

Dragonfly Rotor Optimization using Machine Learning Applied to an OVERFLOW Generated Airfoil Database

Jason Cornelius
Aerospace Engineer
NASA Ames Research Center
Moffett Field, CA

Sven Schmitz
Boeing/A.D. Welliver Professor
The Pennsylvania State University
University Park, PA

ABSTRACT

NASA's 4th New Frontiers Mission is the Titan Dragonfly relocatable lander. This coaxial quadrotor vehicle will be launched on a rocket to Titan in 2028. Following a gravity assisted Earth flyby and an approximate 6-year transit, Dragonfly will enter the Titan atmosphere around 2034 with the goal of exploring Titan's pre-biotic chemistry and habitability. The multirotor design for this unique application has continually evolved since 2016 with constraints such as Titan's cryogenic atmosphere at 95 Kelvin (-288 F), gravity 14% that of Earth's, atmospheric density 440% of standard sea-level air, and the inability to test the entire system together under all these conditions until the first flight on Titan. This paper focuses on rotor design aspects of the Dragonfly lander and introduces a novel framework for multirotor design optimization considering multiple flight conditions. The methodology leverages machine learning methods and is demonstrated in the context of Dragonfly. A new OVERFLOW Machine Learning Airfoil Performance (PALMO) database is first presented. PALMO is then wrapped inside a Bayesian optimization framework and applied to a 4-rotor system (one side of the Dragonfly lander). Training data is generated on each iteration of the optimization using the CAMRAD-II comprehensive analysis software to evaluate successive rotor designs in multiple relevant flight conditions. An optimal design for the 4-rotor system was found with approximately 900 rotor designs analyzed in CAMRAD-II, which required 9 million queries of the PALMO surrogate models. This demonstration case evaluated 10,000,000 potential candidate rotor designs in 5.5 hours on 114 CPU cores using uniform inflow, and in 27.8 hours using the prescribed wake model. This work thus enables mid-fidelity rotor design optimization without requiring access to high-performance computing.

NOTATION

B*	Dragonfly Phase B 'Star' Rotor Design
CP	Control Point (Rotor Design Parameter)
c_i	Chord at CP _i [m]
N_b	Number of Blades (per rotor)
P	Rotor Power Requirement [W]
R	Rotor Blade Radius [m]
r	Radial Location [m]
th_i	Thickness to Chord Ratio at CP _i [%]
tw_i	Twist at CP _i [deg]
x_i	Taper CP _i Location [r/R]
y_i	Twist CP _i Location [r/R]
z_i	Thickness CP _i Location [r/R]
α	Angle-of-Attack [deg]
c_l	Airfoil Lift Coefficient
c_d	Airfoil Drag Coefficient
c_m	Airfoil Pitching Moment Coefficient

BEMT	Blade Element Momentum Theory
CFD	Computational Fluid Dynamics
CST	Class Shape Transformation
EI	Expected Improvement
FNN	Feed-Forward Neural Network
GPR	Gaussian Process Regression
HECC	(NASA) High-End Compute Capability
HPCMP	High Performance Computing Modernization Project
LaRC	NASA Langley Research Center
MAE	Mean Absolute Error
ML	Machine Learning
MOGA	Multi Objective Genetic Algorithm
NACA	National Advisory Committee for Aeronautics
PALMO	OVERFLOW Machine Learning Airfoil Performance Database
PCA	Principal Component Analysis
POD	Proper Orthogonal Decomposition
UIUC	University of Illinois at Urbana Champaign

ACRONYMS

APL	Johns Hopkins University Applied Physics Laboratory
ARC	NASA Ames Research Center

INTRODUCTION

Saturn's moon Titan was discovered by Dutch astronomer Christian Huygens in 1655. Many years later, the Voyager flybys of 1980-1981 identified it as one of the most interesting

planetary bodies in the solar system with a thick atmosphere, low gravity, and signs of pre-biotic organic chemistry on its surface, Ref. [1]. These early observations were confirmed and expanded upon by the Cassini-Huygens mission circa 2005. The Huygen’s probe photographed the surface of Titan and Cassini data has been used to create composite maps of Titan as shown in Figure 1. To investigate Titan further scientifically, NASA awarded the 4th New Frontiers mission to the Johns Hopkins University Applied Physics Laboratory (APL) to send a multi-mission radioisotope thermoelectric generator (MMRTG) powered rotorcraft the size of a small car to explore Titan’s surface and lower atmosphere, Ref. [2]. Figure 2 shows this relocatable lander called ‘Dragonfly’ that will launch in 2028 with an Earth gravity assist and subsequent outbound departure for a Titan arrival around 2034. As an ocean world with abundant carbon rich chemistry, Titan promises a treasure trove of planetary science data that may help us better understand the formation of life here on Earth, Ref. [3].




Figure 1. Titan Surface (left), Huygen’s Probe (top-right), Titan from Cassini (bottom-right), NASA.




Figure 2. Titan Dragonfly Relocatable Lander, NASA’s 4th New Frontiers Mission, APL.

As a multirotor vehicle, a coaxial quadrotor to be exact, Dragonfly’s rotor aerodynamic performance is a critical enabler for overall mission success. The mission requires an efficient rotor design tailored for the unique atmospheric and environmental conditions on Titan, which include low gravity (14% that of Earth), high atmospheric density (440% of sea-level standard air), and a cryogenic atmosphere (95 Kelvin or -288 Fahrenheit). The rotor design is also constrained by stowage in the aeroshell for its transit through space, and by its deployment upon Titan arrival. This work will discuss the Dragonfly rotor development in the context of these unique design constraints specific to rotary-wing flight on Titan.

MOTIVATION

In the rotorcraft community, airfoil performance look-up tables commonly referred to as C81 tables are used in various tools across almost all aeromechanics disciplines to yield fast results for rotor simulations. The characterization of airfoil performance across a range of Mach numbers, Reynolds numbers, and angles-of-attack remains a key aspect for a vast array of these blade-modeled rotorcraft analysis methods. Whether it be lifting-line modeling (CAMRAD-II, CHARM, RCAS), actuator disk modeling (RotCFD, OVERFLOW, HPCMP CREATE-AV Helios), or viscous vortex-particle methods (FLIGHTLAB, RCAS), the generation of C81 airfoil performance tables for these tools is laboriously carried out by countless engineers throughout the community each year. Furthermore, achieving accurate results demand that each new rotor conceptual design should use C81 tables based on the appropriate airfoil, Reynolds number, and Mach number distributions along the blade. Although these mid-fidelity tools rely on these tables to reduce their computational cost, the process of creating accurate C81 tables is still a formidable and computationally intensive task. Due to the computational cost of high-order accurate computational fluid dynamics (CFD) solvers such as OVERFLOW and ARC2D, airfoil performance tables are often generated with lower-fidelity methods or using the high-order accurate methods at proximal but miss-matched conditions. The authors have developed best practices over the years for generating C81 table input decks supporting Dragonfly’s rotor design with XFOIL, MSES, ARC2D, and OVERFLOW. Still, the computational cost associated with supporting each successive rotor design iteration is high.

This work leverages machine learning (ML) to enable 1) massive speedups in both airfoil and rotor performance predictions and 2) optimization over high-dimensionality and non-convex rotor design parameter spaces. This has led to the development of the OVERFLOW Machine Learning Airfoil Performance (PALMO) database enabling real-time high-order accurate airfoil performance estimation for airfoils within the database. This tool will enable rotorcraft designers and engineers to obtain OVERFLOW-quality C81 tables without the need for high performance computing. PALMO also vastly increases the efficacy of the rotorcraft conceptual designer by enabling the real-time updating of these OVERFLOW based C81 tables in each successive rotor

design iteration, capturing changes in airfoil, Mach number, and Reynolds number as a function of planform and RPM.

The PALMO database will be made publicly available soon via a NASA Technical Memorandum by Dr. Jason Cornelius at NASA Ames Research Center (ARC). Additionally, the AIAA Surrogates Modeling Technical Committee plans to use PALMO as an international benchmark database for aerospace machine learning research and development.

To demonstrate this novel capability, PALMO is coupled with CAMRAD-II in a new surrogate-model based rotor design optimization framework created in this work. This framework is demonstrated in the context of the Dragonfly lander rotor design optimization to yield candidate future rotor designs with improved aerodynamic performance on Titan. The methodology is general and can be extended to rotor design optimization for vehicles on Earth and other planetary bodies.

The primary objective of this work is to show an approach for machine learning leveraged rotor design optimization and analysis. This work first presents results from the PALMO database made from high-order accurate OVERFLOW CFD simulations. The database consists of approximately 60,000 National Advisory Committee for Aeronautics (NACA) 4-Series airfoil simulations generated using NASA’s High-End Compute Capability (HECC). The methodology will be shown to efficiently find global optimum designs in a high-dimensionality and non-convex parameter space. The approach simultaneously supports uncertainty quantification and the use of mixed-fidelity simulation data.

BACKGROUND

Given the highly multi-disciplinary nature of this work, it is necessary to provide a thorough background section on a few of the disciplines that will be discussed. Those include methods in supervised machine learning, airfoil surrogate modeling, rotor optimization, and surrogate-model based optimization for rotorcraft applications.

Methods in Supervised Machine Learning

Even as a subset of artificial intelligence, ML has many different approaches and methods, see Figure 3. This includes many common approaches such as linear regression, support vector machines, regression trees, ensemble of trees, Gaussian process regression (GPR), and feed-forward neural networks (FNN). These all fall into the category of supervised ML approaches. Previous work has identified GPRs and FNNs as the most suitable approaches for interpolating the types of databases used in this work, Ref. [4]. At a high level, GPR is a probabilistic modeling approach. It is effective for small datasets with a low number (typically less than 20) of input parameters. GPR has the large advantage that it simultaneously predicts both an output value along with an associated uncertainty. The scalability of it, however, is poor with computational cost to train the model increasing with a cubic dependency on the number of datapoints. FNNs can be readily applied to very large and complex parameter spaces but are data hungry and can be more challenging to create reliable models than simpler ML methods. Both approaches are leveraged in this work.

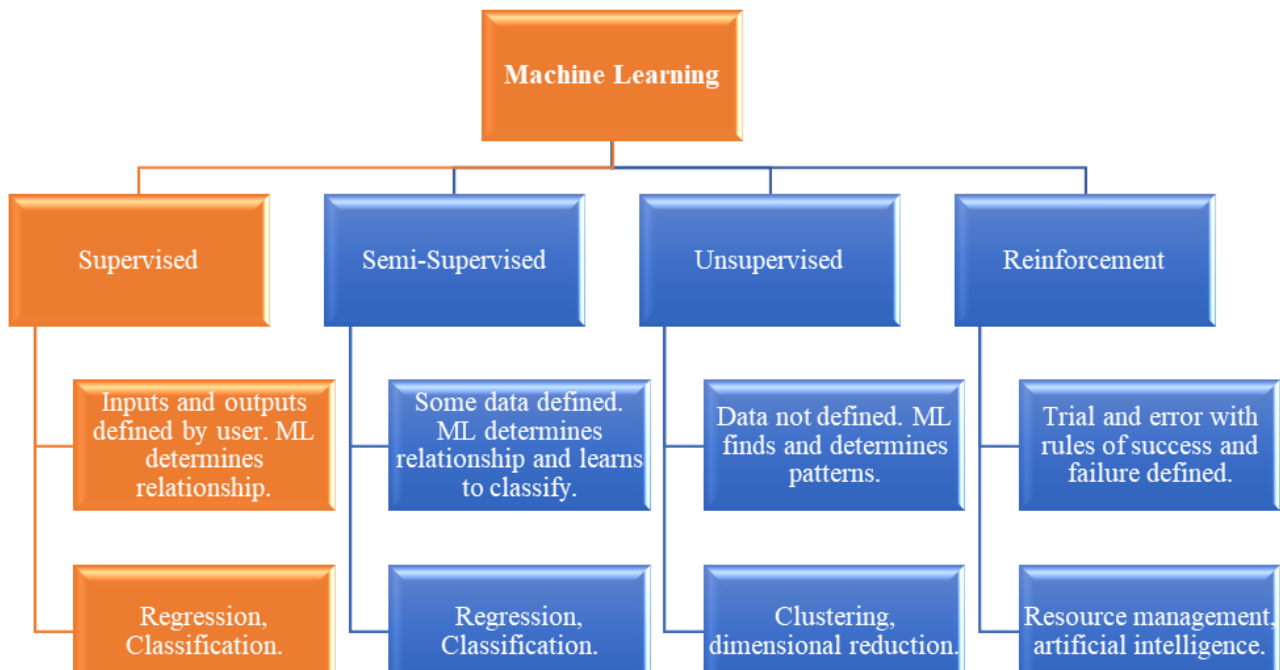


Figure 3. Summary of Machine Learning Approaches and Some Common Applications.

Airfoil Surrogate Modeling

State-of-the-art approaches for C81 airfoil performance table generation involve two-dimensional airfoil CFD analyses such as OVERFLOW or ARC2D. These tables are sometimes created with existing experimental data, however, but are then limited by the Mach and Reynolds numbers tested. Other times they are generated using a lower-fidelity approach such as XFOIL or MSES. The engineer must balance the need for increasing the accuracy of the C81 tables used against the computational cost and time required to create improved tables. Past studies by Patt and Youngren, Ref. [5], and Cornelius and Schmitz, Ref. [6], document both the need for higher refinement implementations of C81 tables and the improvements obtained through their implementation. Creating these high-density C81 tables with CFD in each iteration of conceptual design studies would be cost prohibitive, however, which currently restricts their implementation to preliminary or even detailed design. Conceptual designers often make simplifying assumptions as a result, such as using a set of C81 tables at constant Reynolds number even as the chord-based Reynolds number changes in successive design iterations. The alternative approach is to use fast but lower-fidelity methods such as XFOIL, or the XFOIL generated University of Illinois at Urbana Champaign (UIUC) database, Ref. [7], to update the airfoil performance tables between iterations. Given the high cost of creating these C81 tables, there has been a growing interest in the aviation community to leverage various ML approaches to derive highly accurate, low-cost surrogate models for predicting airfoil performance.

Some recent studies have used neural networks to create highly efficient airfoil performance surrogate models to update the C81 tables as the conceptual design progresses, but they have typically been based on lower-fidelity training data, Ref. [8]. Sridharan and Sinsay did apply neural networks to data from the thin-layer Navier stokes flow solver ARC2D using the wrapper C81Gen, Ref. [9], albeit with a coarse Mach discretization and at a single Reynolds number. The use of these surrogate models for airfoil performance prediction has recently received much attention, Refs. [10-13]. Li et al. recently provided a review on this topic, Ref. [14]. These studies, however, typically rely on Class Shape Transformations (CST) such as Bernstein and Chebyshev polynomials, Refs. [15-17]. Although this process has been adopted as the leading research approach for airfoil shape optimization, this parametrization of the airfoil shape introduces some discrepancies as compared to the original CFD calculations for airfoils in the training datasets.

Rotor Optimization

Aerospace optimization has been an active field of research for many years. Multi-disciplinary optimization methods including gradient-based and gradient-free descent have been proposed by a multitude of groups with current capabilities documented by NASA, Ref. [18], and benchmarks provided by Gray and Martins, Ref. [19]. The NASA Vision 2040

report, however, documents the need for still improved capabilities in design optimization, stating that current methods are restricted by prohibitive computational cost when using mid- to high-fidelity computational data, Ref. [20]. The report cites the advent of machine learning and surrogate modeling as a potentially transformative tool to be used in aerospace design. Some groups, such as the Army's High Performance Computing Modernization Program (HPCMP) are investing heavily to investigate applications of surrogate modeling, Ref. [21].

A few researchers have been specifically studying design optimization as it pertains to rotary-wing aircraft. Collins, Sankar, and Mavris used a mixed-fidelity simulation approach bridging dynamic inflow in RCAS with computational fluid dynamics (CFD) from GT-HYBRID, Ref. [22]. Surrogate models were generated to correlate and scale the lower-fidelity data to the higher-fidelity, achieving fast yet accurate optimization. The work resulted in a rotor with optimal efficiency and constrained vibrations. Another study used CAMRAD-II with parametric equations to represent blade planform changes from the baseline UH-60A rotor, Ref. [23]. This work achieved reductions in power of 1% in hover and up to 17% in high advance ratio cruise, albeit with increases in steady chord bending moment and pitch link loads. The author cited that a larger parameter space and combination with CFD data would likely improve the results.

Several groups have studied design optimization of compound rotorcraft. Lim et al. used a two-tier optimization to achieve an optimized rotor blade cross-sectional design and improved rotor aerodynamic performance with reduced vibrations, Ref. [24]. The study used a fixed aerodynamic outer mold line and cited the inclusion of airfoils in the optimization as a future area of improvement. Hersey et al. combined nonlinear finite element models with a free-vortex wake model for multi-objective optimization, Ref. [25]. They split the problem into a sequence of optimization tasks using radial basis functions to interpolate the results. This limited the study, however, to eight design variables or less, which was cited as likely too few for rotor optimization studies. Another more recent study by Sridharan and Govindarajan applied the Simultaneous Analysis and Design (SAND) approach to parallelize the optimization of a lift-augmented quadrotor biplane tailsitter, Ref. [26]. They successfully optimized the design with 39 variables, albeit using low-fidelity models.

Lim et al. and Allen et al. more recently presented results on an airfoil surrogate-model based framework for rotor design optimization, Refs. [27-28]. The framework uses the Dakota Multi Objective Genetic Algorithm (MOGA) with RCAS and ARC2D, via C81Gen wrapper, to optimize airfoils for three outboard stations of the UH-60A rotor blade. The framework first generates a large set of ARC2D airfoil simulation data at the beginning of the optimization, which required 70,000 CPU hours, or 2.5 weeks on 1100 CPU cores. Four airfoil shape design parameters were used resulting in 12 design

parameters in the rotor optimization. The studies achieved a 6-9% reduction in rotor power for hover and cruise, respectively, with 50,000 RCAS evaluations. This required 3-4 days on 10 nodes of their supercomputer resulting in a full optimization cycle of about three weeks using up to 1100 CPU cores. The authors cited the need for a larger design space, including rotor blade planform, acoustics, and structural constraints, to support future design applications.

Koning has also recently published work on the Evolutionary aLgorithm for Iterative Studies of Aeromechanics (ELISA) tool, Ref. [29]. ELISA also applies a MOGA approach leveraging CAMRAD-II as the function evaluator, but to the aerodynamic optimization of Martian rotors in hover. Koning showed a 19% increase in figure of merit as compared to the Ingenuity Mars helicopter that flew 72 successful flights on Mars, Refs. [30-32], which has been extensively analyzed with CAMRAD-II. This work notes the heuristic MOGA approach is not ideally suited for quickly finding a global optimum but is well suited to explore a wide design space to increase understanding for future design.

Other recent studies explored adjoint-based parameter sensitivity studies for use in gradient-based optimization, Refs. [33-34]. The studies leveraged these approaches to reduce the required parameter space enabling the use of higher-fidelity simulation data.

Surrogate-Model Based Optimization in Rotorcraft

Most of these works cite the curse of dimensionality of the rotor design optimization problem as the largest challenge to overcome. Advanced rotor optimization including airfoil shape, planform, and structural load considerations, for example, could easily require tens of design parameters, which poses challenges for many of the existing optimization methods. A few researchers have recently worked to combat these issues using surrogate-model based optimization.

Sridharan recently used a 70,000 airfoil performance database generated using ARC2D CFD simulations, Ref. [35]. He created several hundred Gaussian Process Regression (GPR) models to predict airfoil performance using 8 design parameters for Chebyshev CST polynomials. Several hundred airfoils from the UIUC database were simulated at 126 flow conditions (-5 to 15 deg angle-of-attack, 6 Mach numbers from 0.3 to 0.9, and a constant Reynolds number to Mach scaling). These surrogate models were used in the Tool for Optimization of Rotorcraft Concepts (TORC) to carry out airfoil optimization. The implementation of the surrogate-model based airfoil database enabled hundreds of thousands of ARC2D-based airfoil design evaluations. These airfoils were then used in a 6-parameter planform (bi-linear twist and bi-linear taper) optimization using low-fidelity blade element momentum theory (BEMT) rotor performance evaluations. Another similar study by Shalu, Govindarajan, Sridharan, and Singh used a nearly identical approach by wrapping airfoil performance surrogate models of XFOIL data with BEMT rotor performance evaluations, Ref. [36]. Airfoils were first

optimized and then used in bi-linear twist and taper proprotor optimization with 6 design variables for hover and cruise.

Peters recently showed large computational cost reductions in high-fidelity rotor performance calculations using proper orthogonal decomposition (POD), Ref. [37]. A several order-of-magnitude computational cost reduction was achieved, and the resulting PODs were then successfully used for parametric rotor design optimization.

Another recent work by Anusonti-Inthra used GPR modeling to achieve surrogate-model based multi-fidelity aerodynamic prediction of a tiltrotor pylon, Ref. [38]. Fifty low-fidelity simulations using the Reduced-Order Aerodynamic Model (ROAM) and five high-fidelity CFD simulations using Helios with FUN3D were combined to create a single surrogate model for predicting pylon lift, drag, and pitching moment as a function of flight condition with a predictive speed of milliseconds. A comprehensive survey of multi-fidelity methods for uncertainty quantification and optimization was recently published by Peherstorfer, Willcox, and Gunzburger, Ref. [39].

Erhard and Alonso recently presented a novel approach using Bayesian optimization together with multi-fidelity modeling for coaxial rotor optimization in both hover and forward flight, Ref. [40]. The work parametrized the rotor blade design using Bezier curves, similar to the work by Koning, which effectively minimized the number of required design variables. The results found several candidate designs near the global optimum with very different rotor planforms, suggesting a highly non-convex and multi-modal solution space. This suggests that gradient-free approaches may be best suited for rotor optimization studies with many free design variables.

Summary and Path Forward

The works cited in this brief survey of airfoil surrogate modeling, rotor optimization, and surrogate-model based optimization represent the current state-of-the-art in these disciplines. Many of the authors cite the curse of dimensionality and cost-prohibitive high-fidelity training data as the major barriers to further progress. This work presents a methodology for machine learning leveraged rotor optimization by coupling high-order accurate based airfoil performance surrogate models directly into the rotor optimization. Previous studies that have done a similar approach then use a gradient-based or evolutionary algorithm to optimize the rotor. This study, however, further leverages machine learning by wrapping the rotor design framework with a Bayesian optimizer. This approach allows for the evaluation of several million candidate rotor designs on each iteration of the optimization and selects the most promising rotor designs for the next batch of analysis in CAMRAD-II.

To provide additional context on the constraints of the rotor design optimization carried out in this work, a brief history of the Dragonfly rotor development will now be presented. The

surrogate-model based design optimization framework, including the PALMO database, will then be discussed in detail with sample results presented.

DRAGONFLY ROTOR DESIGN

Several generations of Dragonfly rotors have been developed to date including the Phase A, Phase B, Phase B* (B Star), and Phase C rotors. (NASA SMD programs are staged in spacecraft development ‘phases’; accordingly, these updated rotor designs reflect an increased maturity in the overall spacecraft development effort.) Each of these rotor design iterations was arrived at after the analysis of many possible rotor designs to support the program requirements and constraints of flight on Titan. As such, each major rotor design iteration could have several various rotor designs associated with it, but a final variant was used for various rotorcraft and entry and descent analyses across the program.

Dragonfly Phase A Rotor Design

From the very beginning of Dragonfly’s rotor design, constraints on the rotor radius were driven by the size of the aeroshell that will protect Dragonfly on its seven-year journey to Titan. Beyond this simple geometric constraint, the previously mentioned intricacies of the Titan atmospheric environment led to several key design decisions. For example, Titan’s cryogenic atmosphere at 95 Kelvin (-288 F) largely precluded the use of an articulated rotor system. One could theoretically use heated hinges, but thermal heat-leakage into the cryogenic atmosphere is one of the largest technical challenges for a Titan-bound lander. As such, all Dragonfly rotor design iterations use fixed-pitch variable-speed (RPM-controlled) rotor systems. Also related to Titan’s cryogenic atmosphere is the choice of material. Although rotor mass is a small fraction of Dragonfly’s total weight, carbon fiber composites were still considered and investigated. A lack of available strength and fatigue data at cryogenic temperatures, however, has up to this point constrained the design iterations to aluminum alloys and titanium.

The Phase A rotors were designed using the NACA 5- and 6-Series airfoils due to the abundance of available experimental data and their widespread use as ‘rotorcraft’ airfoils. One key benefit of the NACA 5-Series is its low airfoil pitching moment, which is useful in rotorcraft applications to alleviate rotor blade torsion and high pitch-link loads. For Dragonfly’s stiff fixed-pitch rotors, however, the time-varying rotor blade pitching moment about the quarter chord is nearly negligible.

Dragonfly Phase B (and B*) Rotor Designs

The complications that Reynolds number scaling poses to the program engendered the Phase B rotor design iteration, which uses NACA 4-Series airfoils. These are considered 1st generation rotorcraft airfoils and have consistent performance across a very wide range of Reynolds number. The 4-Series is also largely insensitive to surface roughness, which is advantageous to mitigate performance degradation from dust impingement on the rotor blade leading edge. In addition, this

airfoil exhibits low suction peaks reducing the chance of condensation, Ref. [41], a large trailing edge angle for manufacturing, and docile stall behavior relevant for descent conditions.

The move to the Phase B rotor system also transitioned from a hyperbolic to a linear chord distribution. This was done for manufacturing and flight performance considerations. High rotor torque was also an issue with the Phase A rotor design, which led to the Phase B rotor having a slightly smaller radius to trade operating speed for torque. One of the unique design considerations of Dragonfly’s rotor system is the complex vibration environment encountered by using fixed-pitch rotors in edgewise flight. This requires careful tailoring of the Dragonfly rotor blade natural frequencies to mitigate large vibrations. The fact that Dragonfly employs variable-speed (RPM controlled) rotors adds further complication over conventional rotorcraft in that the rotor excitation frequencies constantly vary over a wide range throughout each flight. This ultimately led to a structural requirement for very stiff rotor blades to place the 1st flap frequency resonance crossover above the 1st and 2nd rotor blade passage excitation frequencies.

The Phase B* rotor was created to meet this new requirement. It has higher solidity to increase rotor blade cross-sectional area. This resulted in a large increase in the blade’s 1st flap frequency and enabled the rotor to support a higher thrust level, which created additional thrust margin. The increase in cross-sectional area also, somewhat counterintuitively, permitted a reduction in total rotor weight and inertia. This was because the efficiency of a larger cross-sectional area allowed a smaller skin thickness to be used. A full-scale Dragonfly B* coaxial rotor system was tested in the NASA Langley Research Center’s (LaRC) Transonic Dynamics Tunnel (TDT), as shown in Fig. 4, under both r134-a heavy gas and air conditions, Ref. [4, 42].



Figure 4. NASA LaRC Transonic Dynamics Tunnel Test, Dragonfly Phase B* Coaxial Rotor System.

For testing of the full-scale Dragonfly rotors, the United States only has a handful of facilities that can approximate the

aerodynamic conditions a rotor will experience on Titan. The TDT is one such facility and can replace the test section's air with r134-a heavy-gas. TDT testing achieves a chord-based Reynolds number about 1/3rd that of Titan, which is much closer than the order of magnitude discrepancy with one would experience with testing in air. Other U.S. assets that can more closely approximate Titan conditions exist, such as the LaRC National Transonic Facility, albeit with minimum test section speeds far exceeding the planned flight envelope of Dragonfly or with too small a tunnel cross-sectional area.

Dragonfly Phase C Rotor Design

The Phase C rotor is very similar to the B* rotor, but with a slightly increased solidity, an improved structural design, and a further improved twist distribution complementing current (at the time) planned flight conditions on Titan. Both the B* and Phase C rotor designs have solidity in-between the Phase A and B rotors, which is a balance to accommodate the available motor torque and structural requirements. The Phase C rotor design implements linear thickness and taper distributions of NACA 4-Series airfoils, which is congruent with the Phase B* rotor tested in the NASA LaRC TDT. The current Phase C Dragonfly rotor design is used across the program for rotor performance analysis, mission planning, and requirements specification. Rotor performance analyses inform requirements for various systems such as the lander's motors, power distribution, and guidance, navigation, and control systems.

Dragonfly Interactional Aerodynamics Testing

Another key aspect of designing a rotorcraft for flight on Titan is the ability to experimentally test the rotor design under representative conditions. The B* rotor tested under heavy-gas conditions in the LaRC TDT has already been shown. On the other hand, the much easier air testing presents challenges with aerodynamic conditions miss-matched from the Titan atmosphere. For example, if a rotor on Titan has a chord-based Reynolds number of three million, that number is closer to 300,000 in air. This large Reynolds number discrepancy between the testing environment and real conditions on Titan means that airfoil families sensitive to Reynolds number changes can behave very differently in the two different environments. As an example, an NACA 6-Series rotor may behave according to design on Titan, but then stall prematurely in an air test environment and at notably reduced lift-curve slope. This would make the CFD model validation process arduous and increase the risk carried by the program as no test is possible at true full-scale Titan conditions.

Still, several Dragonfly wind tunnel tests in air have been conducted to build confidence in the CFD modeling approaches used on the program, especially in the realm of interactional rotor-rotor and rotor-fuselage aerodynamics. The program has made entries in the NASA LaRC 14-by 22-ft. Wind Tunnel and the NASA ARC 80-by 120-ft. Wind Tunnel to test scaled models of the complete Dragonfly lander, Figs. 5-6, Refs. [4, 43-44].




Figure 5. NASA LaRC 14-by 22-ft. Wind Tunnel Test, Dragonfly Earthbound Integrated Test Platform.



Figure 6. NASA ARC NFAC 80-by 120-ft. Wind Tunnel Test, Dragonfly Scale Model with Entry Backshell.

Dragonfly Rotor Design Summary

As the Dragonfly program continues to evolve, another iteration of the rotor design could become necessary. Factors far outside the rotor performance alone influence this need, such as the motor, power system, and flight control system capabilities. The rotor design and subsequent performance on Titan, however, have a highly multi-disciplinary nature with a multi-objective design space. As such, this work presents a novel framework created by the authors for a multi-tiered ML leveraged multi-objective rotor design optimization. This framework is used to optimize potential future candidate rotor designs in the context of Dragonfly. The approach optimizes the rotor design by analyzing a 4-rotor CAMRAD-II model (one side of the Dragonfly lander) in several flight conditions relevant to the nominal planned mission on Titan. This is one of the first works using mid-fidelity models to optimize a system with rotor-rotor interferences across multiple flight conditions. OVERFLOW based C81 table generation is coupled directly into the optimization framework using ML.

The framework and methodology described in the subsequent sections is general and can be extended to rotorcraft design problems on Earth and other planetary bodies.

The first step in achieving this multi-disciplinary and multi-objective rotor optimization in a large parameter design space involves airfoil surrogate modeling. As such, the next section describes the application of machine learning to an OVERFLOW generated airfoil performance database.

THE OVERFLOW MACHINE LEARNING AIRFOIL PERFORMANCE DATABASE

Recent design and analysis efforts on the program have led to C81 table generation for the Phase B* and C rotor designs using 1) standard-air, 2) r134-a heavy-gas, and 3) Titan aerodynamic conditions. These C81 tables were created using the AFTGen software, Ref. [45], as a wrapper for OVERFLOW, Ref. [46]. Recent improvements in modeling approaches for hybrid BEMT-CFD solvers have led to highly dense C81 input decks for modeling Dragonfly’s fixed-pitch variable-speed rotors, Ref. [47]. An example rotor blade C81 discretization is shown in Fig. 7 for the off-the-shelf rotors used on the Dragonfly sub-scale model tested at NASA LaRC and ARC. Although these rotors are not of a similar planform to the Titan Dragonfly rotors, they have successfully been used for model validation of rotor-rotor and rotor-fuselage interactional aerodynamics. Large radial changes in chord, multiple airfoils, and the variable-speed operation necessitate this high-density modeling. Each permutation of rotor design iteration and atmospheric conditions requires roughly 3,000 OVERFLOW simulations parametrized over angle-of-attack, Mach number, and Reynolds number. Generating C81 tables for the B* and C rotor designs under the previously mentioned three atmospheres requires nearly 20,000 high-order accurate OVERFLOW airfoil simulations, which is an extremely time intensive and computationally demanding task.




Figure 7. High-Density Airfoil C81 Table Discretization for Variable-Speed Rotors with Large Chord Changes.

To alleviate these computational demands, the creation of an airfoil performance database is now presented using the high order accurate OVERFLOW CFD solver. Machine learning methods are then used to create surrogate models of the database for interpolation across the non-linear and high dimensionality training dataset. This enables real-time airfoil performance prediction (towards C81 table generation) at intermediate Mach and Reynolds numbers based on the high-order accurate OVERFLOW data. The surrogate models also enable C81 table generation for intermediate airfoils not explicitly included in the training dataset.

PALMO Database Generation

The airfoil performance database was generated using the NASA High-End Compute Capability (HECC). The first set of airfoils included in the database are the NACA 4-Series, Ref. [48], which have been used in the Dragonfly Phase B* and C rotor designs with an operational Reynolds number range of 1-3 million and a blade-tip Mach number of 0.2-0.4. This is a well understood flow regime for NACA 4-digit airfoils, but experience on the Dragonfly development effort has established the necessity of employing high-fidelity CFD to generate C81 airfoil characteristics for the project.

OVERFLOW simulations second-order accurate in time and fourth-order accurate in space are run with Spalart-Allmaras turbulence closure to develop the airfoil performance training datasets, Ref. [49]. The airfoil simulations are run using the previously mentioned OVERFLOW wrapper AFTGen. Grid studies have been carried out to ensure grid convergence to less than 1% in the linear region of the lift-curve slope following approaches documented by Cornelius et al., Ref. [43]. A typical simulation has 501 wrap-around points, 601 grid points normal to the airfoil, and 41 points on the blunt trailing edge yielding approximately 325,000 cells.

An example grid for the NACA 4418, a mid-blade Dragonfly airfoil, is included in Figure 8. For each airfoil, 3,280 OVERFLOW simulations are run at all possible combinations of the Mach number, Reynolds number, and angles-of-attack reported in Table 1.




Figure 8. OVERFLOW O-Grid for NACA 4418.

Table 1. Parametrization of 3,280 OVERFLOW Simulations per Airfoil in the PALMO Database.

Characteristic	Discretization
Mach Number	0.25, 0.35, 0.45, 0.55, 0.65, 0.70, 0.75, 0.80, 0.85, 0.90
Reynolds Number	75k, 125k, 250k, 500k, 1M, 2M, 4M, 8M
Angle-of-Attack	-20, -19, -18, -17, -16, -15, -14, -13, -12, -11, -10, -9, -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20

The foundation of the in-development PALMO methodology is the airfoil base cube shown in Figure 9, which is a high-density parametrization of OVERFLOW simulation data using the conditions from Table 1. A log scale is used to display the Reynolds number since it varies exponentially. Each base cube is generated using 16 Broadwell compute nodes on the NASA-HECC supercomputer, which is 448 CPU cores. A typical cube runs for approximately 5 days of wall-clock time, resulting in a total computational cost of 54k CPU hours.




Figure 9. PALMO Base Cube, Parametrization of Mach Number, Reynolds Number, and Angle-of-Attack.

This base cube is the first level of the database. PALMO is then increased in complexity by adding additional dimensions with the inclusion of airfoil parametrization. Figure 10 shows a parametrization of the NACA 4-Series with variations in percent camber and thickness being added to the database.

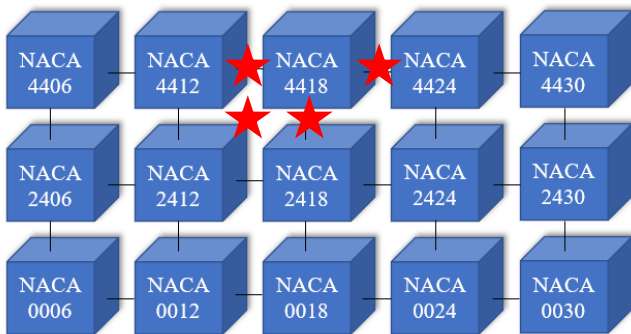


Figure 10. NACA 4-Series PALMO Database, 62k OVERFLOW Simulations, (*Red stars are test data).

This allows C81 tables to be generated for any arbitrary combination of camber and thickness within the bounds of the training data, i.e., from an NACA 0006 to an NACA 4430. The red stars on Fig. 10 are examples of where the database could be queried and represent additional test data that was generated off-axis from the original 4-Series parametrization. The PALMO 4-Series database has the 15 base-cubes shown in Fig. 10 along with additional off-axis test cubes for the NACA 3415, 3418, 4415, and 4421. These 19 base-cubes

represent approximately 62,000 high-order accurate OVERFLOW simulations and required approximately 1,000,000 CPU hours. Surrogate models developed with this database would have five input parameters, the original base cube plus the 4-Series sectional camber and thickness.

PALMO Surrogate Modeling

The current literature in airfoil surrogate modeling focuses on airfoil shape optimization and inverse design routines towards creating new airfoils. PALMO, on the other hand, is targeting real-time airfoil performance estimation with high accuracy for existing and commonly used airfoils. This database will enable more accurate conceptual design and analysis. For the current study, PALMO is restricted to the 4-Series database. PALMO will be expanded in the future through the addition of various airfoil families relevant to the rotorcraft community, starting with the NACA 5-Series and 6-Series.

Following the OVERFLOW database generation, surrogate models were developed to enable the real-time querying of any arbitrary combination of input parameters within the bounds of the training data. The bounds were shown in Table 1 along with the NACA 4-Series thickness from 6% through 30% and camber from 0% through 4%. Surrogate models were developed using both GPR and FNNs. Due to the size of the database with 62,000 simulations, 5 input parameters, and three output parameters (life, drag, and pitching moment coefficients), FNNs were selected due to their lower computational cost when using large datasets.

As an early example, a single set of FNNs, one for each airfoil performance coefficient, were generated for the NACA 4418 base cube. These surrogate models have Reynolds number, Mach number, and angle-of-attack as inputs. To assess the predictive accuracy at the various operating conditions, Table 2 reports test metrics for low, nominal, and high angles-of-attack (α). These test statistics cover the full range of conditions reported in Table 1. These data were generated by segmenting the base-cube simulation output into training and testing portions and then using the testing holdout for these calculations. The predictive accuracy is within 1% for much of the dataset, with high negative angles-of-attack still within approximately 3%.

Table 2. Prediction Accuracy of Example PALMO Surrogate Model (Neural Network) across AOA Range.

MAE	c_l	c_d	c_m
$\alpha < -5$	2.5%	2.1%	3.2%
$-5 < \alpha < 15$	1.0%	0.76%	0.6%
$\alpha > 15$	0.32%	0.63%	0.65%

Taking this a step further, FNNs were trained using a large subset of the PALMO 4-Series database. The NACA 0006 through the NACA 4424 (12 base-cubes) were used to train an FNN with five inputs. Results for the MAE are reported in Table 3 using two different test sets. The first used a 5% holdout from the database, which is about 2,000 OVERFLOW

simulations. The second test, which was much more rigorous, was to predict all 3,280 conditions of a base cube that was not included in the training data. Airfoil performance for the NACA 3415 was predicted, and subsequently compared to an OVERFLOW base cube run after the fact. The same MAE test statistics for the NACA 3415 base cube are also reported for MATLAB’s multi-dimensional linear interpolation. To display the PALMO surrogate-model accuracy visually, comparisons of the ML neural network predictions are plotted against the true OVERFLOW airfoil simulation data in Fig. 11 for the full angle-of-attack range.

Table 3. PALMO 4-Series Surrogate Model MAE.

Mean Absolute Error (MAE)	c_l	c_d	c_m
5% Hold-out	0.0088	0.0022	0.0026
NACA 3415 (3,280)	0.0157	0.0053	0.0055
NACA 3415- Linear Interpolation	0.0193	0.0040	0.0042




Figure 11. PALMO Predictive Accuracy, NACA 3415 Base-Cube, Reynolds Number: 500k, Mach: 0.25.

In general, all test statistics fall well within typical requirements for conceptual design. If separated into low, mid, and high angles-of-attack as reported in Table 2, the predictive accuracy would likely be even higher for the middle angles-of-attack where a rotor design would most likely be operating under nominal conditions. As for the comparison of the surrogate model to linear interpolation, the two are similar for all three metrics, with the lift coefficient better predicted by the FNN and the drag and pitching moment better predicted by linear interpolation. The linear interpolation performs this well only because of the very high point density of the underlying OVERFLOW training data. This can be thought of as a ‘tangent to the curve’ being a good local approximation. Further investigations in adaptive sampling methods will be used to reduce the number of simulations required for future PALMO base cubes. This will likely result in the FNN surrogate-model prediction accuracy far surpassing linear interpolation. Further such analyses will be presented in a more thorough future documentation of the PALMO database.

The comparisons in Fig. 11 are for a Reynolds number of 500,000 and a Mach number of 0.25. The linear interpolation results are also plotted. As a reminder, this OVERFLOW data for the NACA 3415 was generated after the surrogate model had been trained, i.e., it was not included in the surrogate-model training data. All metrics including lift, drag, and pitching moment coefficient are predicted well by the surrogate model. The drag and pitching moment coefficient comparisons highlight some discrepancy with linear interpolation at the lower-left corner of the drag-bucket. The surrogate model, however, captures these features well.

Figure 12 reports comparisons for a Reynolds number of two million and a Mach number of 0.25 to highlight the range of the PALMO database. Similar observations can be made with the surrogate model predicting the OVERFLOW data well.

ROTOR OPTIMIZATION FRAMEWORK

To show the utility of machine learning (ML) for rotor optimization, this work created a new framework to tie together several different components of the rotor design optimization process. The entire framework including the PALMO database module was custom scripted by the authors using Python, and this is the first publication of the work.

ML is leveraged throughout several different portions of the framework to achieve massive searches of the design space with little computational cost. This framework is depicted in Figure 13, which uses Python to couple the PALMO database with CAMRAD-II, Ref. [50], for rotor optimization studies. The various portions of the framework will be briefly described. This paper is the first published use of CAMRAD-II as applied to the Dragonfly rotor aerodynamic design and analysis effort; previous work has been performed using various mid- to high-fidelity CFD solvers. CAMRAD-II was selected for this rotor optimization work because of its relative computational efficiency as compared

to mid-fidelity CFD, while still being able to capture rotor-to-rotor interactional aerodynamics.

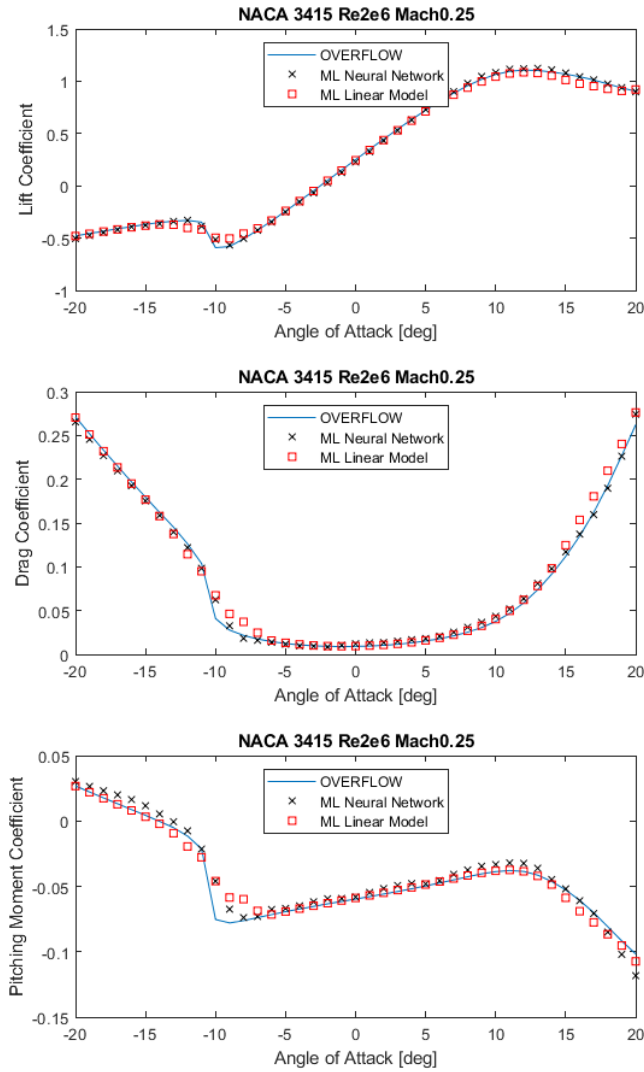


Figure 12. PALMO Predictive Accuracy, NACA 3415 Base-Cube, Reynolds Number: 2M, Mach: 0.25.

CAMRAD AUTO

This Python script is the main module that organizes the framework. It calls various other Python modules and manages the running of all these scripts on the NASA High-End Compute Capability (HECC).

Rotor Optimization: Initial Seeding

The rotor optimization script is the first module called by CAMRAD AUTO. The user sets the rotor design parameters, their bounds, and any desired constraints. This work has used 15 free parameters and 10 fixed parameters to fully define the outer mold line of each rotor design. The fixed parameters are documented in Table 4 and the free parameters are reported in Table 5. The framework handles these 15 free design parameters well and will be expanded in the future to include rotor radius, airfoil camber, and the number of blades per

rotor. Adding in a bi-linear camber distribution, rotor radius, and the number of blades per rotor would increase the number of free-parameters to 21.

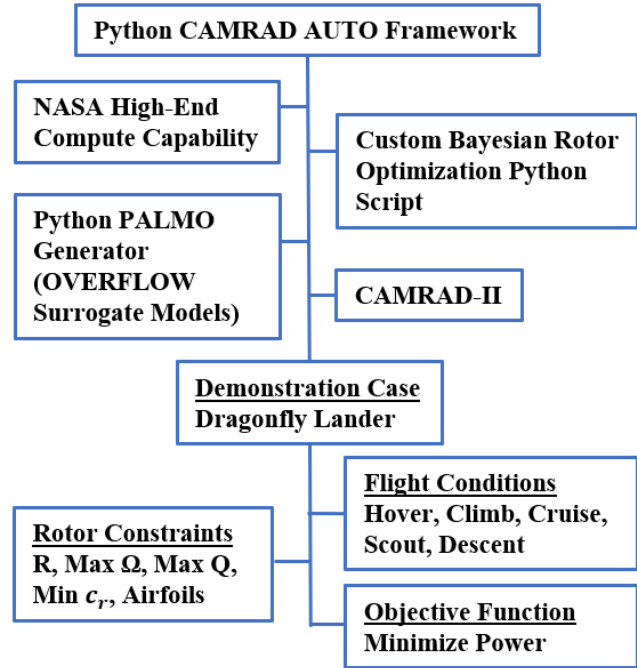


Figure 13. Python Framework for Rotor Optimization.

Table 4. Fixed Design Parameters.

Parameter	Description
x1, y1, z1	Taper, Twist, and Thickness Control Point 1, fixed at r/R = 10%
x3	Taper Control Point 3, fixed at r/R = 0.94
x4, y4, z4	Taper, Twist, and Thickness Control Point 4, fixed at r/R = 1
R	Rotor Radius
Nb	Number of Blades per Rotor
Airfoil % Camber	NACA 44### Series Airfoils currently used. (4 % camber)

Table 5. Free Rotor Design Parameter Ranges.

Parameter	Description (*CP = Control Point*)
1) x2	Taper CP 2, [0.2 - 0.8] % r/R
2) y2	Twist CP 2, [0.2 - 0.7] % r/R
3) y3	Twist CP 3, [0.25 - 0.85] % r/R
4) z2	Thickness CP 2, [0.2 - 0.8] % r/R
5) c1	Chord at x1, [0.03 - 0.1414] m
6) c2	Chord at x2, [0.03 - c1] m
7) c3	Chord at x3, [0.03 - c2] m
8) c4	Chord at x4, [0.03 - c3] m
9) tw1	Twist at y1, [5 - 30] deg
10) tw2	Twist at y2, [5 - tw1] deg
11) tw3	Twist at y3, [4 - tw2] deg
12) tw4	Twist at y4, [2 - 12] deg
13) th1	Thickness at z1, [15 - 24] %
14) th2	Thickness at z1, [12 - tw1] %
15) th4	Thickness at z, [6 - tw2] %

The first iteration of the optimization process requires an initial seeding, or sampling, of rotor designs to be evaluated in CAMRAD-II. The author’s ideal parametrization of this design space, which has anywhere from six to twelve discrete values spanning the range of each design parameter, results in several billion possible rotor combinations. Even after restricting each parameter to 5 values and incorporating constraints, there are still 55 million possible solutions. Various sampling methods have been applied in the rotorcraft literature such as uniform random sampling, grid sampling, Latin hypercube sampling (LHS), and Centroidal Voronoi tessellation (CVT). These methods have various advantages and disadvantages, but the main drawbacks in the context of this work are their inability to achieve one of the following objectives: 1) fill the solution space evenly, 2) sample non-rectangular regions, 3) use an arbitrary number of samples, or 4) apply progressive incremental sampling. Some of these methods, such as LHS and CVT, become increasingly computationally expensive with increases in dimensionality of the problem and the number of desired samples.

This work employs the Greedy Farthest Point (GreedyFP) algorithm to combat these issues and generate an initial seeding of this large parameter design space, Ref. [51]. This approach is well suited to achieve the sampling objectives with a low computational cost compared to other methods. It generates a user-defined number of samples filling the solution space with Euclidean distance between samples maximized. This process is used to generate the first batch of rotor design candidates to be analyzed in CAMRAD-II.

PALMO Generator

This module was developed to automate the process of C81 table generation using surrogate models from the PALMO database. The initial rotor design samples from the GreedyFP algorithm are fed into this module, which then calculates the Mach number and Reynolds number as a function of the chord distribution, blade radius, and rotor speed. The rotor blade C81 table discretization approach shown in Fig. 7 is used in this work to create thirteen unique C81 tables radially along the blade. This high-density parametrization of the airfoil performance look-up tables has been found to improve accuracy for variable-speed rotors with large radial variations in chord.

The surrogate models of the NACA 4-Series database are then queried. This work used feed-forward neural networks (FNNs) with three hidden layers each having one hundred neurons. Separate FNNs were used for each of the airfoil lift, drag, and pitching moment coefficients. The surrogate model evaluation is carried out using TensorFlow. Each rotor design requires approximately 10,000 queries of the FNNs since each of the thirteen C81 tables uses predictions for thirty-five angles-of-attack at seven discrete RPMs. These OVERFLOW based predictions span an angle-of-attack range from -20 to +20 degrees and are combined with experimental NACA 0012 data beyond that. As such, rotor design candidates operating at high lift coefficient should be

reviewed. The initial batch of 100 rotor designs, and each subsequent set, use roughly one million airfoil performance predictions queried from the PALMO surrogate models.

CAMRAD-II

The function evaluator implemented in this work is CAMRAD-II, Ref. [50], which analyzes the performance for each design. For this work, a 4-rotor CAMRAD-II model was created, representing one side for symmetric ‘semispan’ modeling of the overall Dragonfly lander. Figure 14 depicts a CFD flowfield of this 4-rotor configuration in a forward flight ‘cruise’ condition from prior work by the authors, Ref. [47]. The upper-to-lower and front-aft rotor-rotor interactions are exhibited by streamlines, which highlights the need to optimize the rotor design while directly accounting for effects of multirotor interference. Conducting the optimization using an isolated rotor or even a single coaxial rotor system would miss large contributions to the induced inflow of the various rotors, which also change as a function of flight condition.

To this end, several relevant flight conditions are analyzed in each CAMRAD-II job to optimize rotor performance over Dragonfly’s anticipated Titan mission. The upper and lower RPMs within a single coaxial rotor system are constrained to match each other, but the front coaxial RPM and aft coaxial RPM are trimmed to achieve pre-specified thrust targets for each flight condition. For wake modeling, the user can select uniform inflow, prescribed (rigid) wake, or free wake. This work has up to this point implemented both uniform inflow and prescribed wake. Studies were carried out to identify the most robust relaxation factors, convergence tolerance criteria, and iteration count for the various parts of the solution.

This work has used up to four NAS-HECC Pleiades Broadwell nodes, each having 28 physical cores. A timeout is set on each case to kill any evaluations that have not converged in a reasonable amount of time. A fifteen-minute limit was implemented for uniform inflow and a 2.5-hour limit for prescribed-wake calculations. Any non-converged cases are discarded. For converged cases, several rotor performance metrics of interest are collected to be used later as objective function values in the optimization.




Figure 14. Dragonfly Semispan 4-rotor Configuration (for illustration only) in Forward Flight Cruise.

ML Leveraged Bayesian Rotor Optimization

The background section of this work summarized rotor optimization studies from several groups. The most common issues cited were the curse of dimensionality preventing high-parameter design spaces, challenges in convergence for large problems with conventional optimization techniques such as gradient-based and evolutionary algorithms, and the still prohibitive cost of high-fidelity simulation data as inputs to these optimization frameworks.

This work applies a machine learning method called Bayesian optimization to combat these issues. This method, as its name implies, uses Bayesian statistics based on probability theory, Ref. [52]. Bayesian optimization is well suited for problems where each evaluation of the objective function is very expensive, e.g., mid- to high-fidelity simulation data. This type of optimization is especially suitable for finding a global optimum in a large parameter design space with a small number of function evaluations and is also robust against stochastic (noisy) training data. Bayesian optimization is a gradient-free method, which allows it to avoid converging to local optima. Although Bayesian optimization is typically used for hyperparameter tuning of neural networks, these characteristics make it extremely well suited for an applied engineering problem such as rotor design optimization using a numerical function evaluator such as CAMRAD-II.

Although many advanced implementations exist, a high-level simple implementation of Bayesian optimization will be described here in the context of rotor optimization. The process starts by training a probabilistic surrogate model such as GPR using an initial set of function evaluations, i.e., training data. This GPR model predicts the value of an objective function given a set of inputs. The objective function could be the rotor power requirement, and the inputs to the GPR model are the rotor design parameters. Step two evaluates this surrogate model over the full range of possible rotor design solutions, with each evaluation likely being orders of magnitude less computationally expensive than directly carrying out the function evaluation. This largely depends on the tool used for function evaluations and is well suited for mid- to high-fidelity methods such as prescribed-wake, free-wake, viscous vortex-particle methods, and CFD. The GPR model predicts the power requirement for each possible rotor design along with an estimate of the uncertainty of that prediction, i.e., the variance. An acquisition function is then used to determine the most likely rotor design in the solution space to yield an improvement on the objective function. One of the most suitable acquisition functions for this work is Expected Improvement (EI), which is shown in Equations 1-2 for the case of minimizing the objective function.

$$(1) \quad EI(x) = (y_{best} - \mu(x)) * \Phi(z) + \sigma(x) * \phi(z)$$

$$(2) \quad z = \frac{(y_{best} - \mu(x) - \zeta)}{\sigma(x)}$$

Where:

$\mu(x)$ is the GPR predicted value,
 y_{best} is the current minimum power,
 $\Phi(z)$ is the cumulative distribution function,
 $\sigma(x)$ is the GPR predicted variance,
 $\phi(z)$ is the probability density function
 ζ is the explore – exploit parameter

The EI calculation in Equation 1 considers the combination of the predicted value and the variance associated with that value. Including the variance accounts for the inherent uncertainty in the GPR surrogate model and is used to strike a balance between exploitation and exploration of the optimization search space. The explore-exploit parameter, ζ , which is in the numerator of the z-score calculation in Equation 2, can be used to tune the optimizer to focus its search more heavily in one direction or the other. An exploitation focused search will select the next candidate based on where the GPR model has predicted a better solution to exist. An exploration focused search selects the next candidate where the GPR predictions have large uncertainty, which is typically where little data exists and is therefore further exploring the design space. Including a non-zero value for the explore-exploit parameter may slow convergence but increases the probability of finding the global optimum.

The optimization process in this study has implemented an initial sampling of 100 rotor designs. These rotor designs are analyzed in CAMRAD-II and then used to train the GPR model. A new set of 1,000,000 randomly generated rotor designs, adhering to any user-defined constraints, are then evaluated using the GPR model. The optimization then uses the EI acquisition function to identify the next batch of rotor designs to analyze in CAMRAD-II. This process continues until a user-defined number of iterations, a total number of CAMRAD-II function evaluations, or a convergence criterion on the objective function is met.

DRAGONFLY SURROGATE MODEL BASED ROTOR OPTIMIZATION

Optimization of the New Frontiers Dragonfly rotor system has been selected as a demonstration case to highlight the efficacy of this novel Bayesian optimization framework. A CAMRAD-II model with front and aft coaxial rotor pairs is used, which captures rotor-to-rotor interference via the calculation of mean induced velocity at each rotor. This 4-rotor setup represents one side of the Dragonfly lander, which is considered adequate for the optimization due to symmetry about the fuselage. Fuselage-on-rotor interactional aerodynamic effects are not captured in this effort but are deemed to be of a secondary importance for the objective of this work.

A set of flight conditions generated by the Dragonfly flight dynamics team representing a nominal ‘Leapfrog’ flight on Titan are used to evaluate rotor performance over the anticipated mission, Ref. [53]. The Leapfrog profile

represents taking off from an initial location, flying over a Titan dune to scout future landing locations, and then returning to a pre-determined new target landing area, Ref. [54]. The flight conditions used in CAMRAD-II performance evaluations are documented in Table 6.

Table 6. Dragonfly Nominal ‘Leapfrog’ Mission Flight Profile Analyzed in CAMRAD-II.

Condition	Description
Hover	Combination of initial takeoff and landing
Climb	Cruise-climb to gain altitude
Cruise	For traversing long distances
Scout	‘Scouting’ potential future sites of interest
Descent	Cruise-descent to come in for landing

Including this diverse set of flight conditions directly in the objective function evaluation in each CAMRAD-II job is a large increase in complexity over most rotor design optimization studies, which tend to optimize to either hover or a forward flight cruise condition. Although this work does capture several planned flight conditions on Titan, some of the unique aspects of Dragonfly’s operational envelope are not included in the optimization. Examples are the transition to powered flight and preparation for powered flight, which only occur during the initial Titan entry, descent, and landing process, Refs. [53, 55-56].

The Bayesian optimization starts with an initial sampling of the potential rotor design space. Ten fixed and fifteen free design parameters were reported in Tables 4-5, respectively. The GreedyFP algorithm is used to intelligently sample this massive solution space, which is well into the hundreds of millions of possible rotor design permutations. Even five values for each parameter yields 55 million designs. An initial batch of 100 samples are selected to evaluate in CAMRAD-II. The results from this first initialization are then used to train a GPR model to correlate the rotor design inputs with the objective function, which in this case is the mission-weighted rotor power requirement over the Leapfrog flight profile. A new sampling of 1,000,000 rotor designs is then randomly generated and evaluated with the trained GPR model. The EI acquisition function is then used to select the next batch of rotors to evaluate in CAMRAD-II. This process iterates with each batch of CAMRAD-II function evaluations being continuously added to the training dataset used to make the GPR model and evaluate the next batch of one million rotor designs.

This framework leverages the new OVERFLOW Machine Learning Airfoil Performance (PALMO) database to increase the accuracy of the optimization process. Surrogate models created from PALMO, which in this work are feed forward neural networks (FNNs), are used to create robust high-density C81 tables using the appropriate Reynolds number, Mach number, and airfoil thickness distributions for each new rotor design. The 100 rotor designs to be evaluated in CAMRAD-II on each iteration of the optimization require approximately 1,000,000 queries of the PALMO surrogate

models. This process requires a few minutes on a single CPU thread, which is driven by current constraints in coding implementation. Future effort will be made to parallelize this process and eliminate coding bottlenecks, which will reduce this step to a few seconds.

Results are presented for nine iterations of the optimization, which produces approximately 900 function evaluations. A few of the CAMRAD-II cases diverge or fail to reach the convergence criteria. These cases are discarded. Figure 15 shows the GPR training time versus the number of CAMRAD-II function evaluations being used in the training dataset. Each iteration of the optimization has a larger training set as the previous batch of CAMRAD-II rotor performance calculations are added.

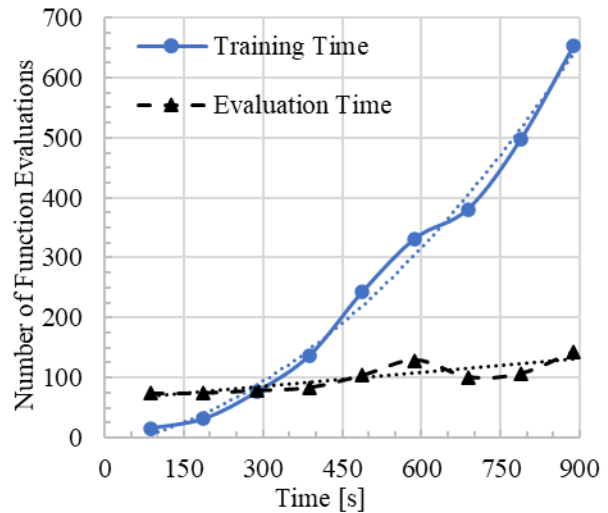


Figure 15. Bayesian Optimization GPR Training Time versus Evaluation Time.

The GPR is optimized with a random search over the hyperparameters and 250 iterations. The largest drawback of GPRs is their poor scalability, with a cubic relationship between computational cost and the number of trainable parameters. Figure 15 shows this rapid increase with the final GPR model training requiring approximately fifteen minutes. This process could be sped up by implementing GPU acceleration in TensorFlow, or by reducing the number of iterations in the GPR hyperparameter optimization. Advanced implementations of Bayesian optimization have even been carried out using neural networks, which is another way to reduce this time requirement, Ref. [57]. The evaluation time in Fig. 15 tracks the amount of time required for the GPR to evaluate the 1,000,000 new rotor designs. This number appears to increase slowly with the size of the training dataset.

Bayesian optimization was carried out for three values of the explore-exploit tunable parameter that was introduced in Equation 2. Figure 16 shows the optimization with an explore-exploit value of 0.01. The x-axis tracks the number of function evaluations, or actual rotor designs analyzed in CAMRAD-II. The y-axis is reporting a few metrics: 1)

current best is the rotor with the lowest power requirement up to that point in the optimization, 2) predicted best is the minimum predicted power in each set of 1,000,000 new rotor designs evaluated using the GPR model, and 3) maximum Expected Improvement (Max. EI) is the maximum value of the EI acquisition function from each set of 1,000,000 rotor designs. The maximum EI after the first initialization step is observed to be large, with the predicted best sitting very far beneath the current best. The optimization targets the regions of highest EI for the next set of CAMRAD-II evaluations, and the results of the first iteration show the current best and predicted best moving closer together.

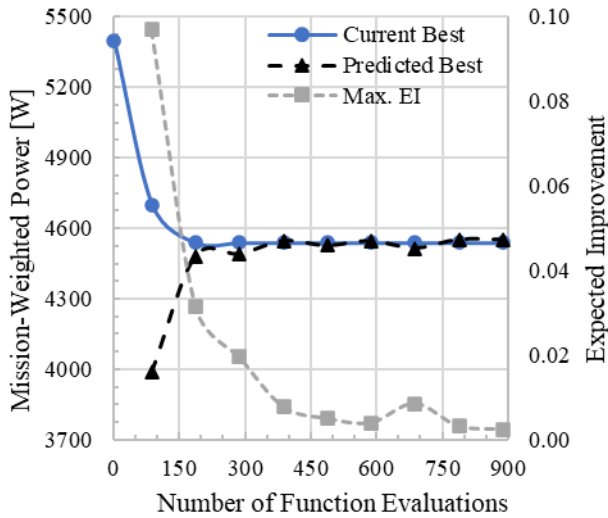


Figure 16. Bayesian Rotor Optimization to Minimize Mission Power, Explore-Exploit=0.01.

The minimum mission-weighted rotor power has already improved even after the initial seeding of the high-dimensionality parameter space. The left-most ‘Current Best’ value corresponds to the current Dragonfly rotor design. It should be noted, however, that this optimization does not yet fully account for all relevant Dragonfly rotor design constraints. Additionally, the current Dragonfly rotor was created using a different set of flight conditions relevant at the time. The minimum mission power continues to decrease as the optimization continues, with the first iteration of the Bayesian optimizer achieving a still further reduced mission power compared to the best value from the GreedyFP initial seeding. The Max. EI on each iteration continues to decrease as the solution space is further explored and the GPR model is improved.

Figures 17 and 18 have the same information plotted, but for explore-exploit values of 0.05 and 0.1, respectively. Similar observations can be made, with the one very noticeable difference being the faster reduction in the EI calculation. This happens because the optimizer is prioritizing rotor designs where the GPR model has higher uncertainty. By adding training data from those regions of high uncertainty, the subsequent GPR model uncertainty is reduced.




Figure 17. Bayesian Optimization, Explore-Exploit=0.05.

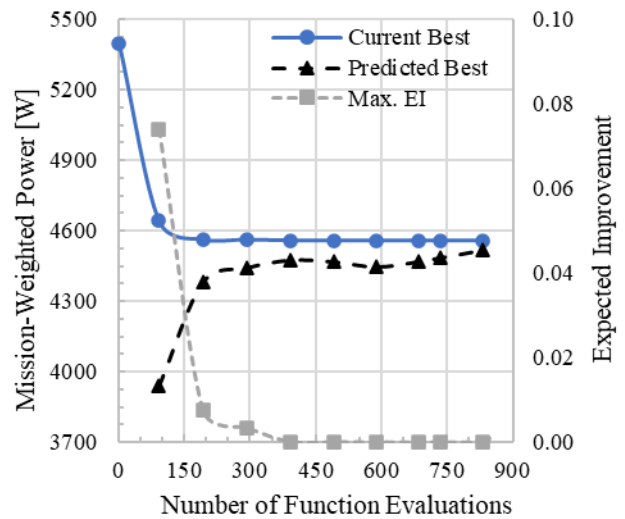


Figure 18. Bayesian Optimization, Explore-Exploit=0.10.

These three optimizations converged to approximately the same objective function value for the rotor power requirement weighted over Dragonfly’s planned Leapfrog flight mission. Each optimization used approximately 900 CAMRAD-II function evaluations, which then trained the GPR surrogate models that collectively analyzed about 30,000,000 possible rotor designs spanning the design space.

Although the three different optimizations resulted in a minimum power within one percent of each other, the three best rotor designs have slightly different chord distributions. The twist and thickness distributions, and the control points, converged to similar values. This indicates that the high-dimensionality design space in non-convex and may yield multiple varied solutions with near-optimal power. These various optimal designs could then be selected from to meet further design constraints such as rotor blade 1st flap frequency, mass, inertia, torque requirement, etc.

These designs do explore large deviations as compared to the existing Dragonfly rotor designs. Bi-linear chord and thickness distributions along with tip taper, for example, are increases in complexity over the baseline. As previously mentioned, however, this rotor aerodynamic optimization does not yet fully capture all relevant constraints that have been used to arrive at the current Dragonfly rotor design. As such, some of the ‘optimal’ solutions generated in this work would likely not be suitable rotor design candidates for Dragonfly. Still, there likely exists some margin of improvement that satisfies all design constraints.

Table 7 tracks how the fifteen free design parameters varied over one of the design optimizations. These data are from the optimization that used an explore-exploit parameter value of 0.05. Five incrementally better rotor designs, i.e., minimum power requirement, were generated throughout the optimization process. Some of the parameters converged quickly to a final value while others did not. This suggests that the application of methodologies such as principal component analysis (PCA) could be used to further target regions of the large design space that would most likely yield improved solutions. Even though this optimization evaluated nearly 30,000,000 rotor designs, that is still only a fraction of the possible permutations of the 15-parameter design space.

Table 7. Best Rotor Designs Generated through Successive Iterations of the Optimization.

Param.	Bayesian Optimization Iteration Number				
	#1	#2	#3	#4	#5
x2 (r/R)	0.225	0.200	0.200	0.250	0.400
y2 (r/R)	0.500	0.600	0.350	0.425	0.350
y3 (r/R)	0.650	0.800	0.450	0.500	0.425
z2 (r/R)	0.500	0.250	0.275	0.300	0.275
c1 (m)	0.086	0.036	0.131	0.126	0.080
c2 (m)	0.079	0.032	0.101	0.080	0.056
c3 (m)	0.047	0.031	0.041	0.034	0.030
c4 (m)	0.041	0.030	0.031	0.030	0.030
tw1 (deg)	27.76	29.55	29.82	29.78	27.93
tw2 (deg)	25.48	15.42	29.45	23.98	26.96
tw3 (deg)	22.98	14.60	18.28	18.76	20.90
tw4 (deg)	10.39	6.69	11.49	9.48	9.63
th1 (%)	22.0	16.0	19.0	22.0	22.0
th2 (%)	14.0	13.0	12.0	13.0	12.0
th4 (%)	12.0	6.0	6.0	6.0	6.0
P (W)	4,810	4,573	4,569	4,559	4,531

This work used the 28-core Broadwell nodes on the NASA-HECC Pleiades supercomputer. Optimization studies were carried out using four compute nodes, or 114 CPUs, and required approximately 5.5 hours for CAMRAD-II evaluations with uniform inflow and 27.8 hours using the CAMRAD-II prescribed wake model. The prescribed wake model implementation is roughly 3,200 CPU-hours, which could be completed in roughly 2-4 days using a powerful desktop workstation, i.e., with a 32-64 core CPU.

Rotor design numbers 2, 3, and 4 from Table 7 are shown in Figs. 19-21 to highlight the differences between various solutions found in the optimization. All three of these rotor designs have approximately the same power requirement over the Dragonfly Leapfrog mission, yet with large deviations in chord, twist, and thickness distributions. Figure 19, which shows Rotor Design #2, exhibits a chord distribution that is likely too low to meet the structural constraints necessary for rotor blade frequency de-tuning. Rotor Design #3 is displayed in Fig. 20 with a large and near-constant inboard root twist. Such a design would likely be susceptible to rotor stall in descent type conditions. Additional constraints on the value of lift coefficient over the disk may rule out this design.

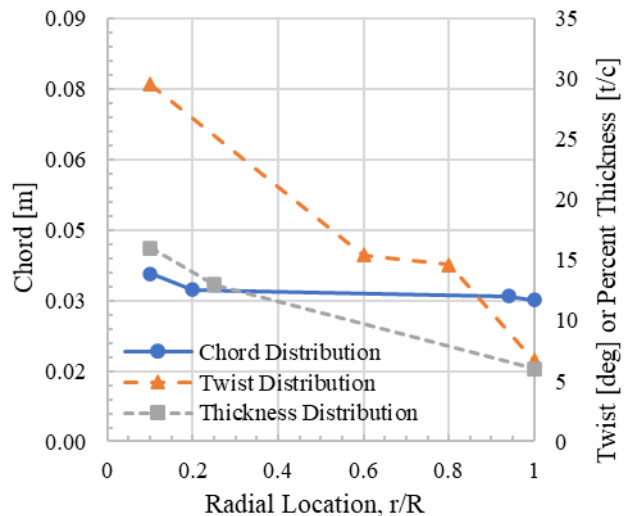


Figure 19. Rotor Design #2 with Low Solidity.




Figure 20. Rotor Design #3 with High Inboard Twist.

Rotor Design #4 is displayed in Fig. 21 with what seems to be a sensible chord, twist, and thickness distribution. Such a rotor design could be a strong candidate meeting Dragonfly design constraints while also achieving large power reductions. Still, further verification checks on the rotor

design structural requirements and other constraints are warranted before deeming a design as satisfactory.

Adding additional constraints inside the acquisition function of the EI calculation is another way to further constrain the problem and more fully evaluate the full solution space. Creating additional GPR models, for example, to evaluate torque, mean lift coefficient, and trim RPMs for each of the one million new candidate rotor designs could be a means to eliminate designs that do not adhere to the system-level design requirements. These additional constraints could either be used to discard rotor designs falling under some threshold or penalize the EI calculation in the acquisition function.

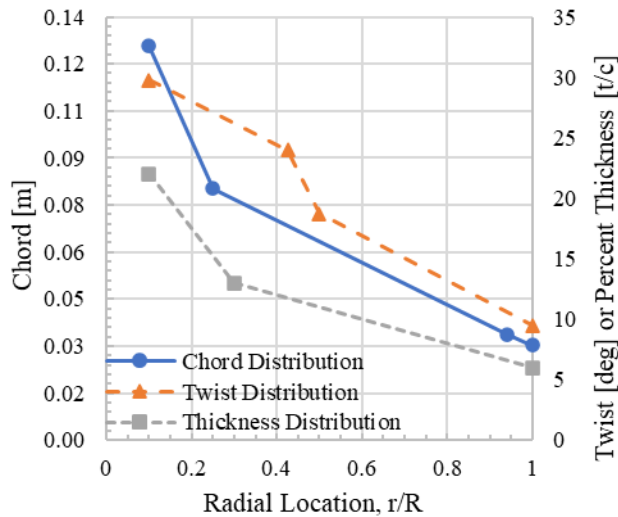


Figure 21. Rotor Design #4, a Potentially Strong Candidate for Optimum Rotor Design.

CONCLUSIONS

This work presented a novel framework for multi-tiered surrogate-model based rotor optimization. The methodology developed was applied in the context of rotor aerodynamic optimization for the New Frontiers Mission Dragonfly lander, but the approach is general and can be applied to both terrestrial and extraterrestrial applications on other planetary bodies of interest.

The OVERFLOW Machine Learning Airfoil Performance (PALMO) database was introduced, which is an airfoil database being developed at NASA Ames Research Center (ARC). The AIAA Surrogates Modeling Technical Committee is planning to use PALMO as a benchmark dataset to share with the international aerospace community. PALMO currently consists of about 62,000 high-order accurate OVERFLOW airfoil simulations. These simulations are a full parametrization of the NACA 4-Series airfoil family spanning the full range of Reynolds number, Mach number, and angle-of-attack relevant to the rotorcraft conceptual designer. Example surrogate models, feed forward neural

networks (FNNs), derived from the database were shown to have predictive accuracy on the order of 0.01 for lift coefficient, 0.002 for drag coefficient, and 0.003 for pitching moment coefficient for a five-percent test holdout. These surrogate model uncertainties are much less than the downstream 15-20% reductions in mission-weighted rotor power requirement. Still, a final rotor design would benefit from the actual creation of OVERFLOW airfoil tables at the specified conditions for a confirmation of the results.

These PALMO surrogate models were built into a new Bayesian rotor optimization framework under development at NASA ARC. PALMO surrogate models were successfully used to generate high-order accurate based C81 airfoil performance tables with the appropriate Reynolds number, Mach number, and airfoil thickness distributions. Each rotor design analyzed in CAMRAD-II queried the PALMO surrogate models about 10,000 times.

The optimization framework leverages methods from the fields of machine learning and image processing to enable massive computational cost reductions in the optimization process. Rotor design optimization was carried out with 15 free parameters defining tri-linear rotor twist, bi-linear rotor taper (with an additional tip-taper), and a bi-linear airfoil thickness distribution. Constraints such as monotonically decreasing airfoil thickness and chord were directly implemented into the rotor design generation process, which limited this solution space to about 55 million possible rotor designs. Since this is still an intractable number to evaluate directly, the Greedy Farthest Point (GreedyFP) algorithm was implemented for intelligent initialization of the design space.

Rotor optimization was carried out using a 4-rotor CAMRAD-II model to calculate the mission-weighted rotor power requirement over Dragonfly’s anticipated Leapfrog flight profile. Three studies were conducted with each optimization analyzing about 900 rotor designs in CAMRAD-II across the various flight conditions and evaluating approximately 10,000,000 rotor designs using the GPR surrogate models trained from the CAMRAD-II output. The three optimizations converged to approximately the same value of minimum mission-weighted power requirement, which was about 15% lower than the baseline, but with slight variations in the resulting rotor designs. This indicates that the high-parameter design space is non-convex and multi-modal, meaning there are likely multiple different designs resulting in a near-optimal power requirement. This is advantageous for when additional constraints or objectives are incorporated in the future. Although this work identified solutions with as much as 15% lower power than the baseline, constraints on the rotor blade flap frequency and mean lift coefficient were not considered in the analysis. These constraints have very much steered the Dragonfly rotor design to date and must still be applied to the optimization to yield suitable candidate rotors relevant for flight on Titan. Furthermore, this optimization used a different set of flight conditions than what the baseline rotor was designed to.

This work demonstrated surrogate-model based rotor design optimization using the newly developed PALMO database. Bayesian rotor aerodynamic optimization was carried out for the Dragonfly lander evaluating roughly 30 million rotor designs. CAMRAD-II was used as the function evaluator, analyzing roughly 2,700 rotor designs in a 4-rotor simulation accounting for rotor-rotor interference in multiple anticipated Dragonfly flight conditions. This process used 26 million evaluations of the PALMO surrogate models. Each rotor optimization study was conducted on 112 CPUs in 5.5 hours using uniform inflow, and 27.8 hours using prescribed wake, which is a massive reduction in computational cost compared to past efforts.

This work posits that Bayesian optimization is well-suited for efficiently finding global optimum design solutions in non-convex and high-dimensionality parameter spaces. The prescribed wake optimization could be completed on a powerful desktop workstation in roughly three days.

Author contact: Dr. Jason K. Cornelius
jason.k.cornelius@nasa.gov

FUTURE WORK

Future work will apply additional constraints directly into the expected improvement calculation, such as rotor torque, mean lift coefficient, and rotor blade structural properties, e.g., mass, inertia, and structural modes. This will be achieved with implementation of the multi-objective parallel expected improvement acquisition function. Principal component analysis will also be explored to traverse the high-dimensionality solution space more efficiently. Future efforts may investigate other types of ML such as reinforcement learning, but these methods have found much less application to date in the rotorcraft field. Lastly, the parametrization of the rotor design using Bezier curves will be explored.

ACKNOWLEDGMENTS

This work was funded by the NASA New Frontiers Dragonfly Mission. The authors would like to thank the NASA Advanced Supercomputing Division for their support in using the High-End Compute Capability for this work and the NASA Science Mission Directorate for the CPU and GPU allocations provided to support the Dragonfly program. The authors also thank Dr. Wayne Johnson, Stephen Wright, Michael Radotich, and Chris Silva for their support in learning CAMRAD-II. Many thanks also to Dr. Nicholas Peters, Tove Ågren, Witold Koning, Kristen Kallstrom, and Carlos Pereyra for countless discussions (lessons) in machine learning and sharing OVERFLOW best practices. The authors are indebted to Ethan Romander for his unwavering support in pushing the limits of high-performance computing. Lastly, the authors would like to acknowledge Larry Young and Dr. William Warmbrodt for their support of this work.

REFERENCES

1. LePage, A. J., "The Mysteries of Titan," *The Space Review*, Nov. 2010.
<https://www.thespacereview.com/article/1722/1>
2. Lorenz, R. D., Turtle, E. P., Barnes, J. W., Trainer, M. G., Adams, D. S., Hibbard, K. E., Sheldon, C. Z., Zacny, K., Peplowski, P. N., Lawrence, D. J., Ravine, M. A., McGee, T. G., Sotzen, K. S., MacKenzie, S. M., Langelaan, J. W., Schmitz, S., Wolfarth, L. S., and P. D. Bedini, "Dragonfly: a rotorcraft lander concept for scientific exploration at Titan," *Johns Hopkins APL Technical Digest* Vol. 34, No. 3, pp. 374-387, 2018.
https://dragonfly.jhuapl.edu/News-and-Resources/docs/34_03-Lorenz.pdf
3. "Why Titan?" The Johns Hopkins University Applied Physics Laboratory.
<https://dragonfly.jhuapl.edu/Why-Titan/>
4. Cornelius, J., Schmitz, S., "Coaxial Rotor CFD Validation and ML Surrogate Model Generation," VFS Aeromechanics Specialists' Conference, Santa Clara, CA, Feb. 2024.
https://rotorcraft.arc.nasa.gov/Publications/files/1686_Cornelius_Final_012324.pdf
5. Patt, D., Youngren, H., "Improving Blade-Element Design Methods for High Speed Propellers," AHS Specialists' Conference on Aeromechanics, San Francisco, CA, Jan. 2010.
6. Cornelius, J., Schmitz, S., "Massive Graphical Processing Unit Parallelization for Multirotor Computational Fluid Dynamics," *AIAA Journal of Aircraft*, Published Online: Jul. 2023.
<https://doi.org/10.2514/1.C037356>
7. Selig, M., "UIUC Airfoils Coordinate Database,"
https://m-selig.ae.illinois.edu/ads/coord_database.html
8. Shalul, H., Govindarajan, B., Sridharan, A., Singh, R., "Blade Shape Optimization of Rotor Using Neural Networks," VFS Forum 79, May 2023.
<https://doi.org/10.4050/F-0079-2023-18006>
9. Sridharan, A., Sinsay, J., "Accelerating Aerodynamic Design of Rotors Using a Multi-Fidelity Approach in TORC: Tool for Optimization of Rotorcraft Concepts," AIAA Aviation Forum, AIAA 2023-4305, June 2023.
<https://doi.org/10.2514/6.2023-4305>
10. Ruh, M., Liu, X., Yu, R., Hwang, J., "Airfoil Shape Parametrization Using Reconstruction-Error-Minimizing Generative Adversarial Networks," AIAA Aviation Forum, AIAA-2023-3722, June 2023.
<https://doi.org/10.2514/6.2023-3722>
11. Paternostro, N., Marepally, K., Anand, A., Lee, B., Baeder, J., "Application of CFD-Trained Neural Networks on the Rotorcraft Airfoil Design Process," VFS Autonomous VTOL Technical Meeting, Mesa, AZ, Jan. 2023.
12. Marepally, K., Paternostro, N., Lee, B., Baeder, J., "Data Puncturing and Training Strategies for Cost-Efficient Surrogate Modeling of Airfoil Aerodynamics," AIAA SciTech Forum, AIAA-2023-2042, Jan. 2023.
<https://doi.org/10.2514/6.2023-2042>

13. Du, X., He, P., Martins, J., “Rapid Airfoil Design Optimization via Neural Networks-based Parametrization and Surrogate Modeling,” *Aerospace Science and Technology*, Vol. 113, Jun. 2021. <https://doi.org/10.1016/j.ast.2021.106701>
14. Li, J., Du, X., and Martins, J., “Machine Learning in Aerodynamic Shape Optimization,” *Progress in Aerospace Sciences*, Vol. 134, Oct. 2022. <https://doi.org/10.1016/j.paerosci.2022.100849>
15. Kulfan, B., “Universal Parametric Geometry Representation Method,” *AIAA Journal of Aircraft*, Vol. 45, No. 1, Jan. 2008. <https://doi.org/10.2514/1.29958>
16. Lu, X., Huang, J., Song, L., and Li, J., “An Improved Geometric Parameter Airfoil Parameterization Method,” *Aerospace Science and Technology*, Vol. 78, pg. 241-247, 2018. <https://doi.org/10.1016/j.ast.2018.04.025>
17. Rajnarayan, D., Ning, A., and Mehr, J. A., “Universal airfoil parametrization using B-splines,” *Applied Aerodynamics Conference*, AIAA 2018-3949, Jun. 2018. <https://doi.org/10.2514/6.2018-3949>
18. Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., Mavriplis, D., “CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences,” NASA/CR-2014-218178.
19. Gray, A. C., Martins, J. R. R. A., “A Proposed Benchmark Model for Practical Aeroelastic Optimization of Aircraft Wings,” *AIAA Science and Technology Forum*, Jan. 2024. <https://doi.org/10.2514/6.2024-2775>.
20. Liu, X., Furrer, D., Kesters, J., Holmes, J., “Vision 2040: A Roadmap for Integrated, Multiscale Modeling and Simulation of Materials and Systems,” NASA/CR-2018-219771.
21. Wissink, A., Hariharan, N., “An Overview of the CREATE Applied Surrogates Institute,” *AIAA Science and Technology Forum*, Orlando, FL, Jan. 2024. <http://doi.org/10.2514/6.2024-0011>
22. Collins, K. B., Sankar, L. N., Mavris, D. N., “Application of Low- and High-Fidelity Simulation Tools to Helicopter Rotor Blade Optimization,” *Journal of the American Helicopter Society*, Vol. 58, No. 4, Oct. 2013. <https://doi.org/10.4050/JAHS.58.042003>
23. Lim, J. W., “Consideration of Structural Constraints in Passive Rotor Blade Design for Improved Performance,” *The Aeronautical Journal*, Vol. 120, No. 1232, Oct. 2016. <https://doi.org/10.1017/aer.2016.77>
24. Lim, J., Shin, S., Kee, Y., “Optimization of Rotor Structural Design in Compound Rotorcraft with Lift Offset,” *Journal of the American Helicopter Society*, Vol. 61, No. 1, Jan. 2016. <https://doi.org/10.4050/JAHS.61.012005>
25. Hersey, S., Sridharan, A., Celi, R., “Multiobjective Performance Optimization of a Coaxial Compound Rotorcraft Configuration,” *Journal of Aircraft*, Vol. 54, No. 4, Jul. 2017. <https://doi.org/10.2514/1.C033999>
26. Sridharan, A., Govindarajan, B., “A MultiDisciplinary Optimization Approach for Sizing Vertical Lift Aircraft,” *Journal of the American Helicopter Society*, Vol. 67, No. 2, Apr. 2022. <https://doi.org/10.4050/JAHS.67.022004>
27. Lim, J. W., Allen, L. D., Haehnel, R. B., Dettwiller, I. D., “An Examination of Aerodynamic and Structural Loads for a Rotor Blade Optimized with Multi-Objective Genetic Algorithm,” *VFS Forum 77*, Virtual, May 2021. <https://doi.org/10.4050/F-0077-2021-16770>
28. Allen, L., Dettwiller, I., Haehnel, R., Lim, J., “Helicopter Rotor Blade Multiple-Section Optimization with Performance Considerations,” *VFS Forum 77*, Virtual, May 2021. <https://doi.org/10.4050/F-0077-2021-16733>
29. Koning, W. J. F., Perez, B. N., Cummings, H. V., Romander, E. A., Johnson, W., “ELISA: A Tool for Optimization of Rotor Hover Performance at Low Reynolds Number in the Mars Atmosphere,” *VFS Aeromechanics Specialists’ Conference*, Santa Clara, CA, Feb. 2024. https://ntrs.nasa.gov/api/citations/20240000766/downloads/1689_Koning_Final_011924.pdf
30. Balam, J., Aung, M., Golombek, M. P., “The Ingenuity Helicopter on the Perseverance Rover,” *Space Science Reviews*, Vol. 217, No. 56, 2021.
31. Dull, C., Wagner, L., Young, L., Johnson, W., “Hover and Forward Flight Performance Modeling of the Ingenuity Mars Helicopter,” *VFS Aeromechanics Specialists’ Conference*, San Jose, CA, Jan. 2022. https://rotorcraft.arc.nasa.gov/Publications/files/Cuyler_Dull_10-Jan-22_06-00-24.pdf
32. Koning, W. J. F., Johnson, W., Grip, H. F., “Improved Mars Helicopter Aerodynamic Rotor Model for Comprehensive Analyses,” *AIAA Journal*, Vol. 57, No. 9, Sept. 2019. <https://doi.org/10.2514/1.J058045>
33. Sinsay, J. D., Alonso, J. J., “A Heuristic Approach to Finding the Preferred Design Variable Parametrization for Optimization,” *AIAA Science and Technology Forum*, San Diego, CA, Jan. 2016. <https://doi.org/10.2514/6.2016-0415>
34. Wang, L., Diskin, B., Biedron, R. T., Nielsen, E. J., Bauchau, O. A., “High-Fidelity Multidisciplinary Sensitivity Analysis and Design Optimization for Rotorcraft Applications,” *AIAA Journal*, Vol. 57, No. 8, Aug. 2019. <https://doi.org/10.2514/1.J056587>
35. Sridharan, A., Sinsay, J. D., “Accelerating Aerodynamic Design of Rotors using a Multi-Fidelity Approach in TORC: Tool for Optimization of Rotorcraft Concepts,” *AIAA Aviation Forum*, San Diego, CA, Jun. 2023. <https://doi.org/10.2514/6.2023-4305>
36. Shalu, H., Govindarajan, B., Sridharan, A., Singh, R., “Blade Shape Optimization of Rotors using Neural Networks,” *VFS Forum 79*, West Palm Beach, FL, May 2023. <https://doi.org/10.4050/F-0079-2023-18006>
37. Peters, N., Silva, C., Ekaterinaris, J., “A Data-Driven Reduced Order Model for Rotor Optimization,” *Wind Energy Science Discussions*, Vol. 8, No. 7, Jul. 2023. <https://doi.org/10.5194/wes-8-1201-2023>
38. Anusonti-Inthra, P., “Data-Driven Modeling of Aerodynamic Loadings for Tiltrotor Pylon using Multi-Fidelity CFD Data,” *AIAA Science and Technology Forum*, National Harbor, MD, Jan. 2023. <https://doi.org/10.2514/6.2023-0234>
39. Peherstorfer, B., Willcox, K., Gunzburger, M., “Survey of Multifidelity Methods in Uncertainty Propagation,

- Inference, and Optimization, Society for Industrial and Applied Mathematics Review, Vol. 60, No. 3, 2018. <https://doi.org/10.1137/16M1082469>
40. Erhard, R. M., Alonso, J. J., “Multi-Fidelity Bayesian Optimization of a Coaxial Rotor for eVTOL Aircraft,” AIAA Science and Technology Forum, Orlando, FL, Jan. 2024. <https://doi.org/10.2514/6.2024-2506>
 41. Lorenz, R. D., Schmitz, S., Kinzel, M., “Prediction of Aerodynamically-triggered Condensation: Application to the Dragonfly Rotorcraft in Titan's Atmosphere,” *Aerospace Science and Technology*, Vol. 114, Jul. 2021. <https://doi.org/10.1016/j.ast.2021.106738>
 42. NASA Langley Research Center Transonic Dynamics Tunnel (TDT) Facility. <https://www.nasa.gov/directorates/armd/aetc/transonic-dynamics-tunnel-tdt-facility/>
 43. Cornelius, J., Schmitz, S., Palacios, J., Juliano, B., Heisler, R., “Rotor Performance Predictions for Urban Air Mobility: Single vs. Coaxial Rigid Rotors,” *Aerospace*, Vol. 11, No. 3, Mar. 2024. <https://doi.org/10.3390/aerospace11030244>
 44. Buckley, M., “Johns Hopkins APL Dragonfly Team Utilizes Unique NASA Facilities to Shape Its Innovative Titan-Bound Rotorcraft,” Johns Hopkins Applied Physics Laboratory, Oct. 2023. <https://www.jhuapl.edu/news/news-releases/231023-dragonfly-wind-tunnel>
 45. AFTGen Airfoil Table Generator, Software Package, Build: 0.10.2, 2023. <http://sukra-helitek.com/>
 46. OVERFLOW 2.3d, Overset Grid Computational Fluid Dynamics Flow Solver With Moving Body Capability (OVERFLOW), NASA Software Catalog, 2023. <https://software.nasa.gov/software/LAR-20095-1>
 47. Cornelius, J., “Mid-Fidelity Performance Analysis for Fixed-Pitch Speed-Controlled Multirotor Aircraft,” Dissertation in Aerospace Engineering, The Pennsylvania State University, Aug. 2023. <https://etda.libraries.psu.edu/catalog/20665joc5693>
 48. Loftin, L., Smith, H., “Aerodynamic Characteristics of 15 NACA Airfoil Sections at Seven Reynolds Numbers from 700k to 9M,” NACA TN-1945, Oct. 1949. <https://ntrs.nasa.gov/citations/19930082618>
 49. Spalart, P., Allmaras, S., “A One-Equation Turbulence Model for Aerodynamic Flows,” AIAA 30th Aerospace Sciences Meeting and Exhibit, Jan. 1992. <https://doi.org/10.2514/6.1992-439>
 50. Johnson, W., “CAMRAD-II Theory,” Johnson Aeronautics, 2022. <http://johnson-aeronautics.com/CAMRADII.html>
 51. Kamath, C., “Intelligent Sampling for Surrogate Modeling, Hyperparameter Optimization, and Data Analysis,” *Machine Learning with Applications*, Vol. 9, 2022. <https://doi.org/10.1016/j.mlwa.2022.100373>
 52. Frazier, P. I., “A Tutorial on Bayesian Optimization,” arXiv: 1807.02811, Jul. 2018. <https://doi.org/10.48550/arXiv.1807.02811>
 53. Foust, R., Schilling, B., Rodovskiy, L., Jenkins, S., Sawyer, C., Villac, B., Boss, C., Lippe, C., Sutton, E., “Dragonfly First Flight – Preliminary Flight Dynamics Analysis,” AIAA Science and Technology Forum, Orlando, FL, Jan. 2024. <https://doi.org/10.2514/6.2024-2119>
 54. McGee, T. G., Adams, D. S., Hibbard, K. E., Turtle, E. P., Lorenz, R. D., Amzajerdian, F., Langelaan, J. W., “Guidance, Navigation, and Control for Exploration of Titan with the Dragonfly Rotorcraft Lander,” AIAA Science and Technology Forum, Kissimmee, FL, Jan. 2018. <https://arc.aiaa.org/doi/pdf/10.2514/6.2018-1330>
 55. Cornelius, J., Opazo, T., Schmitz, S., Langelaan, J., Villac, B., Adams, D., Rodovskiy, L., Young, L., “Dragonfly – Aerodynamics During Transition to Powered Flight,” VFS Forum 77, Virtual, May 2021. https://rotorcraft.arc.nasa.gov/Publications/files/77-2021-0264_Cornelius.pdf
 56. Shapiro, B. N., Boss, C., Putnam, Z. R., Villac, B., Pensado, A. R., Winski, R. G., Wright, M. J., “Dragonfly Preparation for Powered Flight: Lander Separation State Control to Ensure Successful Landing,” AIAA Science and Technology Forum, Orlando FL, Jan. 2024. <https://doi.org/10.2514/6.2024-2118>
 57. Snoek, J., Rippel, O., Swersky, K., Kiros, R., Satish, N., Sundaram, N., Patwary, M., Prabhat, Adams, R. P., “Scalable Bayesian Optimization Using Deep Neural Networks,” 32nd International Conference on Machine Learning, Lille, France, Jul. 2015. <https://dl.acm.org/doi/10.5555/3045118.3045349>