AIAA SciTech Forum
6-10 January 2025, Orlando, FL
AIAA SCITECH 2025 Forum

10.2514/6.2025-0038

PALMO: An OVERFLOW Machine Learning Airfoil Performance Database

Jason K. Cornelius[1] and Nicholas Peters[1]
NASA Ames Research Center, Moffett Field, California, USA

Tove Ågren[1] and Darrell Nieves Lugo[2]
Science and Technology Corporation, NASA Ames Research Center, Moffett Field, California, USA

The OVERFLOW Machine Learning Airfoil Performance (PALMO) database has been created to enable robust modeling of airfoil performance in a variety of applications. The PALMO database uses OVERFLOW simulation data second-order accurate in time and fourth-order accurate in space with Spalart-Allmaras turbulence closure. The foundation of the in-development PALMO database is the airfoil base cube. Each base cube includes simulation data parametrized over a range of Mach numbers, Reynolds numbers, and angles-of-attack. This database includes the NACA 4-series airfoils, with parametrization in airfoil thickness and camber from an NACA 0006 to an NACA 4424. In total, 52,480 NACA 4-series OVERFLOW calculations were run on the NASA High-End Compute Capability (HECC) supercomputer. This provides high-order-accurate simulation data covering a wide range of aerospace design applications, which enables users to develop accurate airfoil performance look-up tables without additional high-performance computing. In addition to engineering design and analysis of aerospace vehicles, PALMO is well suited to be a benchmark dataset for the development and testing of machine learning methods in aerospace engineering. This work presents an example PALMO surrogate model that enables accurate airfoil performance predictions for any arbitrary combination of camber, thickness, Mach number, Reynolds number, and angle of attack within the bounds of the database. Airfoil performance tables predicted for an airfoil not used in training the model are used in three-dimensional OVERFLOW simulations to quantify the downstream accuracy on aggregate rotor performance metrics. For the NACA 3415 airfoil, which had no common thickness or camber with the training data, the surrogate predicted and CFD generated tables were within 2.1% of each other in the forward flight lift to drag metric. This suggests that performance tables generated for airfoils within the bounds of the PALMO database will yield aggregate rotor performance predictions on par with tables generated from directly running OVERFLOW airfoil calculations. The PALMO airfoil performance coefficients are available publicly.

Nomenclature

BEMT	=	blade element momentum theory
CFD	=	computational fluid dynamics
CPU	=	central processing unit
CST	=	class shape transformations
FM	=	rotor figure of merit (hover efficiency)
HECC	=	NASA high-end compute capability

---

[1] Aerospace Engineer, Aeromechanics Office, and AIAA Member.
[2] Student, University of Central Florida, and AIAA Student Member.

LDE	=	rotor lift to drag (forward flight efficiency)
MAE	=	mean absolute error (surrogate model accuracy)
MAPE	=	mean absolute percent error (surrogate model accuracy)
NACA	=	National Advisory Committee for Aeronautics
PALMO	=	OVERFLOW machine learning airfoil performance database
UIUC	=	University of Illinois at Urbana Champaign
α	=	angle of attack, deg
C_T	=	rotor thrust coefficient
C_Q	=	rotor torque coefficient
c_l	=	airfoil lift coefficient
c_d	=	airfoil drag coefficient
c_m	=	airfoil pitching moment coefficient

I. Introduction

The characterization of airfoil performance across a range of Mach numbers, Reynolds number, and angles of attack remains a key aspect for a vast array of mid-fidelity aerospace analysis and design methods. The computation of aerodynamic loads on rotors and wings often relies on previously computed airfoil performance tables. Although the mid-fidelity tools rely on these tables to reduce their computational cost, the process of creating accurate lookup tables is often a computationally intensive, and time-consuming task. Due to the still high computational cost of high-order CFD solvers, airfoil performance tables at proximal but mismatched Reynolds and Mach numbers are often used in conceptual design. These tables are sometimes created either with existing experimental data, which are limited by the test Mach and Reynolds number, or using a lower-fidelity approach such as XFOIL or MSES, Refs. [1-2]. Past studies by Patt and Youngren, Ref. [3], and Cornelius and Schmitz, Ref. [4], document both the need for higher refinement implementations of C81 tables and the improvements obtained by using them. Figure 1 exhibits the high-discretization of airfoil lookup tables needed for analyzing rotors with multiple airfoils, large changes in radial chord distribution, and variable-speed (varying RPM) operation, Ref. [4].

Creating high-fidelity lookup tables in each iteration of conceptual design is currently cost prohibitive. As a result, conceptual designers often make simplifying assumptions such as using a set of tables at constant Reynolds number even as the chord-based Reynolds number changes in successive design iterations, as done in Ref. [5]. An alternative approach is to use fast but lower-fidelity methods such as XFOIL, or the XFOIL generated UIUC database Ref. [6], to update the airfoil performance tables between iterations. Given the high cost of creating these lookup tables, there has been a growing interest in the aviation community to leverage various Machine Learning (ML) approaches to derive sufficiently accurate, low-cost surrogate models for predicting airfoil performance.

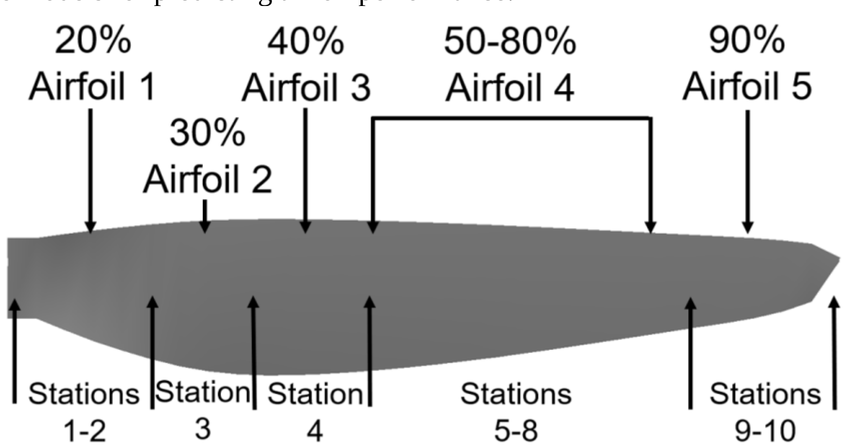

Figure 1. Example of High-Density Airfoil Lookup Table Discretization, duplicated from Ref. [4].

Some recent studies have used neural networks to create highly efficient airfoil performance surrogate models to update the lookup tables as the conceptual design progresses, but they have typically been based on lower-fidelity training data, Ref. [7]. Another study applied neural networks to data from the thin-layer Navier stokes flow solver ARC2D using the wrapper C81Gen, Ref. [8], 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. [9-13], with various approaches used to define the airfoil geometry such as the ParFoil tool and class shape transformations. Li et al. recently provided a review on this topic, Ref. [14]. These studies, however, typically rely on Class Shape Transformations such as Bernstein and Chebyshev polynomials, Refs. [15-17]. Although this process has been adopted as a leading research approach for airfoil shape optimization, the surrogate models can introduce some inaccuracy as compared to the original CFD calculations for airfoils in the training datasets due to imperfect representations of the airfoil shape.

This work pursues the creation of an airfoil performance database with the high order accurate OVERFLOW CFD solver. Machine learning (ML) is then leveraged to create surrogate models of the database for interpolation across the non-linear and high dimensionality OVERFLOW training datasets. This allows for rapid airfoil performance table generation at intermediate Mach and Reynolds numbers that are based on high-order accurate OVERFLOW training data. This work first develops a massive airfoil performance database of OVERFLOW CFD calculations run using the NASA HECC. The core airfoils in the database consist of first- and second-generation rotorcraft airfoils created first by the National Advisory Committee for Aeronautics (NACA), with continued development at NASA, Refs. [18-22].

OVERFLOW simulations that are 2^{nd} order accurate in time and 4^{th} order accurate in space are run with the Spalart-Allmaras turbulence closure to develop the airfoil performance training datasets, Refs. [23-24]. The airfoil simulations are run using the OVERFLOW wrapper AFTGen, Ref. [25]. Grid studies have been carried out to ensure grid convergence to less than 1% in the linear region of the lift-curve slope following the approaches documented by Cornelius, Koning, and Kallstrom, Refs. [26-28]. A typical simulation has 501 wrap-around points, 601 grid points normal to the airfoil, and 41 points on the blunt trailing edge giving approximately 325,000 cells. For each airfoil, 3,280 OVERFLOW simulations are run at all possible combinations of the Mach numbers, Reynolds numbers, and angles of attack reported in Table 1. Whenever possible, subsets of the OVERFLOW generated data will be compared to existing experimental datasets for validation of the training data used in the surrogate model generation, Refs. [18-22, 29-30].

II. PALMO Database Generation Version 1.0 NACA 4-Series

The PALMO database was generated using the NASA High-End Compute Capability (NASA-HECC). The foundation of the in-development PALMO database is the airfoil base cube, which is a high-density parametrization of OVERFLOW simulation data for a single airfoil. The data covers the ranges of Mach number, Reynolds number, and angle of attack reported in Table 1. This first version of the PALMO database has 16 base cubes, with each base cube representing a single NACA 4-series airfoil. The database currently contains 52,480 OVERFLOW simulations (that is, 3,280x16 simulations). Each base cube required five days of wall-clock time on sixteen 28-core Broadwell compute nodes. This resulted in a total computational cost of around 860,000 CPU hours.

The first set of airfoils included in the database are the NACA 4-series. This first-generation airfoil family was selected for the first release of the PALMO database for the following reasons:

1) ample publicly available experimental data in NASA and NACA documents, Refs. [18, 29, 30],
2) complete parametrization of the airfoil coordinates using just thickness and camber, and
3) relevance to a wide variety of aerospace design applications.

These conditions include a wide range of anticipated operating conditions from subsonic to transonic, Reynolds numbers of 75,000 to 8 million, and angle of attack values from -20 to +20 degrees. This is expected to bound many rotorcraft relevant applications that CFD simulation data would generally be used for. Below a Mach value of 0.25, the flow can be assumed subsonic, and thus the values from the lowest

Mach in the database can be used. Below the minimum Reynolds number, which may be encountered by micro unmanned aerial systems (MUAS) or extraterrestrial Mars helicopters, specially tailored OVERFLOW simulations beyond the scope of this work are required. For conditions in deep-stall, higher-fidelity simulations are required that are again beyond the scope of this database.

Table 1. The 3,280 Parametrized Conditions in a PALMO Base Cube.

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 to +20, 1-degree increments

The database is expanded beyond a single base cube with the addition of airfoil parametrization. Figure 2 shows a parametrization of the NACA 4-series with variations in percent camber and thickness being added to the database. For each airfoil, all possible combinations of the Mach number, Reynolds number, and angle of attack from Table 1 are simulated. The database includes the symmetric 4-series airfoils as well as airfoils with two and four percent camber. The maximum percent thickness also varies from six to twenty-four percent. This allows predictions for any arbitrary combination of camber and thickness within the bounds of the training data, i.e., from the NACA 0006 to the NACA 4424.

The red stars in Figure 2 represent additional test data generated beyond the original 4-series parametrization. These test data were generated using the same approach in OVERFLOW but are meant to assess the accuracy of downstream surrogate models. This first release of the full PALMO 4-series database thus has 16 base cubes. The twelve cubes shown in blue with four additional test cubes, denoted by the red stars, for the NACA 3415, 3418, 4415, and 4421. All 52,480 simulations could be used directly or as training data for surrogate models. For surrogate model development and testing, the 12 blue base cubes could be used as training data while holding out some or all the red stars for test data.

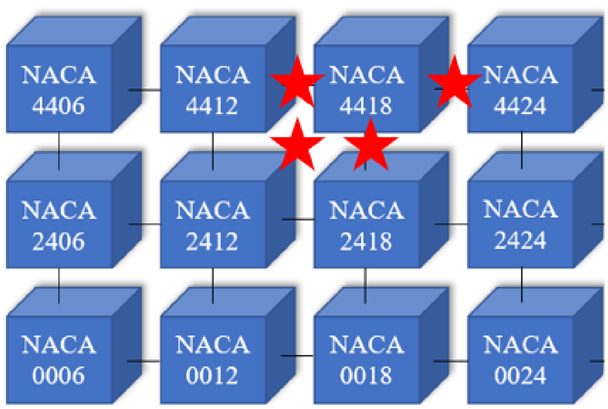

Figure 2. PALMO Database Version 1, NACA 4-series (Red stars for test data cubes).

Additional details about the database generation, including airfoil coordinate file generation, grid convergence studies, and OVERFLOW setup have been documented as a NASA Technical Memorandum, Ref. [31].

III. Surrogate Model Demonstration Case

An example surrogate model will now be presented to demonstrate use of the PALMO database. Although users can directly use the data as is, creating surrogate models from the data allows for airfoil performance prediction at intermediate airfoils and conditions not explicitly included in the database.

Notably, surrogate model interpolation of the database is shown to be more accurate than linear interpolation. Although not shown in this work, the surrogate model prediction speed for the neural networks used is faster than linear interpolation. This enables the user to rapidly generate airfoil lookup tables of similar accuracy to directly running OVERFLOW.

Throughout this work, all models were trained on the 12 baseline airfoils (the blue boxes from Figure 2). Three single-output models featuring three hidden layers with 200 neurons per layer were trained separately to predict the airfoil performance coefficients c_l, c_d and c_m. Each model was trained for 200 epochs. For the PALMO database, shallow and wide networks have consistently proven more accurate and less prone to over-fitting than deeper networks. A mean absolute error (MAE) was chosen as the loss function in training the networks. To assess the predictive capabilities of the models, airfoil performance was predicted for the NACA 3415 airfoil. The 3415 airfoil was one of the test base cubes from Figure 2 represented as a red star, which has different thickness and camber values from all the grid of blue box base cubes, making it a completely "unseen" set of datapoints.

Predictions from the neural network model were compared against both the OVERFLOW CFD data and a linearly interpolated model, which has traditionally been used in the absence of exact CFD and experimental data. These predictions were made using a multi-dimensional linear interpolation approach on the CFD training data. Predictions of airfoil lift and drag coefficient from both the neural network model and the linear interpolation are plotted against the OVERFLOW CFD data in Figure 3. The condition shown is for a Mach number of 0.25 and a Reynolds number of one million. The neural networks are seen to consistently correlate better with the CFD data, especially in the linear regime.

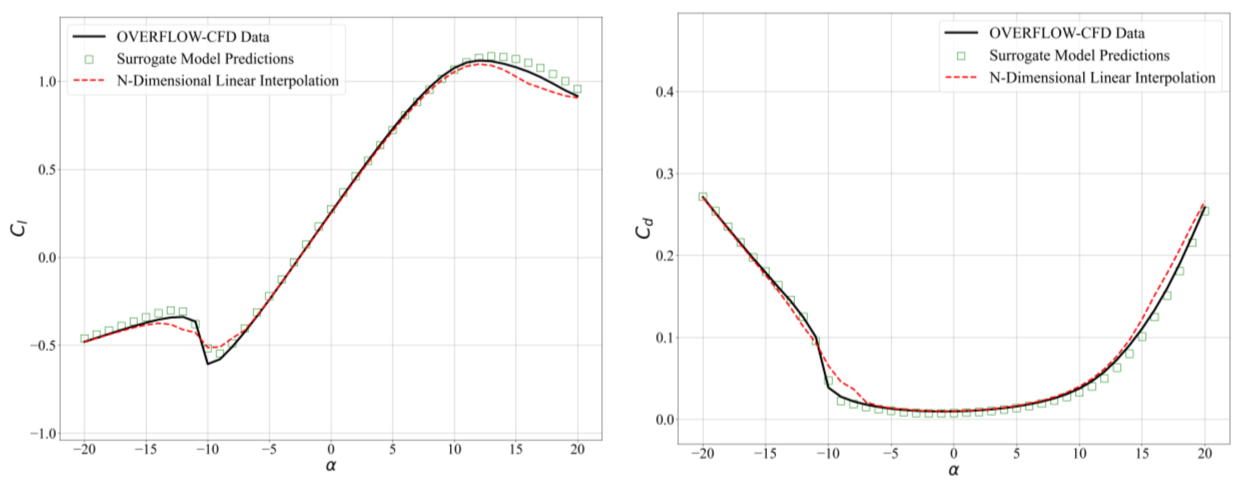

Figure 3. Single-output Model Predictions, Surrogate Model versus Linear Interpolation, NACA 3415, Mach Number = 0.25, Reynolds Number = 10^6.

In contrast to the three separate single-output models, a multi-output approach was also considered in this work. This model simultaneously predicts the three aerodynamic performance coefficients, offering compactness and less overhead, but with increased model size and longer training time. The multi-output model architecture in this work has 7 hidden layers with 100 neurons each. The model was trained for one thousand epochs. The same model training loss function, MAE, was used and compared to the MAE from the three single-output "expert" models. The test MAEs generated when predicting the NACA 3415 airfoil data are reported in Table 2. In this example, the model accuracy is comparable across the single- and multi-output approaches. The multi-output model performs better for each of the three airfoil performance coefficients. Although this is not always the case, observing this outcome for this data could be because the multi-output model benefits from the additional information yielded by including the relationships between lift, drag, and pitching moment coefficients. Still, this is a valuable sanity check and comparison study that should be conducted for each new set of data.

Table 2. Test Prediction Error Results on NACA 3415.

	Single Output MAE	Multi Output MAE
c_l	0.0087	0.0069
c_d	0.0098	0.0011
c_m	0.0094	0.0015

When these airfoil performance lookup tables are used in downstream applications, predictive accuracy around the nominal operating angle of attack will drive accuracy of the analyses. Thus, in addition to the full test data set, prediction errors were also computed for the angle of attack range from 2 to 8 degrees. In Table 3, the MAE as well as the mean absolute percentage error (MAPE) are reported for the multi-output model. The MAPE values are informative but can become inflated due to prediction values close to zero. Considering MAE, lift coefficient is predicted within 0.01 and drag is predicted within eight drag counts.

Table 3. Multi-output Model Errors on NACA 3415 Test Data, Angle of Attack: $2 \leq \alpha \leq 8$.

	MAE	MAPE, %
c_l	0.0072	3.6
c_d	0.0008	1.3
c_m	0.0020	3.4

The associated coefficient predictions and OVERFLOW test data for a Mach number of 0.25 and Reynolds number of one million are shown in Figure 4, with the angle of attack range of particular interest highlighted in green. The predictions overlay the CFD data almost perfectly across the entire angle of attack range and especially the highlighted region that is likely encountered in nominal rotor operation.

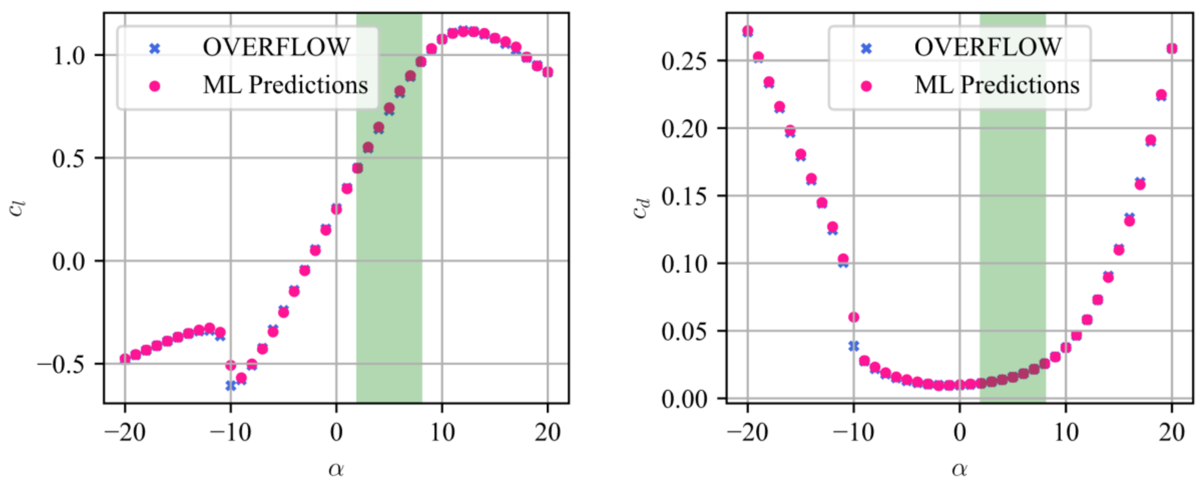

Figure 4. Multi-output Model Predictions on NACA 3415 Test Data, Mach Number = 0.25, Reynolds Number = 10^6, c_l (left), c_d (right).

To further characterize the predictive accuracy in the context of rotor optimization problems, Figure 5 displays parity plots of the multi-output model prediction of the maximum lift-to-drag ratio and minimum drag. These comparisons are plotted for all Mach and Reynolds numbers included in the PALMO dataset, as reported in Table 1. The predictions fall nearly on top of each other with an MAPE of 1.5% for the maximum lift to drag ratio and an MAPE of 2.5% for the minimum drag coefficient. These numbers quantify the predictive performance of the surrogate model at the likely operating points of interest for the airfoil. Accuracy within 1-3% for an airfoil unseen in the training data is a strong correlation suitable for conceptual and preliminary design.

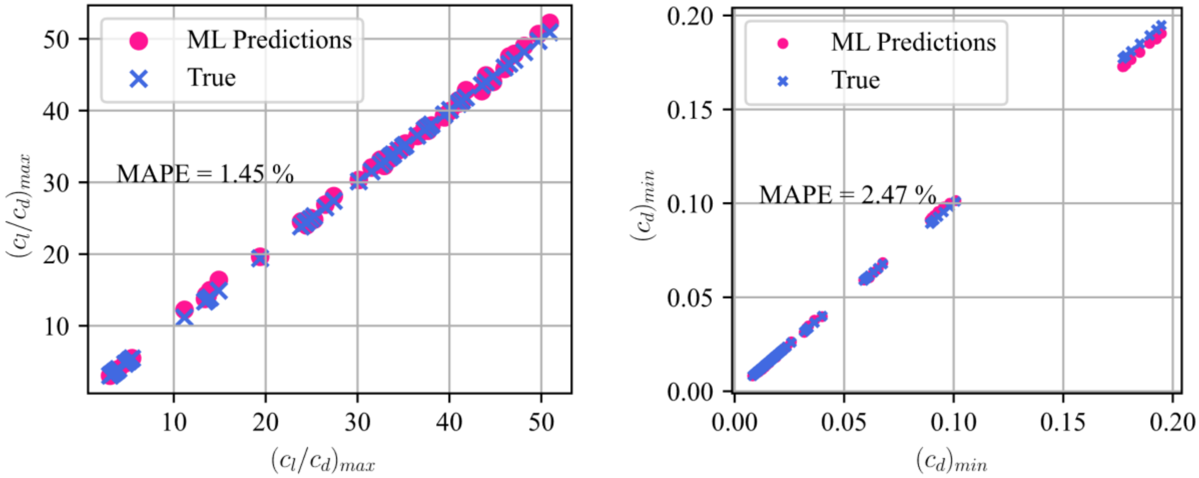

Figure 5. Multi-output Model Predictions of NACA 3415 Airfoil Performance. $(c_l/c_d)_{max}$ (left) and $c_{d,min}$ (right).

IV. PALMO Accuracy in Downstream Use Case – Rotorcraft CFD

This section demonstrates the downstream accuracy of aggregate rotor performance metrics when using airfoil lookup tables generated with the previously discussed surrogate models. A generic rotor model was created using OVERFLOW's actuator disk implementation to compare rotor performance results when using the CFD airfoil performance lookup tables versus tables generated from the surrogate models. For rotorcraft applications, the airfoil performance lookup tables are formatted as C81 tables. A rotor was created with radius of 5 meters, a constant chord of 0.25 meters, and blade-tip Mach number of 0.6. The C81 airfoil performance lookup tables were created for both the NACA 0012 and NACA 3415 profiles using 1) the actual data from the raw OVERFLOW simulations, and 2) data predicted by the surrogate model. The tables use Mach numbers from the database and the closest Reynolds number values to the true chord-based Reynolds number at each Mach number. This was done to isolate the impact on predictive error to only the error between the direct CFD and predicted values. Since the NACA 0012 airfoil was included in the surrogate model training dataset, the lookup tables are expected to be in close agreement. The NACA 3415 airfoil has a unique camber and thickness, neither of which is explicitly included in the training dataset. The effect of camber and thickness, however, is indirectly accounted for in the training database and this is thus a true demonstration for a PALMO use case. A sample forward flight simulation of the OVERFLOW rotor model is shown in Figure 6.

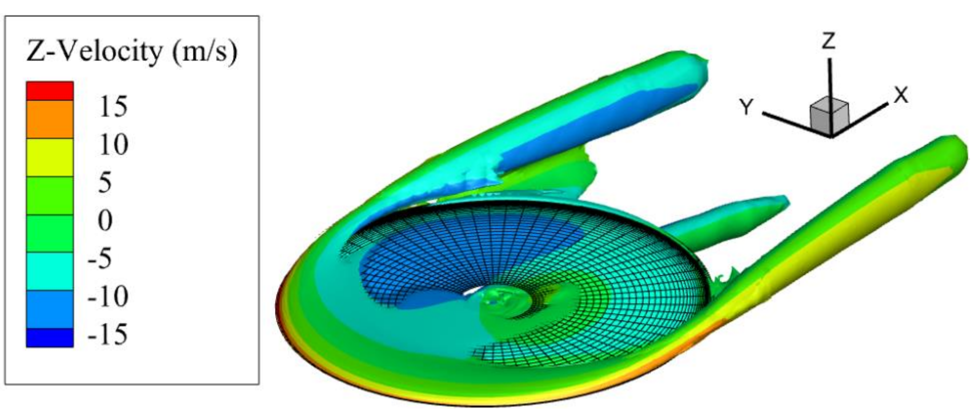

Figure 6. CFD Simulation of OVERFLOW Actuator Disk Model in Forward Flight, Advance Ratio 0.17, Rotor Shaft Angle 0 deg, Rotor Collective 4 deg.

Rotor performance metrics are first compared using the NACA 0012 airfoil. Plots of rotor Figure of Merit (FM) and thrust coefficient versus torque coefficient are reported in Figure 7. The dashed red line is using C81 tables from the direct OVERFLOW data while the solid black line is using data from the surrogate models. The two lines are essentially indistinguishable from each other. Forward flight values are reported in Table 4. The maximum calculated discrepancy in rotor effective lift to drag ratio is 0.45% and occurs at a rotor collective setting of 6 degrees. Still, the performance is predicted within 1% over the entire range.

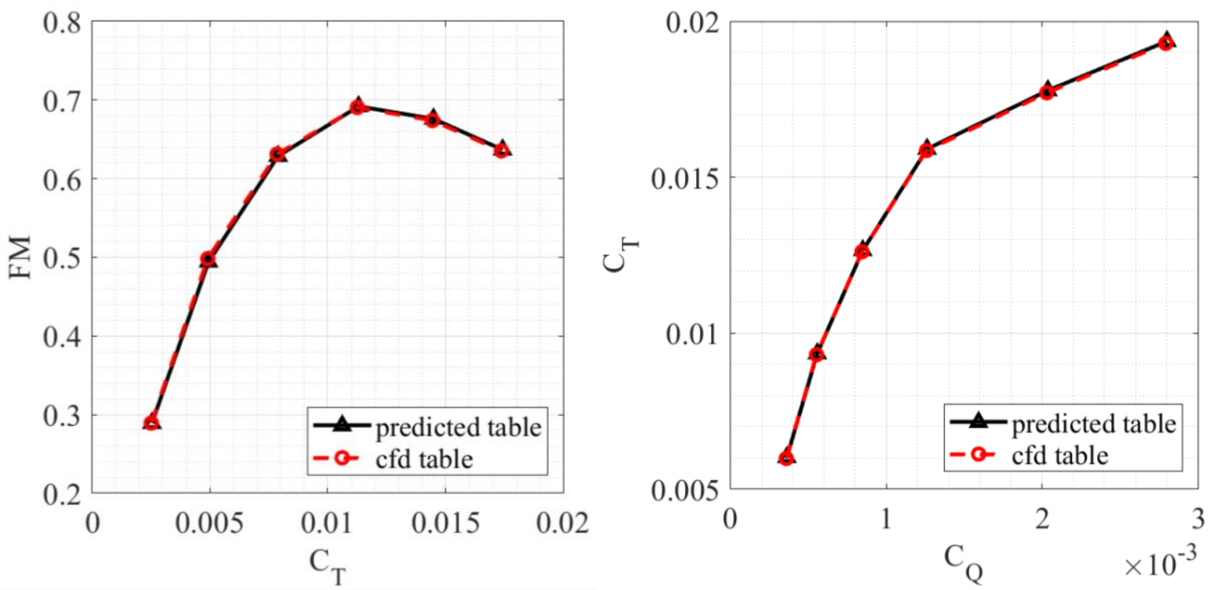

Figure 7. NACA 0012 Rotor CFD vs PALMO Predicted C81 Tables, Hover (Left), Forward Flight (Right).

Table 4. NACA 0012 Forward Flight Comparison, Advance Ratio 0.17, Rotor Shaft Angle 0 deg.

Collective [deg]	C_T Predicted	C_T CFD	C_Q Predicted	C_Q CFD	LDE Predicted	LDE CFD	LDE %Error
4	0.0060	0.0060	0.00036	0.00036	1.66	1.67	0.40
6	0.0094	0.0093	0.00056	0.00055	1.68	1.69	0.45
8	0.0127	0.0126	0.00085	0.00084	1.49	1.49	0.25
10	0.0159	0.0159	0.00126	0.00126	1.26	1.26	-0.03
12	0.0178	0.0177	0.00204	0.00203	0.87	0.87	-0.07
14	0.0194	0.0193	0.00280	0.00279	0.69	0.69	-0.07

The more meaningful test of the PALMO methodology is a comparison of rotor predictions using the NACA 3415 airfoil, since that data was not included in the surrogate model training data. Equivalent comparisons for the held-out NACA 3415 airfoil are plotted in Figure 8 and the forward flight metrics are reported in Table 5. The comparisons are again very close, and the effective rotor lift to drag ratio is predicted within 2.1% over the entire rotor collective range. The accuracy is observed to improve with increasing collective values, which appears to be a result of the same approximate absolute errors with increasing predicted values. The result suggests that any arbitrary NACA 4-series airfoil within the bounds of the database can now be modeled with accuracy similar to the CFD predictions.

Actuator rotor disk data was pulled from the NACA 3415 hover simulations to take a closer look at radial distributions of some of the metrics used in calculating rotor thrust and torque. The radial values are an azimuthal average over the entire rotor disk. The rotor blade angle of attack is plotted in Figure 9 for two values of rotor collective. The predicted and true values are nearly indistinguishable. In the following plots, “true” is used to denote the direct OVERFLOW CFD data as ‘truth’ data. The “Predicted” values are derived from the surrogate models attempting to predict what would be the OVERFLOW CFD truth data.

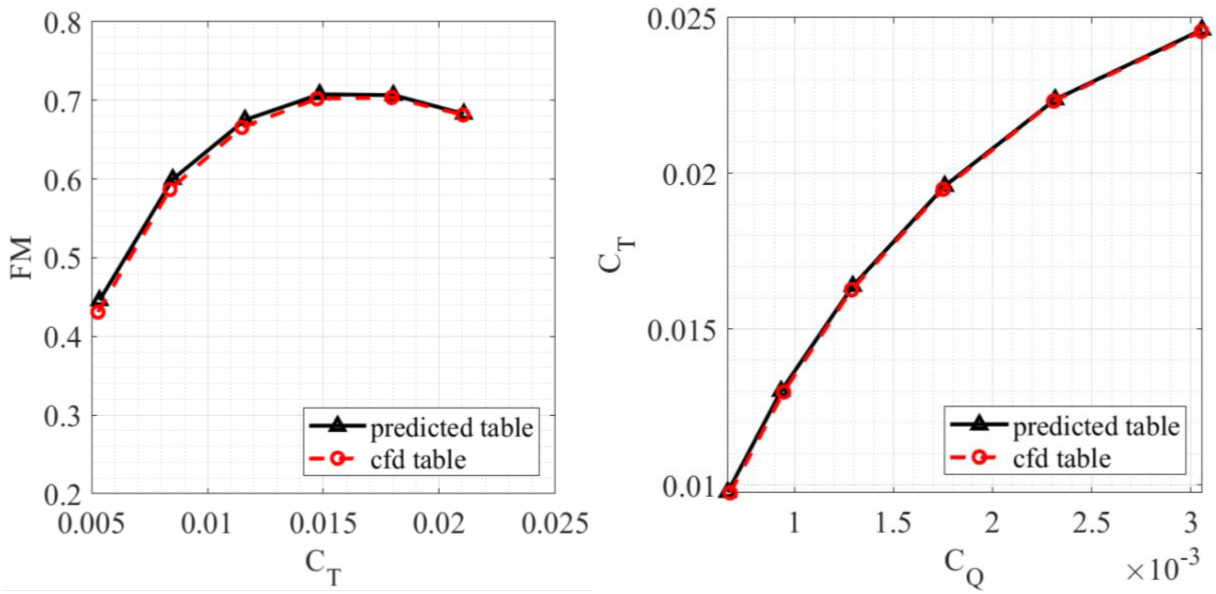

Figure 8. NACA 3415 Rotor CFD vs PALMO Predicted C81 Tables, Hover (Left), Forward Flight (Right).

Table 5. NACA 3415 Forward Flight Comparison, Advance Ratio 0.17, Rotor Shaft Angle 0 deg.

Collective [deg]	C_T Predicted	C_T CFD	C_Q Predicted	C_Q CFD	LDE Predicted	LDE CFD	LDE %Error
4	0.0098	0.0097	0.00067	0.00068	1.47	1.44	-2.11
6	0.0130	0.0130	0.00093	0.00095	1.39	1.37	-1.56
8	0.0164	0.0163	0.00130	0.00129	1.26	1.26	-0.23
10	0.0196	0.0195	0.00176	0.00175	1.11	1.11	-0.05
12	0.0224	0.0223	0.00232	0.00231	0.97	0.97	0.09
14	0.0246	0.0246	0.00306	0.00305	0.81	0.80	-0.12

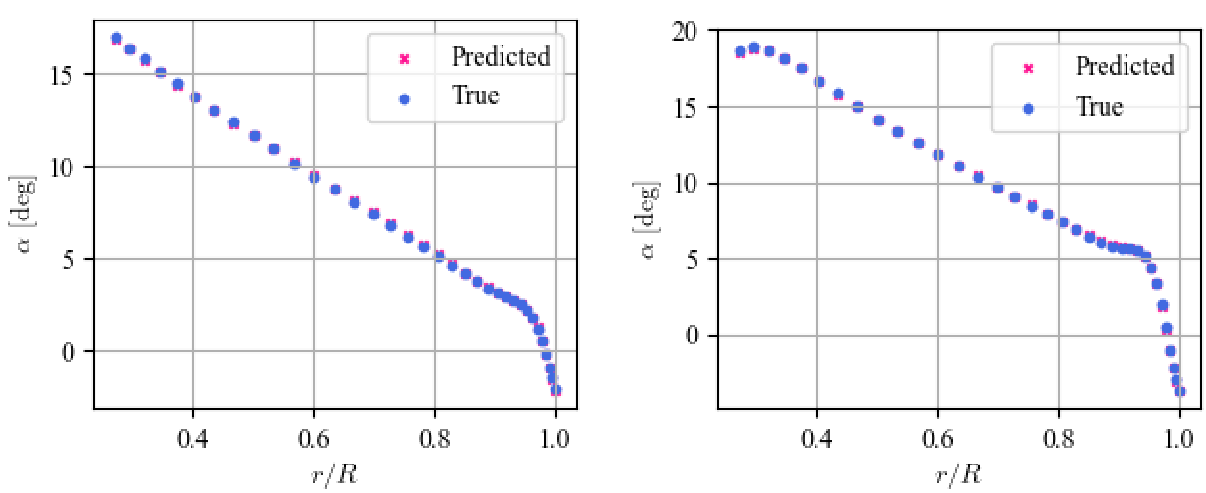

Figure 9. Angle of Attack Distribution, 6 Deg. Collective (Left), 10 Deg. Collective (Right).

Figure 10 plots the relevant lift coefficient values being pulled at each of the angles of attack reported in Figure 9. For the results to match, the actuator disk must be calculating the same inflow velocity and thus angle of attack over the rotor, in addition to the C81 tables matching at the same value of alpha. Figure 10 shows a near perfect agreement between the simulations using the true OVERFLOW C81 tables and the surrogate model predicted C81 tables. A slight discrepancy is observed at the highest angles of attack for the lower collective condition, but from Figure 9 these values are observed at the inboard radial station and thus have a small contribution to the aggregate rotor thrust and torque metrics.

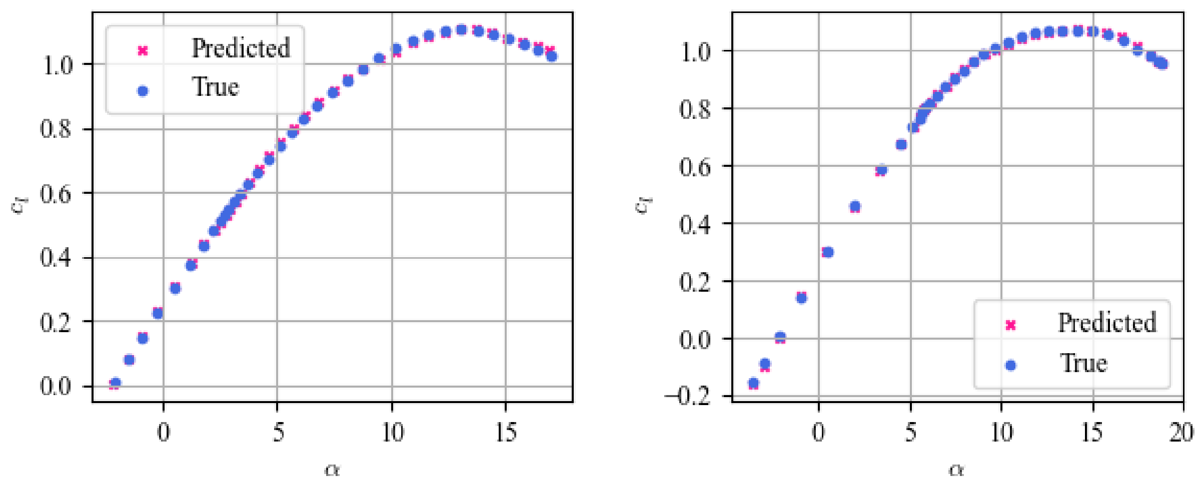

Figure 10. Lift Coefficient Comparison, 6 Deg. Collective (Left), 10 Deg. Collective (Right).

The azimuthally averaged lift coefficient is plotted in Figure 11 versus the normalized rotor blade radius. The slight deviation at the lower collective setting's most inboard station is observed, but otherwise the predicted and true tables have great agreement. Figure 12 further quantifies this agreement by plotting the mean absolute percent error of the rotor mean lift coefficient. The largest discrepancy is roughly 3% at 8 degrees of collective.

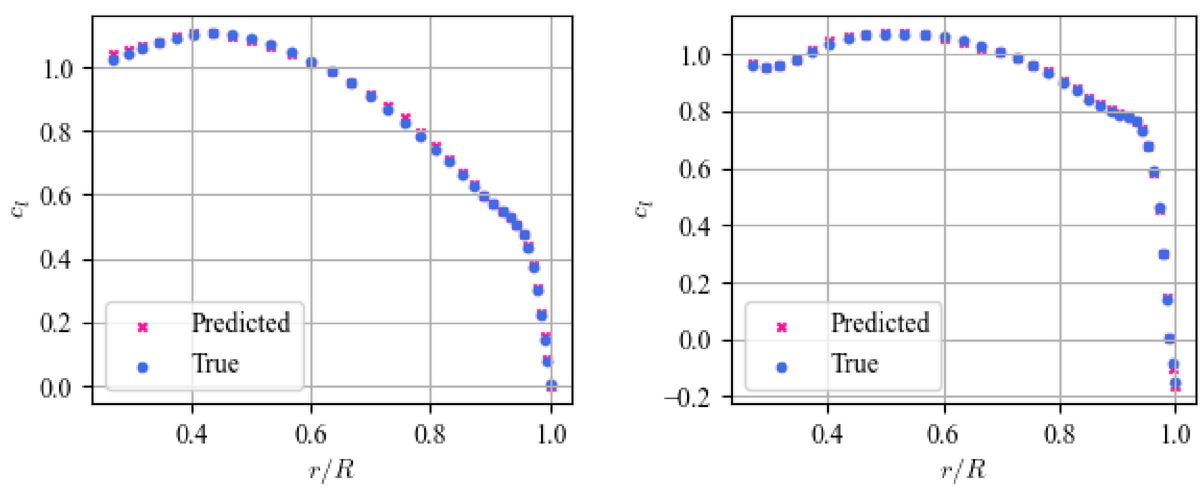

Figure 11. Lift Coefficient Distribution, 6 Deg. Collective (Left), 10 Deg. Collective (Right).

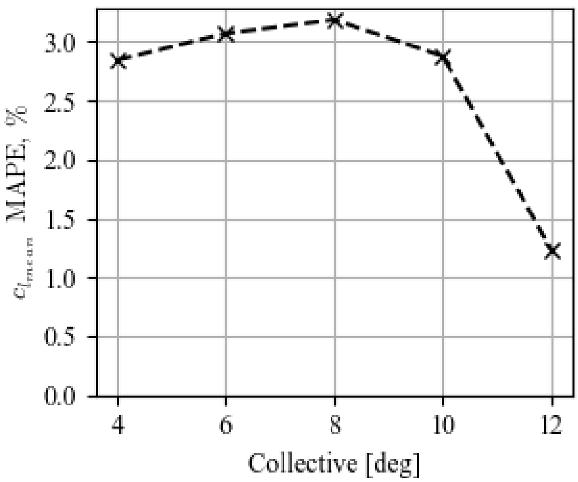

Figure 12. Mean Lift Coefficient Comparison, Mean Absolute Percent Error.

V. Conclusion

This work presented and discussed the PALMO database. This first release of the database, Version 1.0, includes predictions from 52,480 OVERFLOW airfoil simulations. The database spans ranges of Mach number, Reynolds number, and angle of attack relevant to many aerospace problems, especially rotorcraft applications. PALMO was created to enable surrogate model development and fast yet accurate airfoil performance prediction. This enables engineers to rapidly generate airfoil performance predictions with similar accuracy to OVERFLOW generated data for any airfoil and conditions within the bounds of the database without high performance computing.

Example surrogate models were trained and compared. Both single-output models and a multi-output model were created. The multi-output model, which is more compact, was found to outperform the single-output models. A true hyperparameter optimization was not carried out, but it is expected that such an optimization would likely yield even better surrogate model predictive performance.

Airfoil performance lookup tables were created to use in OVERFLOW actuator disk simulations. Simulations using the surrogate predicted and CFD generated tables for the NACA 0012 airfoil, which was included in the surrogate model training data, were within 1% of each other for forward flight rotor effective lift to drag calculations. For the NACA 3415 airfoil, which had no common thickness or camber with the training data, the surrogate predicted and CFD generated tables were within 2.1% of each other in the forward flight lift to drag metric. This suggests that performance tables generated for airfoils within the bounds of the PALMO database will yield aggregate rotor performance predictions on par with tables generated from directly running OVERFLOW airfoil calculations. The prediction speed using the surrogate models is negligible, on the order of microseconds depending on specific computing hardware, which makes it a great tool for applications such as real-time closed loop simulation and rotor design optimization.

Acknowledgments

The authors would like to thank Witold Koning and Kristen Kallstrom for their assistance in brainstorming early ideas related to the creation of the PALMO database. The authors also thank Ethan Romander for his assistance in using the NASA High-End Compute Capability computational resources. Also, the authors thank Stephen Wright, Allen Ruan, and Sesi Kottapalli for their feedback and review of this work. Lastly, the authors would like to express their gratitude to William Warmbrodt, Larry Hogle, Christopher Silva, and Wayne Johnson for their continued support of this work. This work was supported by the NASA Revolutionary Vertical Lift Technology Project.

References

[1] Drela, M., "XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils," *Low Reynolds Number Aerodynamics*, Springer, Berlin, Heidelberg, 1989, pp 1-12. https://web.mit.edu/drela/Public/web/xfoil/

[2] MSES Multielement Airfoil Design/Analysis System, Software Package, Build: 3.11, Massachusetts Institute of Technology, 2015. https://web.mit.edu/drela/Public/web/mses/

[3] Patt, D., and Youngren, H., "Improving Blade-Element Design Methods for High Speed Proprotors," *AHS Specialists' Conference on Aeromechanics*, San Francisco, CA, Jan. 2010.

[4] Cornelius, J., and Schmitz, S., "Single vs Coaxial Rigid Rotors," *VFS Transformative Vertical Flight Conference*, San Jose, CA 2022.

[5] Sridharan, A., and 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

[6] Selig, M., "UIUC Airfoils Coordinate Database." https://m-selig.ae.illinois.edu/ads/coord_database.html

[7] Shalul, H., Govindarajanl, B., Sridharan, A., and Singh, R., "Blade Shape Optimization of Rotor Using Neural Networks," *VFS Forum 79*, May 2023. https://doi.org/10.4050/F-0079-2023-18006

[8] C81 Airfoil Table Generator using ARC2D (Efficient Two-Dimensional Solution Methods for The Navier-Stokes Equations), Sukra Helitek, Inc., Software Package, Build: July 2018, Ames, IA. http://sukra-helitek.com/

[9] Allen, L., Lim, J., Haehnel, R., and Dettwiller, I., "Rotor Blade Design Framework for Airfoil Shape Optimization with Performance Considerations," *AIAA SciTech Forum*, AIAA 2021-0068, Jan. 2021. https://doi.org/10.2514/6.2021-0068

[10] Ruh, M., Liu, X., Yu, R., and 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., and 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., and 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., and 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] Loftin, L., and 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

[19] Noonan, K., and Bingham, G., "Two-Dimensional Aerodynamic Characteristics of Several Rotorcraft Airfoils at Mach Numbers from 0.35 to 0.90," - NASA TM-X-73990, Jan. 1977.

https://ntrs.nasa.gov/citations/19770008056
[20] Bingham, G., and Noonan, K., "Two-Dimensional Aerodynamic Characteristics of Three Rotorcraft Airfoils at Mach Numbers from 0.35 to 0.90, NASA TP-2000, May 1982. https://ntrs.nasa.gov/citations/19840020674
[21] Noonan, K., "Aerodynamic Characteristics of Two Rotorcraft Airfoils Designed for Application to the Inboard Region of a Main Rotor Blade," NASA TN-3009, Jul. 1990. https://ntrs.nasa.gov/citations/19900014923
[22] Noonan, K., "Aerodynamic Characteristics of a Rotorcraft Airfoil Designed for the Tip Region of a Main Rotor Blade," NASA TM-4264, May 1991. https://ntrs.nasa.gov/citations/19930020261
[23] 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
[24] Spalart, P., and 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
[25] AFTGen Airfoil Table Generator, Software Package, Build: 0.10.2, 2023. http://sukra-helitek.com/
[26] Koning, W., Johnson, W., and Grip, H., "Improved Mars Helicopter Aerodynamic Rotor Model for Comprehensive Analyses," *AIAA Journal*, Vol. 57, No. 9, Sept. 2019. https://doi.org/10.2514/1.J058045
[27] Koning, W., "Generation of Performance Model for the Aeolian Wind Tunnel (AWT) Rotor at Reduced Pressure," NASA/CR–2018–219737, published online Dec. 2018. https://ntrs.nasa.gov/citations/20180008699
[28] Kallstrom, K., "Exploring Airfoil Table Generation using XFOIL and OVERFLOW," *VFS Transformative Vertical Flight Conference*, San Jose, CA 2022.
[29] Abbott, I., Von Doenhoff, A., Stivers, S., "Summary of Airfoil Data," NACA Report No. 824, 1945. https://ntrs.nasa.gov/citations/19930090976
[30] Abbott, I., and Von Doenhoff, A., "*Theory of Wing Sections*," Dover Publications, 1959. ISBN: 9780486605869.
[31] Cornelius, J., "PALMO: An OVERFLOW Machine Learning Airfoil Performance Database Version 1.0 NACA 4-Series," NASA/TM-20240014546, 2024.