

Long-Range Mars Rotorcraft Design Optimization using Machine Learning

**Jason Cornelius, Nicholas Peters,
Jeremy Aires**
NASA Ames Research Center
Moffett Field, CA

Tove Ågren
Analytical Mechanics Associates
NASA Ames Research Center
Moffett Field, CA

**Darrell Nieves Lugo,
Anthony Comer, Zachary Miles**
Science and Technology Corporation
Moffett Field, CA

ABSTRACT

Simulation data consisting of multiple fidelity levels were generated using Graphical Processing Unit (GPU) resources on the NASA supercomputers. First, two large aerodynamic simulation databases were generated for geometric perturbations over a range of flight conditions for a hex-rotor bi-plane tailsitter aircraft. Results were visualized using the NASA Advanced Supercomputing Division's Hyperwall to improve the geometric design constraints. More than 3,000 full aircraft aerodynamic simulations were run using GPU enabled OVERFLOW with an actuator disk model to generate the airframe aerodynamic database. These simulations were completed in roughly 1.5 weeks on 32 GPU nodes using 128 NVIDIA V100 GPUs. Surrogate modeling techniques including Gaussian Process Regression (GPR), sparse GPRs, and a variety of Neural Networks (NNs) were used to create surrogate models to predict airfoil aerodynamic performance as well as airframe aerodynamics as a function of flight condition, airframe geometry, and rotor control input. These surrogate models were combined with additional Python modules predicting aircraft mass and inertia to generate another 3,000 aircraft simulations in CAMRAD-II in less than six hours using 240 Central Processing Unit (CPU) cores. Stability and control derivative matrices were obtained from the output to evaluate the open-loop characteristics. Lastly, the optimization framework was setup to allow simultaneous optimization of the aircraft and flight condition. The framework can now be used to optimize the aircraft for various objectives such as maximum range, endurance, or payload while satisfying constraints on controllability. This work brings higher-fidelity simulation data into the earlier stages of conceptual design, improving accuracy of the results and reducing the risk of missing critical design issues.

NOTATION

P	Rotor Power Requirement [W]
R	Rotor Blade Radius [m]
r	Radial Location [m]
α	Angle-of-Attack [deg]
C_T/σ	Blade Loading
c_l	Airfoil Lift Coefficient
c_d	Airfoil Drag Coefficient
c_m	Airfoil Pitching Moment Coefficient

LRMR	Long-Range Mars Rotorcraft
MSL	Mars Science Laboratory
MAE	Mean Absolute Error
MFM	Multi-fidelity Model
ML	Machine Learning
NN	Neural Network
ROAMX	Rotor Optimization for the Advancement of Mars eXploration
SMD	Science Mission Directorate
VTOL	Vertical Takeoff and Landing

ACRONYMS

ARC	NASA Ames Research Center
BEMT	Blade Element Momentum Theory
CFD	Computational Fluid Dynamics
CPU	Central Processing Unit
EI	Expected Improvement
ELISA	E Evolutionary aLgorithm for Iterative Studies of Aeromechanics
GPR	Gaussian Process Regression
GPU	Graphical Processing Unit
HECC	(NASA) High-End Compute Capability

INTRODUCTION

Recent NASA missions as part of the Science Mission Directorate (SMD) have developed vertical take-off and landing (VTOL) aircraft for exploration of other planetary bodies in our solar system. The Ingenuity Mars Helicopter was the first aircraft to fly on another world with a total of 72 flights over its 1.5 year mission, Ref. [1]. The New Frontiers Dragonfly lander is now under development as the next rotorcraft to be launched from Earth and will explore Saturn's largest moon Titan, Ref. [2]. Such rotorcraft explorers enable

science campaigns to be conducted across geographically diverse regions and in locations that rovers cannot easily access. Renderings of Ingenuity and Dragonfly are included in Figure 1.




Figure 1. The Ingenuity Mars Helicopter (top) and the Titan Dragonfly Lander (bottom).

Although these aerial explorers provide unique mission capabilities, their design and pre-flight qualification is a formidable task. The Titan atmosphere consists of cryogenic Nitrogen and Methane, and the atmospheric density is 440% that of Earth. On Mars, the thin atmosphere at 1-3% that of Earth leads to a low Reynolds number, high-Mach number operating regime. Although these conditions can be partially simulated in facilities at NASA Ames Research Center, the Jet Propulsion Laboratory, and the Johns Hopkins University

Applied Physics Laboratory, Refs. [3-5], getting accurate experimental flight data remains a challenge. The 25-foot Space Simulator at JPL was used to conduct hover testing for Ingenuity and select forward flight testing, while the NASA Langley Transonic Dynamics Tunnel (TDT) has been used to test Dragonfly's full-scale coaxial rotor system at roughly 75% of Titan's atmospheric density. These facilities do not, however, simultaneously recreate the true environment including atmospheric properties, temperature, and gravity. Due to these constraints, there is little opportunity for 'free-flight' testing of these aircraft until they are deployed on their target planetary bodies.

These challenges in experimental validation and verification result in a heavy reliance on modeling and simulation. Despite advancements in modern day supercomputing, the highest fidelity simulations are still far too computationally expensive to support the extensive analysis required for designing and understanding these aircraft. Engineers must thus use lower-fidelity models that could miss critical aspects of the associated physics and/or complex multi-disciplinary interactions.

The objective of this work is to leverage machine learning (ML) to bring multi-disciplinary, higher-fidelity simulation data forward in the conceptual design process. This has two main goals: 1) mitigate the chance of selecting aircraft design features with critical issues, and 2) increase confidence in the modeling and simulation results of the aircraft's predicted performance in its intended environment. This work develops a methodology for machine learning leveraged design optimization. The methodology is demonstrated here for the design of a Long-Range Mars Rotorcraft (LRMR) given the previously stated challenges of planetary aerial vehicle design. The ML methodology, however, is general and can be applied to a variety of aerospace design problems such as atmospheric entry vehicles, rockets, airplanes, and hypersonic vehicles, as well as other data-driven problems relevant to NASA.

OVERVIEW

This paper details a methodology for multidisciplinary multi-fidelity full aircraft design optimization. The methodology is a continuation from previous work by Cornelius and Schmitz, Ref. [6], that used a Bayesian optimization approach to find optimal rotor aerodynamic designs in a high dimensional non-convex design space. The methodology found optimal solutions with substantially reduced computational cost compared to the previous state-of-the-art methodology. These computational efficiency increases were attained using surrogate modeling to create




Figure 2. Machine Learning Leveraged Design Optimization Approach.

lower-order representations of OVERFLOW generated airfoil data and CAMRAD-II generated rotor performance data.

This work advances that methodology to include full aircraft optimization using varying fidelity simulation data from multiple disciplines in rotorcraft engineering (see Figure 2). Higher-fidelity aerodynamic forces and moments for the airframe, including rotor interactional aerodynamics, are first calculated using OVERFLOW with an actuator disk implementation. A database spanning the entire geometric design space and flight envelope is generated using an intelligent sampling technique referred to as the Greedy Farthest Point algorithm, Ref. [7]. A surrogate model of this OVERFLOW generated airframe aerodynamic data is then used in CAMRAD-II to calculate aircraft performance, vehicle trim, and stability and control derivative matrices as a function of geometric design parameters and the flight condition. This output is then built into a Bayesian optimization framework to enable optimization for any user-defined objective functions.

CONCEPTUAL DESIGN

The Long-Range Mars Rotorcraft

The successful flight campaign of the Ingenuity technology demonstrator paved the way for future aerial explorers to be used on Mars. As a technology demonstrator and add-on payload to the Perseverance Rover mission, Ingenuity was only 1.8 kg and had no payload capacity. The aircraft's maximum range was 705 meters (Flight 69), the maximum flight time was 170 seconds (Flight 12), and the maximum ground speed was 10 m/s (multiple flights), Ref. [8].

Past studies have analyzed an array of potential aircraft configurations for Mars exploration, and a few works have even been proposed as follow-on explorers to Ingenuity. A few of these configurations include a conventional single main rotor helicopter, tiltrotor, tailsitter, and multirotor. Each of these aircraft configurations come with unique challenges requiring extensive simulation and experimental testing. This work aims to enable advanced concept aircraft through the inclusion of higher-fidelity simulation data in the early stages of conceptual design. As such, an advanced vehicle configuration was selected to demonstrate the ML-based methodology being developed.

Given the desire for VTOL capability while simultaneously achieving good cruise efficiency for a long range, the quadrotor biplane tailsitter (QBiT) configuration was explored. QBiT aircraft have been extensively investigated by several groups for Earth applications in both the commercial drone industry and as a defense technology, Refs. [9-13], Figure 3.

An aircraft that could traverse a large range over multiple sorties would enable exploration of geographically diverse regions of Mars. Traversing a large range in a single flight creates opportunity for unique missions such as aerial exploration of otherwise inaccessible regions.



Figure 3. Quadrotor Biplane Tailsitter (QBiT).

Aircraft Sizing for Mars Entry Vehicles

The aircraft was constrained to fit within existing Mars atmospheric entry vehicles. The Mars Science Laboratory (MSL) entry vehicle's 4.5-meter diameter aeroshell was used in this work; a similar aeroshell carried Perseverance and Ingenuity, Ref. [14]. Figure 4 depicts the MSL aeroshell. A rotor sizing study was conducted to determine the maximum rotor radius for a given number of rotors, as shown in Figure 5. More rotors yield higher total rotor swept area decreasing power requirement, but also add mechanical complexity. To strike this balance and increase fault tolerance, a Hexrotor Biplane Tailsitter (HBiT) was chosen since it can still be controlled with one to two rotor or motor failures.

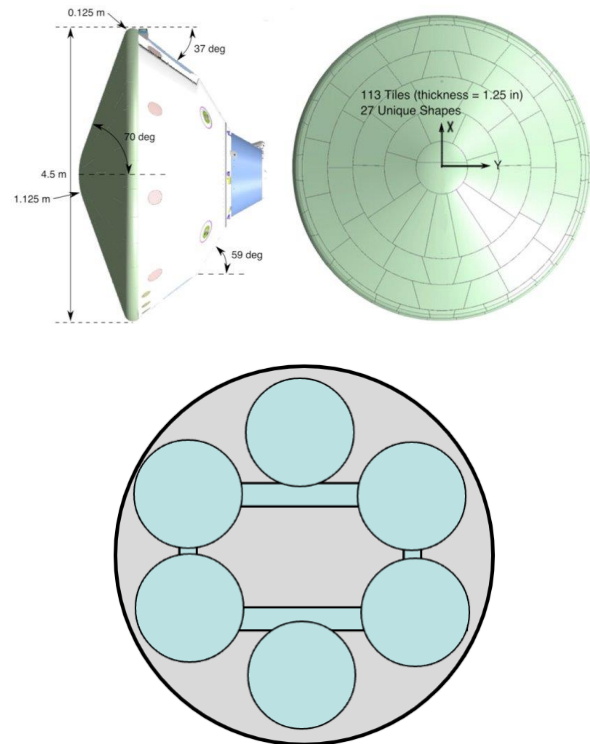


Figure 4. Mars Science Laboratory Entry Vehicle (top), HBiT Rotor Layout in Aeroshell (bottom).

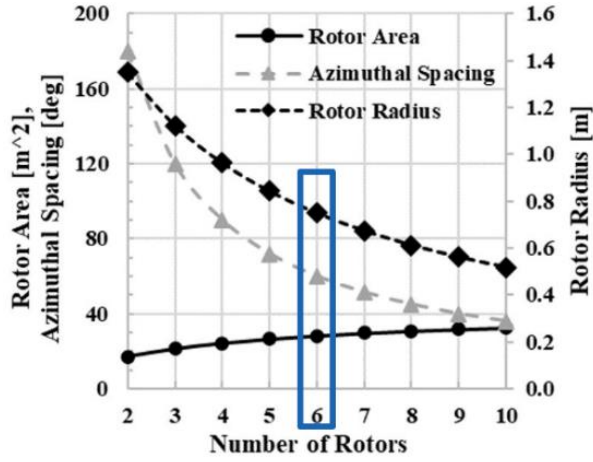


Figure 5. Rotor Disk Area Study.

Based on six rotors, preliminary rotor sizing calculations were conducted using rotor design information from the Ingenuity Mars Helicopter and continued Mars rotor development efforts by the NASA Rotor Optimization for the Advancement of Mars eXploration (ROAMX) team. Rotor solidity from 0.2 to 0.3 with blade-tip Mach numbers from 0.8 to 0.9 were explored. Mars airfoils and rotors optimized by the ROAMX team cover the region of interest for these rotors. OVERFLOW aerodynamic data were also obtained from the ROAMX team to inform the study. A rotor radius of 0.75 m was fixed given the entry vehicle constraints.

Table 1 reports the aircraft mass supported in hover (from momentum theory) for various combinations of blade-tip Mach number, solidity, and normalized rotor thrust coefficient (C_T/σ). An aircraft mass of 50-70 kg appears viable, and designs on the lighter end of that spectrum have substantial thrust margin remaining for control. The corresponding chord-based Reynolds numbers are reported at the 75% blade radial location in Table 2. The rotor design space has a Reynolds number range from 20,000 to 40,000, which has been extensively analyzed for the ROAMX project. This work leverages airfoil performance data for the ROAMX rotor generated using the Evolutionary aLgorithm for Iterative Studies of Aeromechanics (ELISA) optimization toolset, Refs. [15-17]. Table 3 reports the estimated hovering power for the entire aircraft using momentum theory. The power requirement in cruise is highly dependent on the vehicle trim and will be reported with higher fidelity analyses in a subsequent section.

Table 1. HBiT Hover Target Mass, Radius = 0.75m.

Aircraft Weight Supported [kg]		Tip Mach Number			C_T/σ $\sigma=0.3$
		0.80	0.85	0.90	
C_T	0.025	39.5	44.6	50.0	0.083
	0.030	47.4	53.5	60.0	0.100
	0.035	55.3	62.4	70.0	0.117
	0.040	63.2	71.4	80.0	0.133
	0.045	71.1	80.3	90.0	0.150

Table 2. Rotor Blade Reynolds Number at $r/R=0.75$.

Rotor $Re@r/R = 0.75$		Rotor Solidity		
		0.20	0.25	0.30
Tip Mach Number	0.600	18,176	22,720	27,265
	0.700	21,206	26,507	31,809
	0.800	24,235	30,294	36,353
	0.850	25,750	32,187	38,625
	0.900	27,265	34,081	40,897

Table 3. Estimated HBiT Hover Power Requirement for Various Blade-tip Mach Numbers.

Aircraft Power [kW]		Tip Mach Number		
		0.80	0.85	0.90
C_p	0.0020	2.25	2.70	3.21
	0.0040	4.50	5.40	6.41
	0.0060	6.75	8.10	9.62
	0.0080	9.01	10.80	12.82
	0.0085	11.26	13.50	16.03

Bounds for the airframe were also set based on the aeroshell geometry. An ellipsoid was placed in the center of the six rotors, between the wings, to house the electronics, flight battery, and scientific payload. Maximum values for the wingspan, wing chord, wing vertical spacing, fuselage semi-major and semi-minor axes were set by the aeroshell geometry. Minimum values were determined to keep the Reynolds number in a range suitable for analysis in OVERFLOW and a total aircraft lift relevant to the target mass of 50-75 kg. Additional constraints on fuselage minimum volume, wing aspect ratio, and fuselage aspect ratio were also applied.

Mass and Moment Estimation

Mass and moment of inertia estimates for the various aircraft designs are needed for higher fidelity analyses, including trim and vehicle stability and control. To this end, a custom routine was developed for analyzing the mass and moment of inertia properties of the individual components comprising the HBiT. This was done to facilitate design changes such as vehicle dimensions, rotor locations, and density of materials. The largest geometries consisted of extruded airfoil profiles for the wings and support structures as well as a prolate spheroid for the fuselage. Most of the internal components were assumed to have simple geometries such as hollow cylinders for spars, solid cylinders for the motors, thin disks for the rotors, a rectangular prism for the avionics, etc.

Material properties of structural components, wing skin, supporting spars, and fuselage components were assumed to be made of carbon fiber or aluminum materials. The thickness of structural elements was determined through classical statics estimates given the component's intended loading. For example, a wing spar was estimated as two cantilever beams with loading estimated from the vertical spar attachment joints.

Staying with common geometries benefits from having readily available mass and inertia equations available in many textbooks and online. For example, estimates for the prolate spheroid were adopted from an ellipsoid with minor modifications (see Equations 1-3).

$$M = \rho \frac{4}{3} a^2 b \pi \quad (1)$$

$$I_x = \frac{2}{5} M b^2 \quad (2)$$

$$I_{y,z} = \frac{1}{5} M (a^2 + b^2) \quad (3)$$

Where a and b are the major and minor axes lengths, respectively, for the prolate spheroid. ρ represents the material density.

Unlike the simpler geometries, the properties for the extruded airfoils were calculated using a generalized subroutine that accepted the chord-length-normalized coordinate files being used for the aerodynamic analyses. This routine scales the airfoil to the appropriate size, divides the profile into cells based on a resolution set by the user, and extrudes the cells to the appropriate span length before summing the weights and using the relative distances to determine mass and inertia properties.

Additionally, the subroutine has a parameter to specify a thickness to be used for shell profiles. This allows for more accurate calculations of the hollowed wings. Examples of the solid and shell distribution profiles for an Eppler E387 airfoil can be seen in Figures 6-7.

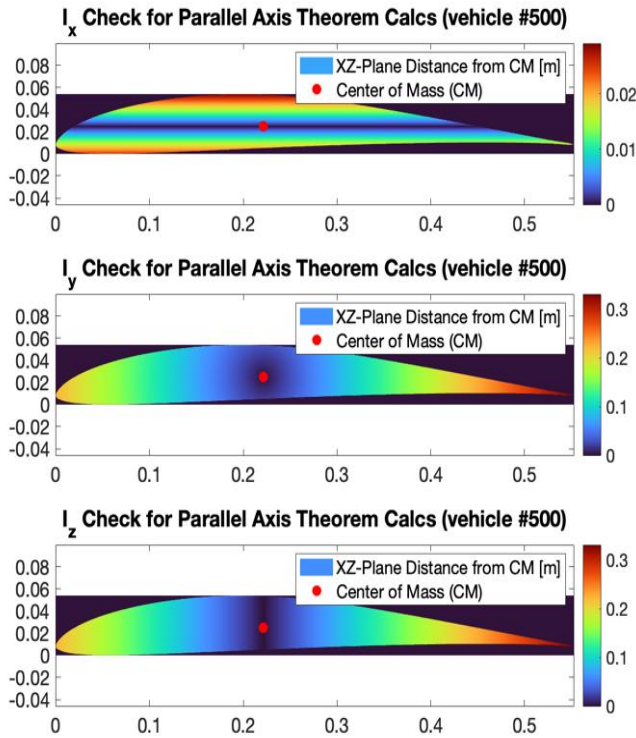


Figure 6. Mass Distribution Estimate for a Solid Airfoil.

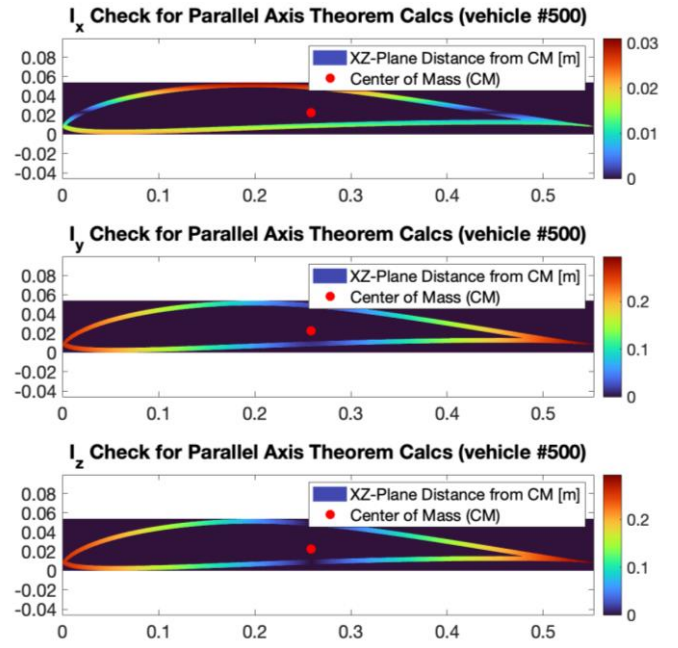


Figure 7. Mass Distribution Estimate for an Airfoil Shell.

The algorithm for mass and inertia estimation also included motor and battery sizing calculations that were constrained to a maximum vehicle mass of 75 kg. Using this target mass, motors were sized such that 80% power setting would provide sufficient thrust to maintain hovering flight in the event of two inoperative motors. This provided a 20% power margin for maneuvering the vehicle in the event of this worst-case-scenario failure situation. Based on this sizing constraint, the mass of the motor begins with an estimated motor power per kilogram (W/kg) of 175 being selected. This value provides an estimate for power needed to lift a specified vehicle mass. Using the 75 kg target mass, spread across four propellers (assuming the two-rotor failure case), this would require each rotor to support 18.75 kg to maintain hover. Using the 175 W/kg conversion, this yields an estimated 3,281 W/motor to sustain hover. Adjusting for the desired 80% throttle setting, this comes out to an adjusted 4,102 W/motor to sustain hover. The motor mass can then be determined using Equation 4 as follows:

$$\text{Motor Mass [kg]} = k \cdot \sqrt{T[N] \cdot P[W]} \quad (4)$$

Where k is a lightweight parameter, set to $k = 0.004$, responsible for estimating the mass of a motor given a known thrust T in Newtons (where $g = 3.71 \text{ m/s}^2$ on Mars) and power, P , is in Watts. Based on these values, this would yield an estimated motor mass of 2.14 kg/motor.

The final item to be sized was the battery. Weights and dimensions of Panasonic NCR2170-M cells were used while assuming a future improved power density of 400 Watt-hr/kg. Minimum voltage requirements for the motors were also used to determine whether a sufficient number of cells could be combined in series and parallel to meet the motor demands and targeted nine-minute operation time. If enough power could be supplied to meet the target operational time, the

remaining mass of the vehicle was left untouched to maximize payload capacity.

With each component individually calculated, the parallel axis theorem was used to find the overall moments of inertia for the aircraft. Due to the symmetric nature of the vehicle, and to save computational processing time, the contributions from misalignments of the components' principal axes were neglected.

AERODYNAMIC DATABASE MODELING

The OVERFLOW CFD Tool

All Computational Fluid Dynamics (CFD) simulations completed in this work used the NASA OVERFLOW solver. The OVERFLOW code is an overset, high-order structured grid-based compressible CFD code often leveraged in the rotorcraft community to obtain high-fidelity predictions for rotorcraft-based applications. In this study, two aerodynamic databases were generated. For the first database, only the fuselage, wing, and pylons were modeled within the CFD code, Figure 8. The generation of this geometry requires 21 overset near-body grids and approximately 30 million grid points. A total of 500 CFD simulations were completed for this preliminary database.




Figure 8. V1 Aerodynamic Database OVERFLOW Grid.

For the second database, the fuselage, both wings, the pylons, and all 6 rotors were modeled in the CFD simulation, Figure 9. The generation of this geometry required 33 overset near-body grids and approximately 57 million grid points. A total of 3,000 CFD simulations were completed. To model the rotors in the CFD solver, the recently implemented Blade Element Theory (BET) source term-based rotor disk model was leveraged to simulate all presented OVERFLOW-based CFD results. To compute rotor loads using the BET rotor disk model, the high-fidelity ROAMX airfoil performance lookup tables were used. The baseline ROAMX rotor was used in the simulations. Each rotor disk was modeled using 50 spanwise nodes and 36 azimuth nodes for the rotor's CFD mesh. A Cartesian, multiblock off-body mesh was then generated with a mesh refinement of 10% wing mean chord length and a refinement box that extended four rotor diameters downstream of the rotor's hub. Both the near-body rotor and off-body background Cartesian grids used second-order temporal and fourth-order spatial discretization with a constant global Courant–Friedrichs–Lewy (CFL) prescribed.

Simulations used third-order central differencing for the Euler terms and Beam-Warming for the viscous terms.




Figure 9. V2 Aerodynamic Database OVERFLOW Grid.

One notable contribution of this study is its investigation into deriving meaningful and highly generalized machine learning-based surrogate models from high-fidelity CFD simulation data. In developing such a surrogate model from high-fidelity overset curvilinear CFD simulation data, two key challenges arise.

First: Efficient automatic generation of parametrically defined, high-quality overset curvilinear grids. NASA's Chimera Grid Tools (CGT) code, a curvilinear grid generation tool, was used to create all the necessary grids for CFD simulations in this study. The CGT code is commonly employed for generating the grids of pylons, fuselages, and rotors for simulations in OVERFLOW. However, generating high-quality overset grids for the wings of the vehicle proved to be challenging, leading this study to utilize the recently implemented Bladegen program. Bladegen is an internal NASA grid generation tool designed to automatically create high-quality grids for wings and rotor blades. Although Bladegen is not currently available for public release, it will be included in the upcoming CGT 2.3 release.

Second: Developing surrogate models from high-fidelity CFD simulations requires a sufficient number of simulations. This study required several thousand CFD simulations to create a meaningful surrogate model. Traditionally, completing such a large number of three-dimensional high-fidelity simulations would be infeasible within a single study. However, recent developments in OVERFLOW that enable it to run on GPUs, combined with NASA's investment in GPU hardware, made it possible to complete all required simulations efficiently. Despite each simulation requiring tens of millions of cells and thousands of iterations, the use of GPUs allowed each simulation to be completed in roughly one hour. By leveraging 80 NVIDIA V100 GPUs at NASA, this study was able to complete approximately 400 full aircraft CFD simulations per day, thereby generating extensive, high-fidelity CFD databases for machine learning applications.

Model Validation

Prior to executing all 3,500 simulations, a preliminary validation effort was conducted. This involved performing a series of CFD simulations to confirm that the study accurately modeled the lifting surfaces. One of the challenges in this validation process is the scarcity of experimental data available for the low Reynolds number conditions examined in this study. Additionally, limited experimental data exists for the specific vehicle configuration considered. Fortunately, there were experimental measurements available for the lift coefficients of an Eppler E387 airfoil, Ref. [18]. While the focus of this validation effort was for a three-dimensional wing, given a sufficiently high wing aspect ratio, the mid-span lift coefficient should match the experimental measurements for the airfoil. Thus, this effort allowed the validation of the chosen mesh and solver settings employed in the research.

A series of isolated wing simulations were conducted at progressively higher wing aspect ratios. Each wing was simulated at an angle of attack of 5 degrees and a Reynolds number of approximately 100,000. The initial wing aspect ratio started at 3.5 and was increased until a converged solution for the loading at the wing's center was achieved. This converged solution was obtained at an aspect ratio of 14.2, as demonstrated in Figure 10. The final computed lift coefficient at the wing's center was 0.89, which lies well within the uncertainty bounds reported by the experimental data.




Figure 10. Effect of Increasing Wing Aspect Ratio on Spanwise Loading.

Aerodynamic Database V1

To create the OVERFLOW aerodynamic database, the airframe was parametrized using: 1) wingspan, 2) wing chord, 3) wing vertical separation, 4) fuselage semi-major axis, and 5) fuselage semi-minor axis.

The flight condition was parametrized by cruise velocity and vehicle pitch attitude. The Greedy Farthest Point (GreedyFP) space filling algorithm was used to generate 500 intelligently sampled OVERFLOW simulations of the airframe geometric design space and flight envelope. This database was generated using the NASA High-End Compute Capability's (HECC) GPU clusters and a sample flowfield visualization is shown in Figure 11 with lift to drag results in

Figure 12. Each OVERFLOW simulation used approximately 30 million grid points and required 40 minutes on a GPU compute node with four NVIDIA V100s. The database was generated in one weekend.

Aerodynamic Database V2

To create the second OVERFLOW aerodynamic database, the airframe and rotors were parametrized using: 1) wingspan, 2) wing chord, 3) wing vertical separation, 4) fuselage semi-major axis, 5) fuselage semi-minor axis, 6) rotor common collective input, and 7) rotor longitudinal pitch input. The two rotor control inputs fully defined the collective values for each of the six rotors. The flight condition was again parametrized by cruise velocity and vehicle pitch attitude, yielding a total of nine free parameters. The GreedyFP algorithm was again used to generate 3,000 OVERFLOW simulations spanning the airframe geometric design space, control input space, and flight envelope. This V2 aircraft database was generated using the NASA HECC GPU clusters and a sample flowfield visualization is shown in Figure 11. Each OVERFLOW simulation had about 57 million grid points and required 1 hour and 15 minutes on a GPU compute node with four NVIDIA V100s. The database was generated in roughly 1.5 weeks using 20 GPU nodes.




Figure 11. Aerodynamic Database Flowfield Visualizations, Left: V1, Right: V2.




Figure 12. Aerodynamic Database V1 Lift-to-Drag, Chord-based Reynolds Number Varies from 50k – 250k.

Two major outliers are observed in the lift to drag plot, and data pruning is discussed as a pre-processing step before training the surrogate models in the next section. Airframe geometry perturbations and flow visualizations were viewed using the NASA-HECC Hyperwall as seen in Figure 13.

MACHINE LEARNING

This section accounts for the process of deriving surrogate models of the airframe aerodynamics and ROAMX airfoil performance databases, respectively.

Airframe Aerodynamics Models

Given the aerodynamic databases, referred to as V1 and V2 and described in the previous section, the surrogate modeling task at hand was to predict the lift, drag and pitching moments imparted on the airframe, as a function of the geometrical vehicle design parameters and the flight condition. This amounts to seven total input features for the V1 database; the V2 database has nine input features including the two additional rotor control inputs. The lateral force as well as yaw and roll moments were excluded from the regression task, due to their small magnitude relative to the on-axis forces mostly due to vehicle symmetry for the considered flight conditions.

Gaussian Process Regression (GPR) is a popular choice for small to mid-sized data set surrogate modeling, owing to the ability to capture non-linear features, ease of implementation, and probabilistic representation allowing for intrinsic uncertainty quantification of the model predictions. Exact GPR models were initially explored for the airframe surrogate task. In the earlier stages of generating the V1 aerodynamic database, a preliminary study of the dependency of model prediction accuracy on the size of training data set was conducted.

Ultimately, these results guided the generation of the subsequent databases, providing insight into the trade-off between potential increase in accuracy and the cost of generating new samples. An exact GPR model, simultaneously predicting the three aerodynamic loads, was trained on an iteratively larger portion of the first 400 cases of the V1 database and evaluated on a test set of 50 held out datapoints. The Mean Absolute Error (MAE) on the test set as function of number of samples in the training dataset is depicted in Figure 14. The accuracy improvement is shown to level off after 350 samples, implying existence of a convergence limit where there is less benefit to generating more data for this model. Given some accuracy criterion, this

could be used to halt the data generation once sufficient accuracy has been reached.

In many practical scenarios, access to high-fidelity data is limited, be it due to computationally expensive simulations or experimental testing, which prohibits robust surrogate model fitting. In contrast, low-fidelity data may introduce bias and/or noise in the surrogate predictions on unseen conditions. Multi-Fidelity Modeling (MFM) addresses this issue through leveraging the computational efficiency of low-fidelity data generation to learn trends, allowing for a large enough dataset to provide meaningful inference, while increasing accuracy of the predictions through select, high-fidelity data anchoring points.




Figure 14. Prediction Accuracy as a function of Training Dataset Size.

The co-kriging technique, originally developed for geostatistical inference, has shown great success in MFM for small to mid-sized data sets in a vast range of engineering applications, owing to its robustness and straight-forward implementation. The application of MFM to rotorcraft analysis has gained increasing attention in the pursuit of integrating higher-fidelity CFD data with predictions from reduced-order, computationally less expensive aerodynamic models, Refs. [19-20]. In the context of the LRMR



Figure 13. NASA Advanced Supercomputing Division Hyperwall, 128 LRMR Geometric Perturbations.

optimization, the airframe aerodynamic V1 database did not include effects of the rotors, which were later included in the V2 database. Yet, general trends such as those with cruise speed and angle of attack are shared across the two data sets, as shown in Figure 15.

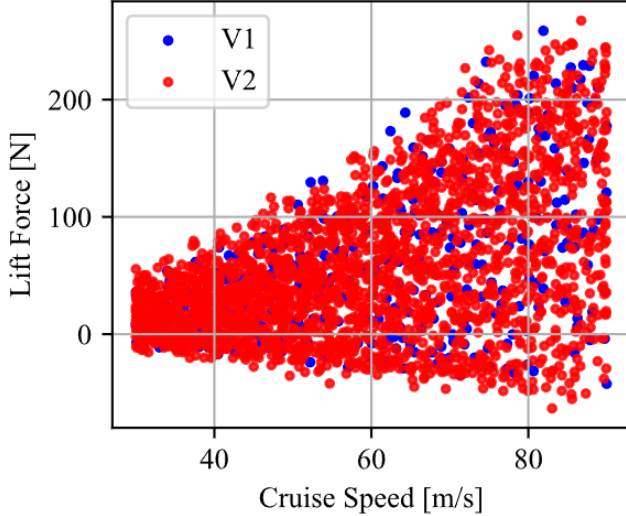


Figure 15. Airframe Multi-fidelity Aero Database.

Here, the cruise-speed-dependent lift force is co-plotted for the V1 and V2 databases, showing it as an appropriate MFM application. The linear, autoregressive model formulated by Kennedy and O’Hagan in Ref. [21], and non-linear generalization by Perdikaris, Ref.[22], as embodied in the open-source Python package Emukit, Ref. [23] was used for a demonstration of MFM here for the airframe aerodynamics databases.

The data was normalized with dynamic pressure prior to training to reduce the dominant dependence on forward flight speed. A linear as well as non-linear model formulation for the cross-correlation between the two “fidelities” were investigated. For reference, the two multi-fidelity models were contrasted against an exact GPR trained exclusively on V2 data, considered as the “high-fidelity” data. All models featured Radial Basis Function kernels, without a noise parameter. The evaluation was performed on a 10% hold-out from the V2 database. The MAE on the predicted lift, drag and pitching moment coefficients are presented in Table 4. Notably, the exact GPR predictions are improved 25 to 35% with the MFM approach, where the non-linear model is shown to slightly outperform the linear model.

Table 4. Mean Absolute Error for MFM and Single-Fidelity Models for Airframe Aerodynamic Databases.

MAE	Non-linear MFM	Linear MFM	GPR HF
C_L	0.33	0.35	0.51
C_D	0.0530	0.0556	0.0736
C_M	0.040	0.042	0.054

Figure 16 displays the correlation of MFM predictions on the hold-out data for the airframe’s aerodynamic performance. Lift normalized by dynamic pressure is plotted on the y-axis versus cruise speed. In almost all cases, the

MFM model is shown to predict the unseen data very closely. These results confirmed the statistical analyses observed in Table 4, and indicate that the MFM model can be used in downstream applications to accurately predict the airframe aerodynamic performance as a function of the geometry and flight condition inputs.

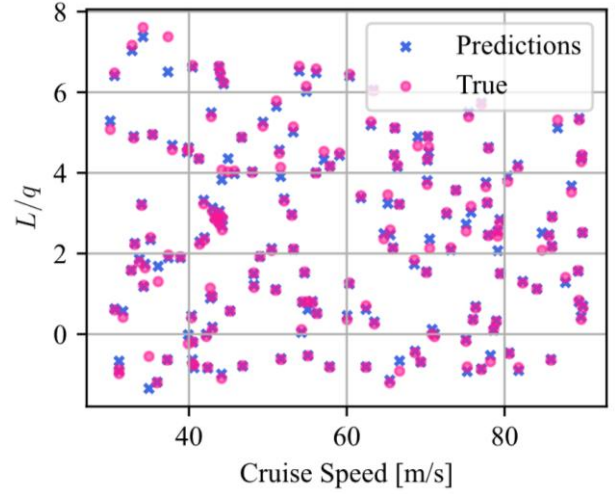


Figure 16. Airframe Surrogate Model Predictions.

ROAMX AIRFOIL PERFORMANCE MODELS

This work leveraged previous efforts on rotor design optimization in Mars conditions as performed under the ROAMX project, Ref. [15-17]. Data from the ROAMX project has been used to create a baseline rotor, and surrogate models of ROAMX Mars airfoil performance data have been created to simultaneously pursue full-aircraft optimization.

The ELISA optimization toolset, developed under the ROAMX project and detailed in Ref. [16], is a multi-objective genetic algorithm approach to airfoil and rotor optimization using OVERFLOW, tailored to low-Reynolds, high Mach number applications typical for Mars operating conditions.

ELISA features a dedicated airfoil parametrization, readily allowing for efficient exploration of unconventional airfoil shapes. Thickness and camber distributions are uniquely defined through Bezier curves; the notation for the parametrization of the ROAMX airfoil classes is $roamx-n_c p_c n_t p_t$, where n denotes the number of nodes and p the order of the curve segments, where subscripts c and t refer to camber and thickness, respectively.

The computed aerodynamic coefficients c_l , c_d , and c_m from the iterations of ELISA optimization runs for the $roamx-0202$ and $roamx-0201$ served as the airfoil performance database used in this work. Including Mach and Reynolds number, the total number of independent variables for the $roamx-0201$ and $roamx-0202$ models is 5 and 7, respectively. The $roamx-0201$ were considered for the outboard section of the blade, with data covering solutions at Mach conditions 0.3, 0.4, 0.6, 0.8, and 0.9. The variable thickness airfoil parametrization, $roamx-0202$, were considered for the inboard segments, ran at Mach number 0.07 and 0.2, to better accommodate structural requirements, in a similar vein as for the ROAMX rotor design as described in Ref. [16]. The ROAMX airfoil data c_l/c_d distribution is shown in Figure 17.

The data set contains more than 60,000 OVERFLOW airfoil simulations. The airfoil data used in this approach is from Reference [17] and will be called the ROAMX data from hereon out. The data for Mach = 0.90 has been generated using ELISA to expand the data set for this project.




Figure 17. ROAMX Data c_l/c_d Distributions.

Deriving models from a data set that is noisy or one containing outliers, which are common occurrences for CFD data, can skew the model predictions unfavorably. To mitigate bias in the fit models, a simplistic data pruning approach was employed for the data sets, leveraging Mahalanobis distance to sift out statistical outliers. This distance measure is a multi-variate extension of Z-score to account for correlation structure between variables in multi-dimensional datasets. In Figure 18, a cross-section of the roamx-0202 data is shown, highlighting data points that are significantly separated from the data distributions and removed automatically through the data pruning. A cut-off percentile of 95% was employed in the data pre-processing throughout this work, unless stated otherwise.




Figure 18. Removing Outliers from Training Dataset, ROAMX-0202, Mach = 0.2.

Due to the large size of the ROAMX training dataset and the poor scalability of exact GPRs, a scalable approach through sparse approximations using inducing variables was explored, Ref. [24]. While exact GPRs experience a cubic training time with the number of data points, n , a sparse approximation with m inducing variables allows for a more tractable time complexity $O(nm^2)$. Using 2,500 inducing points for the ~35,000-point roamx-0201 data set, three separate model for the lift, drag and pitching moment coefficient polars were fit using a Matern kernel. The training time on a single CPU was on the order of one hour. A parity plot for the drag coefficient c_d on a 10% hold-out test data set is presented in Figure 19, showing decent correlation between the predicted and ground truth values. The data has a reported R-squared metric of 0.97, MAE of 0.009, and an 11% mean absolute percentage error (MAPE). The predictive capability of the model tends to decrease with increasing drag.




Figure 19. Sparse GPR Surrogate Model Predictions of ROAMX-0201 Airfoil Drag Coefficient.

While the sparse approximator greatly improves the tractability of the surrogate, the prediction accuracy on the high-dimensional dataset is not entirely satisfactory. Thus, moving away from Kriging approaches, alternative scalable methods to improve the accuracy and computational efficiency, specifically feed-forward neural network approaches, were explored further under this effort.

Dense, feed-forward networks were trained and evaluated on the ROAMX data. Models used for benchmarking featured 7 hidden layers, each with 100 densely connected neurons. As a general observation, relatively narrow and shallow feed-forward networks with a smaller number of trainable parameters were shown to perform better than wider, deep models. Unless stated otherwise, 10% of the training data was held out and used for evaluating the models on unseen test data points. In general, all neural networks were shown to outperform the sparse GPRs, both in terms of prediction accuracy and speed. Initially, six separate neural networks were trained to predict the airfoil performance of the two different airfoil classes. To improve compactness and potentially reduce training time, two multi-output models that predict all three aerodynamic coefficients simultaneously of the two airfoil

classes were also considered. The MAPE as predicted on the outboard section (roamx-0201) by the multi-output model is reported in Table 5, showing that the lift and drag coefficients are predicted to within about a percent.

Table 5. Multi-Output Model Prediction Accuracy on roamx-0201 Test Data.

Mean Absolute Percentage Error		
c_l	c_d	c_m
1.19%	1.28%	3.48%

The larger error in the pitching moment coefficient is likely due to the characteristics of the underlying data set. The ELISA tool objective functions include maximizing lift and minimizing drag, which ensure a certain amount of diversity and structure in the data for the lift and drag coefficients, due to intrinsic properties of the Pareto front. However, less structure and more noise in the data set is observed as it pertains to the pitching moment coefficient, which inevitably propagates to the quality of the predictions. Finally, a “unified” model that predicts across the two airfoil classes was also considered, augmenting the constant thickness airfoil data with two “dummy” variables to have a constant zero thickness distribution (in addition to the baseline, prescribed minimum thickness). With this augmentation, the models were trained on the entire ROAMX data set. A parity plot of the predicted drag coefficient using the unified model is presented in Figure 20, together with the MAPE.

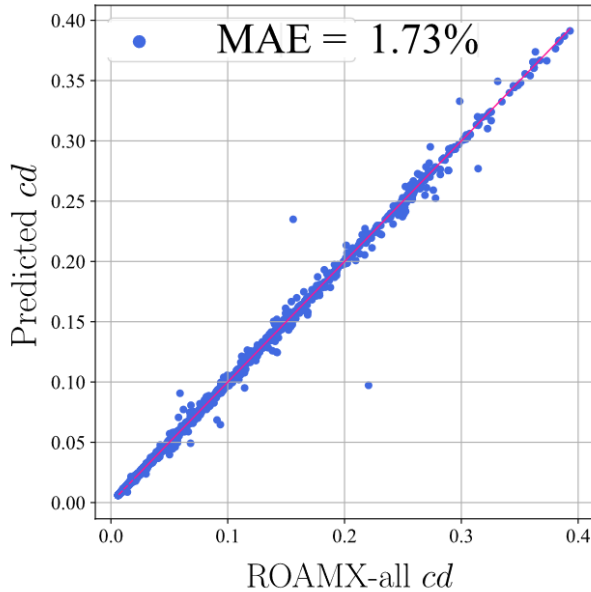


Figure 20. Unified Airfoil Model Prediction Accuracy on Test Data.

Besides a small number of outliers, the model is seen to follow the truth data closely; the discrepancy falls within 2% of the unseen OVERFLOW CFD data. This is a decrease in error of a factor of five from the sparse GPR.

Moving beyond benchmarking models, a framework to optimize the hyperparameters of the neural networks was implemented through a Bayesian Optimization routine, as embodied in the optimization framework Optuna, Ref. [25]. Training and evaluating neural networks are highly parallelizable tasks, and as such, enjoy great computational benefits from employment on GPU high performance compute clusters. Thus, utilizing the NASA HECC GPU’s enables an otherwise infeasible number of training cycles to find an optimal set of hyperparameters, and was used for this study. In the initial development trials, parallelized model training was employed on NVIDIA V100 and A100 GPUs using the NASA HECC resources. For reference, a wall time of two hours on two GPU’s allowed for approximately 45 trained models. The framework can readily be expanded to any arbitrary number of GPU nodes to increase the number of networks tested in a given amount of time.

FLIGHT CONTROLS

Flight Control Allocation

Control allocation for the LRMR vehicle is depicted in Figure 21 with the rotor locations, spin directions (blue and gray opposite), and orientations. The outboard rotors have five degrees of inward cant aligned with the wingspan direction. The rotor cant angle is subject to change based on preliminary analysis of the control authority of the current design. For this configuration, the cant angle is more effective at controlling yaw (in hover conditions) and roll (in cruise conditions) than differential torque alone.




Figure 21. Fault-tolerant Control Allocation for a Long-Range Mars Rotorcraft.

The aircraft is nominally collective control only, but by having cyclic available on rotors 2 and 5, the aircraft can lose any single or diagonal pair of rotors and still be controllable in all axes. The control allocation matrix for differential blade pitch control is depicted in Figure 22. The matrix was selected to contain values of unity for this study, but these values could be further optimized to achieve different design goals (e.g., minimize control power, maximize control authority). In vertical flight mode, roll control is achieved through differential blade pitch of the left (1, 6) and right (3, 4) rotors,

pitch control is achieved through differential blade pitch of the fore (1, 3) and aft (4, 6) rotors, and yaw control is achieved through a combination of increasing thrust on the positive-cant rotors (1, 4) along with differential cyclic of the 2 and 5 rotors while decreasing thrust on the negative-cant rotors (3, 6). In forward flight mode, the roll and yaw axes are swapped in an inertial sense, but the control action remains the same.

This control allocation was used in both OVERFLOW and CAMRAD-II to generate the aerodynamic performance databases and subsequent stability and control derivative matrices.

$$\begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \\ \beta_5 \\ \beta_6 \end{bmatrix} = \begin{bmatrix} 1.00 & 1.00 & 1.00 & 1.00 \\ 1.00 & 0.00 & 0.00 & 0.00 \\ 1.00 & -1.00 & 1.00 & -1.00 \\ 1.00 & -1.00 & -1.00 & 1.00 \\ 1.00 & 0.00 & 0.00 & 0.00 \\ 1.00 & 1.00 & -1.00 & -1.00 \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_\phi \\ \beta_\theta \\ \beta_\psi \end{bmatrix}$$

Figure 22. Control Allocation Matrix.

Control Authority Analysis

Preliminary assessment of the LRMR designs was based on their bare-airframe eigenvalues. From classical control theory, the inherent stability of a system can be analyzed using a complex plane where eigenvalues with negative real parts indicate stability while those with positive real parts represent unstable modes. Analogously, the higher the magnitude of the real part of the eigenvalue, the quicker that mode either stabilizes or destabilizes (depending on if its positive or negative). Using this analysis to drive the vehicle optimization towards more inherently stable dynamics is hypothesized to benefit control design and flying qualities; however, it should be noted that systems with eigenvalues far inside the stable, left-half plane typically require more energy to actuate. To find out, a cost function was developed to quantify stability in a normalized fashion such that it could be easily weighted or manipulated within the global aircraft optimization framework. This was achieved by using the following sigmoid function:

$$\sigma(x) = \frac{1}{1+e^{-0.06x}} \quad (5)$$

Here the exponential component was chosen such that the tapering of the cost function would occur at approximately ± 50 along the real axis. This was done to ensure a sufficient region for differentiating between costs of the various systems and was chosen to be approximately 2.5 times the highest, real-part eigenvalue magnitude for an earlier baseline LRMR vehicle which was around 20.

Oscillatory modes (eigenvalues with imaginary parts) were also included with the cost function. Specifically, the damping ratio, ζ , was chosen as it reflects how vibrations are attenuated. The range of underdamped damping ratios ($0 < \zeta < 1$) are of particular interest because values closer to 1 exhibit very little oscillation while values closer to 0 experience oscillations that decay relatively slowly. Figure 23 shows how damping ratio maps onto the complex plane.




Figure 23. Damping Ratio in the Complex Plane.

To avoid sustained oscillations, the cost function was updated to incorporate additional cost for the left-half plane eigenvalues with low damping ratios. Thus, the mapping seen in Figure 24 is reflected in the contours of the final cost function surface shown in Figure 25.

Lastly, an arbitrary amount of integrator states can potentially be augmented into the system without affecting the dynamics, so poles located at the origin were assigned a cost of zero. This can be seen as a spike beneath the cost function surface in Figure 25. The final cost function is defined by Equation 6.

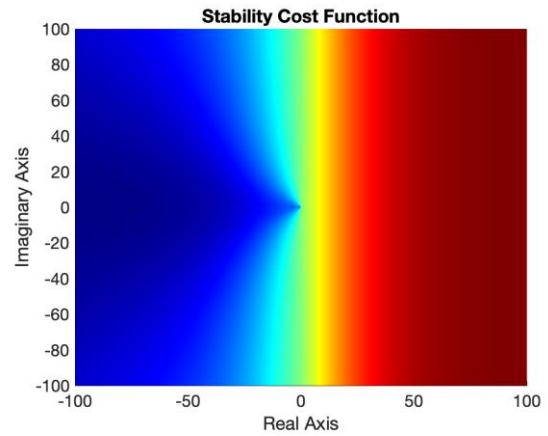


Figure 24. Stability Cost Function.

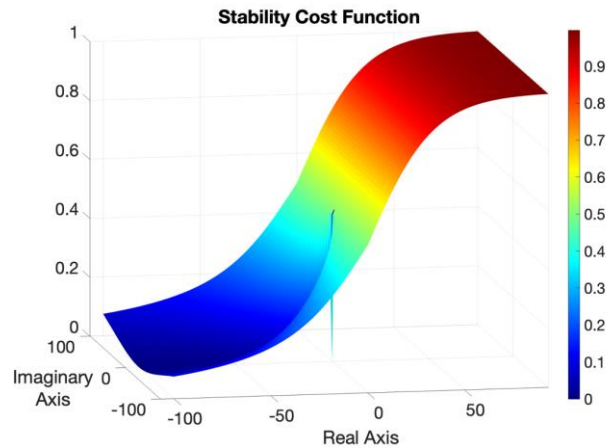


Figure 25. Perspective View of the Stability Cost Function Surface.

$$J = \sum_{i=1}^N \begin{cases} 0 & \text{if: } \lambda_i = 0 \\ \frac{1-\zeta_i}{4} + \frac{0.5}{1+e^{-0.06\sigma_i}} & \text{if: } \sigma_i \leq 0 \\ \frac{1}{1+e^{-0.06\sigma_i}} & \text{if: } \sigma_i > 0 \end{cases} \quad (6)$$

Where i represents the N number of eigenvalues, λ , in each system; σ represents the real part of the eigenvalues; and ζ represents the damping ratio for complex eigenvalues.

Although outside the current scope, without a control law and higher fidelity modeling, it is unclear whether this method may be pushing the system towards an area sensitive to resonant frequencies. This is a topic that will be considered for further analysis in future work. The result from the cost function for stability of each aircraft will be used as a metric to drive the overall optimization.

DESIGN OPTIMIZATION METHODOLOGY

The optimization framework used in this work builds on previous work by the first author, Ref. [26]. Figure 26 shows a high-level layout of the various modules in the optimization pipeline. The entire framework is coded in Python and run using a single shell script which requests resources on the NASA supercomputer and queues the main “Optimization Framework” script.

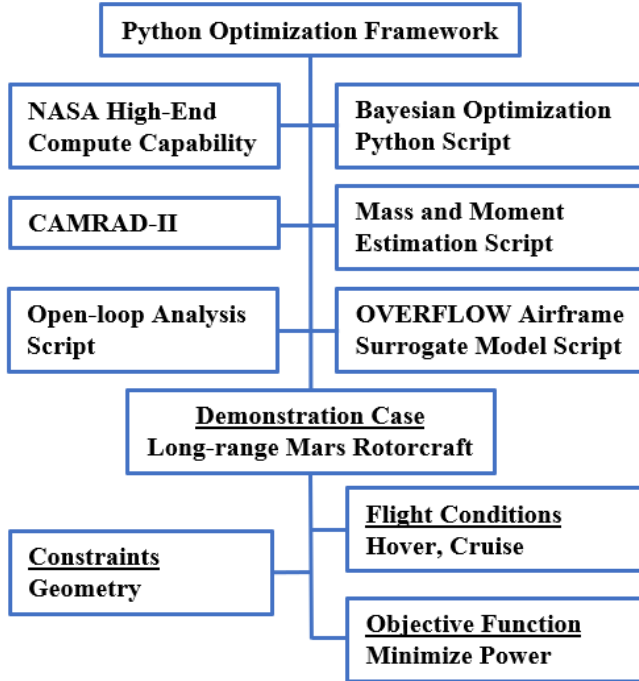


Figure 26. ML-based Optimization Framework.

The framework starts with an initial seeding, which can either be read in from a text file or generated using the GreedyFP space filling algorithm. Implementation improvements of the GreedyFP algorithm enabled a more than 100x speedup over previous work, enabling 10,000 intelligent samples to be generated from a random initial pool of 200,000 in a matter of minutes on a single CPU.

The current approach uses Bayesian optimization, which is a two-step method whereby training data is first generated using a more computationally expensive function evaluator. In this case, CAMRAD-II is the function evaluator. The initial training data is then used to build a surrogate model of the objective function and any constraints, which are then used to evaluate a much larger number of design solutions efficiently, but with lower accuracy than the function evaluator.

An acquisition function is then used to down-select the most promising solutions from this larger set, which are then run in CAMRAD-II. The new output is added to the training data, the surrogate models are improved, and the process continues until a convergence criterion or pre-determined number of iterations completes. This work uses the expected improvement acquisition function detailed in Equations 7-8.

$$EI(x) = (y_{best} - \mu(x)) * \Phi(z) + \sigma(x) * \phi(z) \quad (7)$$

$$z = ((y_{best} - \mu(x) - \zeta)) / \sigma(x) \quad (8)$$

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 framework has successfully run on five NAS GPU nodes, which equates to 240 CPU cores and 20 V100 GPUs. The master script does some minimal pre-processing, including generating the mass and moment of inertia values needed by CAMRAD-II. It then distributes individual cases among the available CPU cores and further continues the pre-processing step by querying the airframe aerodynamic surrogate models generated from the OVERFLOW databases. This step creates an airframe aerodynamic look-up table specific to the individual case’s geometry.

CAMRAD-II Modeling

The CAMRAD-II model consists of six rotors in an equally spaced hex-rotor configuration. The rotor locations and rotation directions match the schematic in Figure 21. Rotors 1, 3, 4, and 6 have 10 degrees of inward cant for enhanced yaw stability in forward flight. Although this hex-rotor configuration with rotors 2 and 5 having cyclic control was chosen to increase system redundancy in the event of a motor failure, the current CAMRAD-II model analyzes the nominal case and uses collective control only for each rotor.

The six rotors are currently fixed as the ROAMX baseline rotor, which has been found to yield improved performance compared to the Ingenuity Mars Helicopter for operation in Mars atmospheric conditions, Ref. [16]. Airfoil performance C81 look-up tables from the ROAMX team are used. The tables were generated using the ELISA tool as a wrapper for two-dimensional OVERFLOW airfoil calculations.

The OVERFLOW Airframe Surrogate Model Script is used to generate an airframe aerodynamics look-up table for

each individual CAMRAD-II case, which is queried as a function of airframe geometry from the high-dimensionality database. The resulting table has airframe performance as a function of flight speed and vehicle pitch attitude.

Each CAMRAD-II job analyzes the given vehicle configuration in two conditions: hover and cruise. The hover condition has a simple one degree of freedom trim with rotor common collective control input used to trim the thrust to be equal and opposite to the weight of the aircraft calculated using the mass and moment estimation script.

The cruise case flips the aircraft 90 degrees and solves the two-degree-of-freedom trim equations in forward flight. Thrust is trimmed equal to drag, lift equal to weight, and pitching moment trimmed to zero. The two control inputs are rotor common collective control and rotor longitudinal cyclic, which fully defines the collective pitch for each of the six rotors following the control allocation described in Figure 22. The same rotor input decks and airframe aerodynamic tables are used. Rotor tip speed is modulated as a function of forward flight speed to reduce the resultant blade-tip Mach number, as was done in the OVERFLOW BET simulations.

The CAMRAD-II flutter task is then performed to calculate the stability and control derivative matrices about that trim condition. This output provides information about the static stability and control of each aircraft geometry under the given flight condition.

Various metrics for aircraft performance (e.g., total power requirement) and the cost metric from the section describing the control authority analysis are then used as inputs to the Bayesian optimization module. The entire framework has been created and the LRMR aircraft can currently be optimized for a single objective at a time, with other metrics applied as constraints in the optimization.

SUMMARY

This work set out to explore several technical challenges required for multi-fidelity surrogate modeling towards aerospace vehicle optimization. An optimization framework was created for a Long-Range Mars Rotorcraft to demonstrate the feasibility of the approach, but the methodology is general and applicable to a wide range of aerospace design problems and other data-driven fields. The list below summarizes a few of this work's major successes and findings:

1. Large-scale aerodynamic simulation databases were created using the NASA OVERFLOW Computational Fluid Dynamics (CFD) flow solver. This was enabled by three major advancements that were tested extensively throughout this work: a.) automated structured grid generation, b.) GPU computing implementation of the tool, and c.) a blade-element theory description of the rotor disk. In total, more than 3,000 high-fidelity CFD simulations were run of the full LRMR aircraft with two wings, a fuselage, pylons, and six rotors. This used approximately 80 GPUs with a completion rate of roughly 400 cases per day. The simulation database spanned airframe geometric perturbations, flight

condition perturbations, and control input perturbations. Simulations were selected using the Greedy Farthest Point sampling algorithm, which saw speed increased by 100x in this work over the prior implementation by the authors.

2. Surrogate models were tested and developed using approaches such as Gaussian Process Regression (GPR) and Neural Network (NN) approaches. Various implementations were tested such as sparse GPRs as well as shallow and deep NNs. Co-kriging multi-fidelity models (MFM) were also created using the combined V1 and V2 OVERFLOW LRMR aerodynamic database. The NN was found to outperform the GPR by a factor of 5 (MAE), and the multi-fidelity database was found to further increase accuracy of the resulting surrogate model. Surrogates were also created for the ROAMX airfoil performance data, and various methods were tested to combine the thin outer radial station airfoils with the thicker inboard radial station airfoils. These surrogate models were used as inputs into the CAMRAD-II comprehensive analysis model.
3. Python scripts were created to estimate the mass and moment of inertia of the LRMR given a set of geometric parameters. A target mass value was selected at 75 kg, and any remaining mass after estimating the aircraft weight was filled with additional battery capacity. This mass and inertia information was used as an input to the CAMRAD-II model.
4. A CAMRAD-II simulation model of the LRMR was created to calculate both aircraft performance and the stability and control derivative matrices for a given aircraft geometry, flight condition, and control input. This model used airframe aerodynamic look-up tables from the OVERFLOW-based surrogate models. It also incorporated the mass and moment of inertia data for each aircraft geometry. The simulation used one degree of freedom trim in hover and two degree of freedom trim in forward flight to calculate the performance metrics for the aircraft. The same 3,000 flight conditions that were run in OVERFLOW to create the V2 aerodynamic database were run through CAMRAD-II, creating another mid-fidelity source of performance data for those conditions to be used in future MFM efforts.
5. A cost function was created to assess the controllability of the aircraft. In the case of the LRMR, this was derived from an open-loop eigenvalue analysis of the stability and control derivative matrix output. The result of this cost function, along with other performance metrics from CAMRAD-II, can then be used to drive and constrain the optimization.

Through the above steps, an optimization framework has been created to allow for iterative Bayesian optimization of the Long-Range Mars Rotorcraft. The methodology makes

use of high-performance GPU computing, machine learning based surrogate modeling, and multi-fidelity simulation data to accomplish the objectives of this work.

The framework can now be used to optimize the LRMR for a user-defined objective function, with additional metrics applied as constraints in the optimization. This work has demonstrated the ability of ML-based surrogate modeling to bring high-fidelity simulation data forward in the conceptual design process.

FUTURE WORK

Follow-on work exists to conduct constrained design optimization and generate optimal LRMR aircraft. Several objectives can be defined to assess how each choice influences the resulting aircraft size, geometry, and optimal flight conditions. Potential multi-objective optimization, and hybrid Bayesian optimization approaches could be explored to assess optimization convergence rates.

Additional studies in multi-fidelity surrogate modeling combining the OVERFLOW and CAMRAD-II datasets would also be highly informative to see how surrogate models built with completely different tools would perform.

Corresponding Author: jason.k.cornelius@nasa.gov

ACKNOWLEDGMENTS

This work was funded by the NASA Digital Transformation Office through a Prototype Project award. The authors would like to thank Patrick Murphy, Hannah Brooks, Tracy Bierman, and Maria Revlett for their support of the project. Additionally, the authors thank Larry Young, Wayne Johnson, Carlos Malpica, Witold Koning, Ethan Romander, Larry Hogle, Tom Norman, William Warmbrodt, and Carl Russell for their support of this work. The authors also thank the NASA Revolutionary Vertical Lift Technology Project, the NASA ROAMX ECI project, and the US Army at Moffett Field for their technical guidance, feedback, and mentorship throughout the project.

REFERENCES

1. Balaram, J., Aung, M., Golombek, M. P., "The Ingenuity Helicopter on the Perseverance Rover," *Space Science Reviews*, Vol. 217, No. 56, 2021.
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. NASA Ames Planetary Aeolian Laboratory, NASA Ames Research Center. <https://www.nasa.gov/thermophysics-facilities-branch-planetary-aeolian-laboratory/>
4. Space Simulator Facility, NASA Jet Propulsion Laboratory, last updated September 27, 2023. <https://www.nasa.gov/setmo/facilities/25-foot-space-simulator/>
5. Buckley, M., "A Taste of Titan," Johns Hopkins Applied Physics Laboratory, November 20, 2023. <https://www.jhuapl.edu/news/news-releases/231120-dragonfly-titan-chamber>
6. Cornelius, J., Schmitz, S., "Dragonfly Rotor Optimization using Machine Learning Applied to an OVERFLOW Generated Airfoil Database," VFS Forum 80, Montreal, Canada, May 2024.
7. 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>
8. Ingenuity Helicopter Flight Log, NASA JPL, June, 2024. <https://science.nasa.gov/mission/mars-2020-perseverance/ingenuity-mars-helicopter/>
9. Hrishikeshavean, V., Bogdanowicz, C., Chopra, I., "Experimental Investigation of Performance of a Wing-Propeller System for a Quad-Rotor Biplane Micro Air Vehicle," AIAA Structures, Structural Dynamics, and Materials Conference, April 2013.
10. Reddinger, J., "Performance and Control Scalability of a Quadrotor Biplane Tailsitter," AIAA SciTech, Jan. 2019.
11. Singh, R., Hrishikeshavan, V., Sirohi, J., "Common Research Configuration for Collaborative Advancement of Scalable VTOL UAS Technologies," VFS Forum 75, May 2019.
12. Govindarajan, B., Sridharan, A., "Conceptual Sizing of Vertical Lift Package Delivery Platforms," *Journal of Aircraft*, Vol. 57, No. 6, 2020.
13. Sridharan, A., Govindarajan, B., Chopra, I., "A Scalability Study of the Multirotor Biplane Tailsitter Using Conceptual Sizing," *Journal of the American Helicopter Society*, Vol. 65, 2020.
14. Hollis, B., "Blunt-Body Entry Vehicle Aerothermodynamics: Transition and Turbulence on the CEV and MSL Configurations," AIAA 40th Fluid Dynamics Conference, June 2010. <https://doi.org/10.2514/6.2010-4984>
15. Cummings, H., Perez Perez, B.N., Koning, W., Johnson, W., Young, L., Haddad, F., Romander, E., Balaram, J.B., Tzanetos, T., Bowman, J., Wagner, L., Withrow-Maser, S., Isaacs, E., Toney, S., Shirazi, D., Conley, S., Pipenberg, B., Datta, A., Lumba, R., Chi, C., Smith, J.K., Cornelison, C., Perez, A., Nonomura, T., Asai, K., "Overview and introduction of the Rotor Optimization for the Advancement of Mars eXploration (ROAMX) project," VFS Transformative Vertical Flight Conference, San Jose, CA, 2022.
16. Koning, W., Perez, B., Cummings, H., Romander, E., Johnson, W., "ELISA: A Tool for Optimization of Rotor Hover Performance at Low Reynolds Number in the Mars Atmosphere," VFS Transformative Vertical Flight Conference, San Jose, CA, 2024.
17. Koning, W., Perez, B., Cummings, H., Romander, E.,

- Johnson, W., "Overview of Rotor Hover Performance Capabilities at Low Reynolds Number for Mars Exploration," 50th European Rotorcraft Forum, Marseille, France, Sept. 2024.
18. Cole, G., and Mueller, T., "Experimental Measurements of the Laminar Separation Bubble on an Eppler 387 Airfoil at Low Reynolds Numbers," University of Notre Dame Aerospace Series (UNDAS)-1419-FR, 1990.
 19. Titsias, M. K., "Variational Learning of Inducing Variables in Sparse Gaussian Processes," Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, 2009.
 20. Kennedy, M. C., O'Hagan, A. "Predicting the output from a complex computer code when fast approximations are available". *Biometrika*, Vol. 87, pp. 1-13, 2000.
 21. Perdikaris, P., Raissi, M., Damianou, A., Lawrence, N. D., Karniadakis, G. E., "Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 473, 2017.
 22. Anusonti-Inthra, P., "Data-Driven Modeling of Aerodynamic Loadings for Tiltrotor Pylon using Multi-Fidelity CFD Data," AIAA SciTech 2023 Forum, January 2023.
 23. Walizer, L., Robert B. H., Luke A., "Application of Multi-fidelity Methods to Rotorcraft Performance Assessment," AIAA SciTech 2024 Forum, January 2024.
 24. Paleyes, A., Pullin, M., Mahsereci, M., Lawrence, N., Gonzalez, J., "Emulation of Physical Processes with Emukit," 2nd Workshop on Machine Learning and the Physical Science, NeurIPS, 2019.
DOI:10.48550/arXiv.2110.1329
 25. Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M. "Optuna: A Next-generation Hyperparameter Optimization Framework," Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, July, 2019.
 26. Cornelius, J., Schmitz, S., "Dragonfly Rotor Optimization using Machine Learning Applied to an OVERFLOW Generated Airfoil Database," VFS Forum 80, Montreal, Canada, May 7–9, 2024.