifidelity surrogate constructed using 12 points of HELIOS *a priori* sampled in a two dimensional design space using Latin hypercube sampling. It can be seen in Figure 8, that the progressively update surrogate is approaching a similar optimal location, with the general relative trends in chord sensitivity in fewer than 12 high fidelity samples. These general trends in figure of merit for chord variables also appear to be similar to those seen by other workshop participants as shown in Figure 9.

Figure 10 shows the history of figure of merit prediction at the predicted optima from the multifidelity surrogate compared to the actual value calculated by HELIOS. By epoch 3 the progressively updated multifidelity surrogate, using 6 samples, has reached a similar accuracy prediction as the 12 Latin hypercube sampled surrogate. The exact values of the design variables found differ slightly, but show a very similar trend of taper at the tip with no increase in chord at the r/R=0.75 station. A comparison of chord laws between the baseline HART–II blade and the multifidelity surrogate optimized blade is shown in Figure 11. The taper found with the updated multifidelity surrogate model matches more closely with the taper found by other InDeWo participants for their blades which considered higher fidelity analysis.

This approach of re-sampling near predicted optimal design locations helps to alleviate the curse of dimensionality by reducing the number of points in each re-sampling to scale with the number of objective function evaluations and not the number of design variables. This reduces the total number of expensive higher fidelity function evaluations required and the total computational resource requirements. However, because multiple epochs of sampling are required, this does impose a potential wall-clock penalty since, unlike with a Latin hybercube sampling approach, all of the desired sample points are not known *a priori* and therefore cannot be all run in parallel. In practical terms on the DSRCs this means completing an epoch of optimization, identifying new high fidelity points and then queuing those points for execution. DSRC queue wait times can in some instances be on the same order of time as the case calculation time and contribute significantly to the total wall-clock time required to complete the optimization.

CONCLUSIONS

A method for incorporating the higher fidelity HELIOS CFD analysis into rotor blade optimization was demonstrated using a multifidelity surrogate approach. Because of the large increase in computational costs of using HELIOS compared to CA, care must be taken to select the right trade-off in solution setup for absolute accuracy versus computational cost. Exploration of the HART–II rotor problem, suggests that coarsening of the near-body mesh may be acceptable, but that the azimuthal time step and off-body mesh sizes need to remain similar to previously recommended best practices. The choice of DDES or RANS approaches to turbulence did not make a noticeable difference in computational cost. Further exploration of opportunities to speed-up time to solution are desirable. In particular, efforts to take advantage of the relative similarity of the various designs searched should be explored. These may include use of ways to morph the near-body mesh such that old design solutions can be used as restarts for new design solutions; reducing the number of revolutions required to achieve a steady condition. Use of a coarser azimuthal time step to start-up simulations, followed by finer resolution near the end will also significantly reduce wall clock time.

A successive refinement approach to updating the multifidelity surrogate was shown to reduce the number of cases needed to reach an optimal design compared with traditional *a priori* Latin hypercube sampling for a single objective function case. Future work will extend this approach to the multiobjective InDeWo design case with more design variables.

The results obtained for a optimization of a hovering HART–II rotor to the InDeWo hover design conditions compare well with those obtained by other participants. Similar trends on the impact of chord on figure of merit were seen, although the absolute magnitudes predicted were universally lower than those found by others. Some of this difference may be attributable to the fact that even the relatively coarse

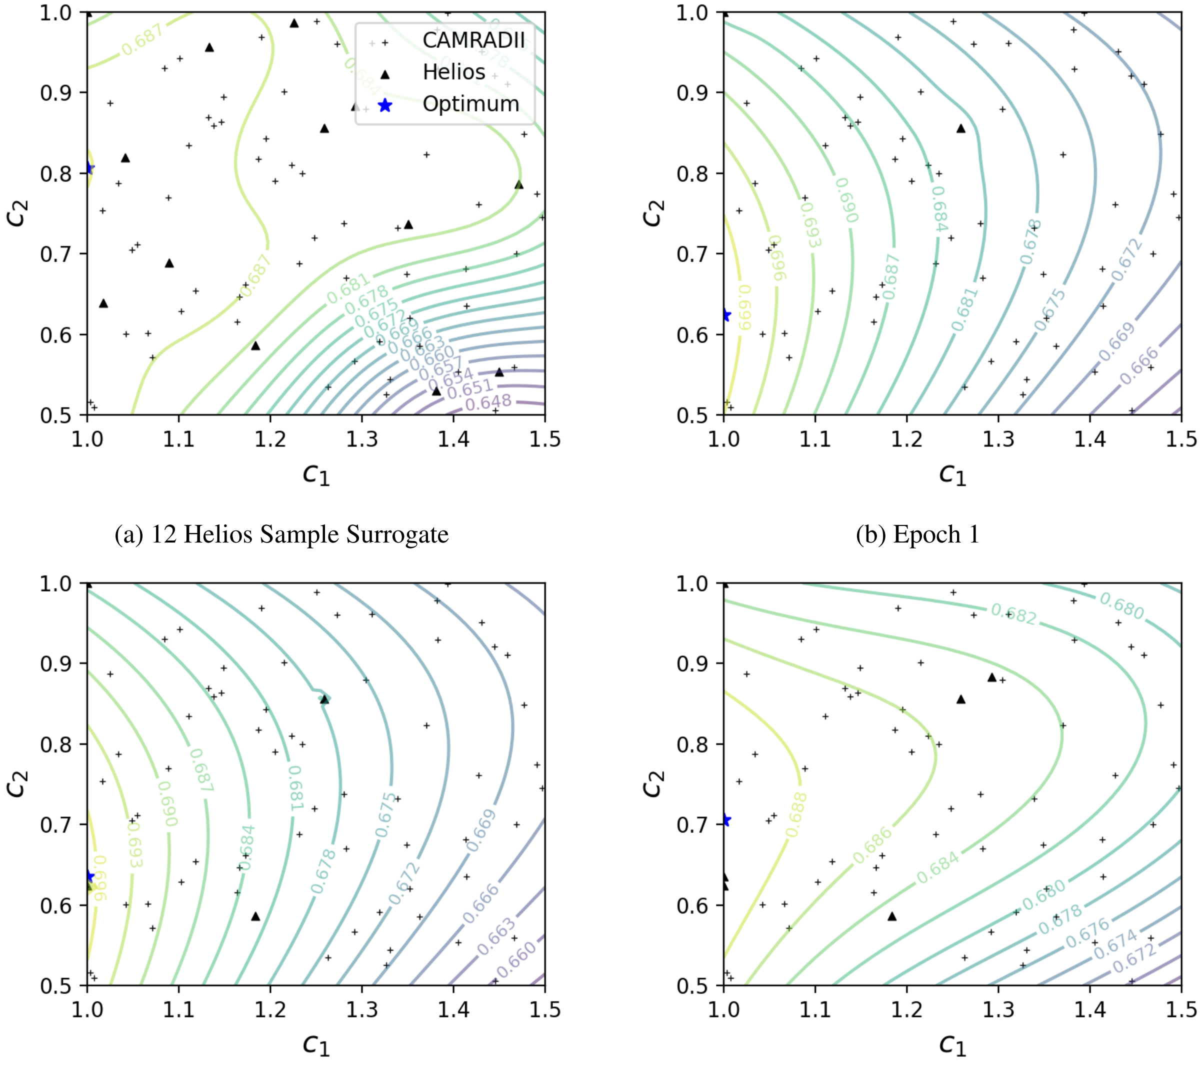

(a) 12 Helios Sample Surrogate

(b) Epoch 1

(c) Epoch 2

(d) Epoch 3

Fig. 8: Evolution of figure of merit surrogate with high fidelity sampling

HELIOS case used for optimization here was higher resolution than any of the other participants' CFD analysis models. The current approach has been successfully implemented in TORC. Future work is planned to extend into multiobjective optimization, and incorporate airfoil design variables using methods demonstrated recently (Ref. 28) as a complementary effort to this work.

ACKNOWLEDGMENT

The authors wish to thank CPT Charles R. Whittaker (USA) who performed the initial optimization in hover using CAMRADII for the 2023 workshop and whose results are briefly reported here in Figure 7. Special thanks also goes out to members of the HELIOS community at Ames: Rohit Jain, Jay Sitaraman, Dylan Jude, and Andy Wissink, in particular, for all of their patience in helping us debug CFD cases and explore best practices for optimization.

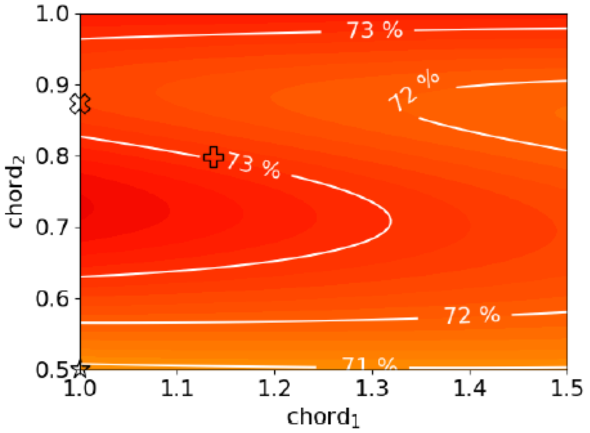

(a) DLR Chord Response Surface

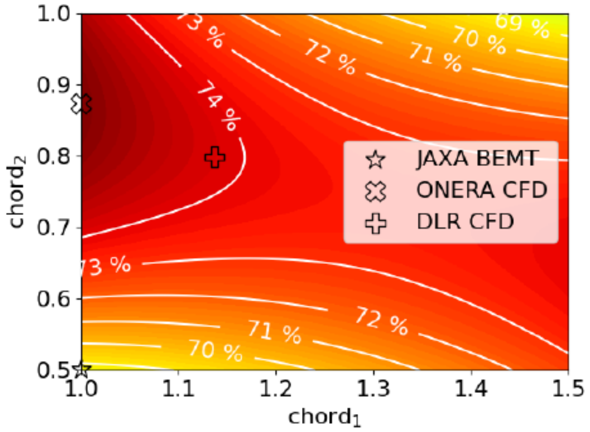

(b) ONERA Chord Response Surface

Fig. 9: Chord design variables figure of merit response surface from other InDeWo participants. (Ref. 14)

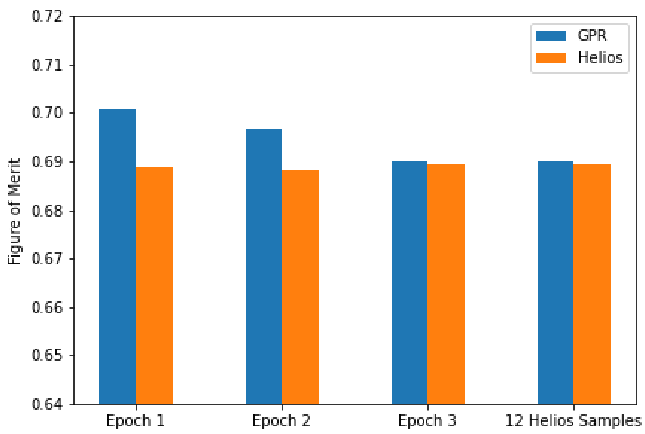

Fig. 10: Comparison of surrogate predicted figure of merit with actual Helios result

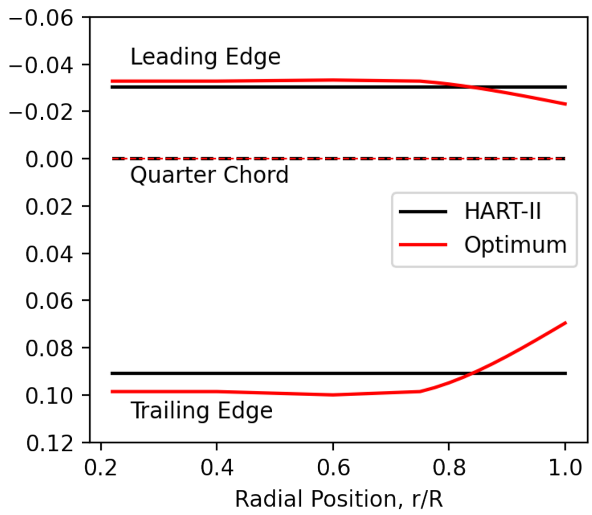

Fig. 11: Comparison of chord planforms for baseline HART–II rotor and optimum from multifidelity surrogate search, showing distinct tip taper in contrast to CA only results of Figure 7b.

REFERENCES

[1]Camarillo, G., “Army Directive 2024-03 (Army Digitial Engineering),” Memoranum, May 2024.

[2]Smith, M. J., “Computational Vertical Lift Aeromechanics and its Future in the Twenth-First Century,” VFS 78th Annual Forum and Technology Display, May 2022.

[3]Johnson, W., “Milestones in Rotorcraft Aeromechanics,” *Journal of the American Helicopter Society*, Vol. 56, (031001-1–24), 2011.

[4]Wood, T., “Fifty Years of Industry Perspective, 36th Alexander A. Nikolsky Honorary Lecture,” *Journal of the American Helicopter Society*, Vol. 65, 2020.
doi: 10.4050/JAHS.65.021001

[5]Min, B.-Y., Kim, J., Klimchenko, V., Neiswonger, J. C., Griffths, D., and Wake, B. E., “Full Vehicle Helios Model Performance Correlation with SB>1 Defiant Flight Test,” 80th Annual Forum of the Vertical Flight Society & Technology Display, May 2024.

[6]Jain, R., “CFD-CSD Pre-test Hover Predictions Validations for the HVAB Rotor,” AIAA SciTech Forum, January 2024.

[7]Jude, D., Sitaraman, J., and Wissink, A. M., “An Octree-based, Cartesian CFD Solver for Helio on CPU and GPU Architectures,” AIAA Scitech 2021 Forum, January 2021.

[8]Hosseinverdi, S., Sitaraman, J., and Jude, D., “A Performance-Portable Flow Solver for Rotorcraft Application,” 6th Decennial Aeromechanics Specialists’ Conference, February 2024.

[9]Leusink, D., Alfano, D., and Cinnella, P., “Multi-Fidelity Optimization Strategy for the Industrial Aerodynamic Design of Helicopter Rotor Blades,” *Aerospace Science and Technology*, Vol. 42, 2015, pp. 136–147.

[10]Collins, K. and Sankar, L., “Application of Low and High Fidelity Simulation Tools to Helicopter Rotor Blade Optimization,” AHS International 65th Annual Forum & Technology Display, 2009.

[11]Wilke, G., “Quieter and Greener Rotorcraft: Concurrent Aerodynamic and Acoustic Optimization,” *CEAS Aeronautical Journal*, 2021.

[12]Johnson, W., “Technology Drivers in the Development of CAMRAD II,” American Helicopter Society Aeromechanics Specialist Meeting, January 1994.

[13]Sitaraman, J., Hosseinverdi, S., Jude, D., Roget, B., Abras, J., Lakshminarayan, V., and Zagaris, G., “Progress in Massively Parallel and Performance Portable Overset Implemenations in HPCMP CREATE™-AV Helios,” AIAA SciTech Forum, January 2024.
doi: 10.2514/6.2024-1940

[14]Wilke, G., Bailley, J., Kimura, K., and Tanabe, Y., “JODROC results in hover extension to the InDeWO,” 1st International Design Workshop on Rotor Blade Optimization, September 2023.

[15]Wilke, G., Bailly, J., Kimura, K., and Tanabe, T., “JAXA-ONERA-DLR Cooperation: Results from Rotor Optimization in Hover,” *CEAS Aeronautical Journal*, April 2022.
doi: https://doi.org/10.1007/s13272-022-00580-8

[16]Bailly, J., Wilke, G., Kimura, K., and Tanabe, Y., “JAXA-ONERA-DLR Cooperation: Results from Multi-Point Aerodynamic Optimization of a Rotor in Hover and Forward Flight,” 74th European Rotorcraft Forum, September 2023.

[17]Hong, Y., Lee, D., Kang, Y.-E., and Yee, K., “Rotor Blade Design Optimization with Airfoil Consideration Using Advanced Reduced Order Models,” 80th Annual Forum of the Vertical Flight Society & Technology Display, May 2024.

[18]Blank, J. and Deb, K., “pymoo: Multi-Objective Optimization in Python,” *IEEE Access*, Vol. 8, 2020, pp. 89497–89509.

[19]Sinsay, J. D. and Sridharan, A., “Tool for Optimization of Rotorcraft Concepts Applied to eVTOL Rotor Design,” VFS Aeromechanics for Advanced Vertical Flight Technical Meeting, January 2022.

[20]Lakshminarayan, V., Wissink, A., Jain, R., and Sitaraman, J., “Improvements to Automated Strand Meshing Capabilities for Rotary Wing Applications,” VFS International Transformative Vertical Flight 2020 Meeting, January 2020.

[21]Küfmann, P., Bartels, R., van der Wall, B. G., Schneider, O., Holthusen, H., Gomes, J., and Postma, J., “The First Wind Tunnel Test of the DLR’s Multiple Swashplate System: Test Procedure and Preliminary Results,” AHS 72nd Annual Forum and Technology Display, May 2016.

[22]Jayaraman, B., Jain, R., and Lim, J., “HPCMP CREATE™-AV Helios Predictions of BVI Loading and Rotor Wakes of HART II Rotor,” AHS Technical Meeting on Aeromechanics Design for Vertical Lift, January 2016.

[23]Wilke, G., “Variable-Fidelity Methodology for the Aerodynamic Optimization of Helicopter Rotors,” *AIAA Journal*, Vol. 57, (8), August 2019, pp. 3145–3158.

[24]Forrester, A., Sóbester, A., and Keane, A., *Engineering Design via Surrogate Modelling*, John Wiley & Sons, Ltd., 2008.

[25]Walizer, L., Haehnel, R., Allen, L., and Wenren, Y., “Application of Mulit-fidelity Methods to Rotorcraft Performance Assessment,” Paper 2024-0012, AIAA SciTech Forum, January 2023.
doi: 10.2514/6.2024-0012

[26]Paleyes, A., Pullin, M., Mahsereci, M., McCollum, C., Lawrence, N., and González, J., “Emulation of Physical Proccesses with Emukit,” Second Workshop on Machine Learning and Physical Sciences, NeurIPS, 2019.

[27]GPy, "GPy: A Gaussian Process Framework in Python," http://github.com/SheffieldML/GPy, since 2012.

[28]Sridharan, A. and Sinsay, J. D., "A Comparative Evaluation of Select Shape Parameterization Approaches for Airfoil Optimization using Neural Networks," AIAA Aviation Forum, August 2024.