Proceedings of the 1st International Symposium on AI and Fluids Mechanics
Paper No S10 P3

Rotorcraft full state-space AI-accelerated predictions of aerodynamic coefficients

S G Paquette, E Laurendeau and D Vidal

Department of Mechanical Engineering, Polytechnique Montréal, 2500 chemin de Polytechnique, Montréal, QC, H3T 1J4, Canada

Abstract. Aerodynamic coefficients, such as the lift, drag, and moment coefficients, are essential in the analysis and design optimization procedures of aircraft. These coefficients are traditionally obtained numerically through conventional Computational Fluid Dynamics (CFD) simulations that are computationally expensive and require several complex workflows (i.e., geometry parameterization, mesh generation, flow solver, and post-processing) that may not be fully automated. AI-accelerated surrogate modeling aims at reducing computational costs and reducing or removing complex workflows at the cost of eventually lower but acceptable engineering accuracy. In the present work, deep learning is employed in the development of a surrogate model capable of full state-space predictions of 2D rotor section aerodynamic coefficients under standard, reversed, deep stall, and transonic rotorcraft flow conditions. Latin Hypercube Sampling (LHS) is used to configure a suitable design space composed of roughly 16,000 unique combinations of flow conditions and Non-Uniform Rational Basis Spline (NURBS) generated airfoils. The flow conditions are sampled across the following ranges: angle of attack from -180° to 180°, Mach number from 0.2 to 0.9, and Reynolds number from 2.5×10^6 to 20.0×10^6. Data points are generated using high-performance computing resources and the CHAMPS finite volume CFD solver configured with Reynolds-Averaged Navier-Stokes (RANS) modeling. Hyperparameter meta-optimization techniques are used to find the Pareto front between prediction performance and model complexity objectives. Optimized data-driven models derived from this framework accelerate predictions by a factor of 0.1 to 1×10^9 compared to conventional CFD computations, mapping the full solution space with an $R^2 > 0.99$ on a chosen testing dataset. The paper also addresses the method's limitations on more complex cases.

Keywords: Aerodynamic Coefficients, Rotorcraft, Deep Learning, Multi-Objective Optimization, Meta-Optimization

1. Context

Medium-fidelity aerodynamic simulations used in early-stage concept design and flight simulators often require the usage of lookup tables [1][2]. This is especially the case for rotorcraft simulations, where tables often formatted in the C81 table style, are composed of a 2D geometry's lift (C_L), drag (C_D), and moment (C_M) coefficients across full-factorial combination ranges for angle of attack (α), Mach number (M), and Reynolds number (Re) [3]. Consequently, a database of lookup tables is necessary in cases where multiple complex geometries are present. For example, the Non-Linear Unsteady Vortex Lattice-Vortex Method (NL-UVLM) [3] is a medium-fidelity method that makes use of lookup table databases. Parametric studies conducted on NL-UVLM have highlighted the importance of higher-fidelity lookup table data, observing significant differences in accuracy between predictions made with low- and high-fidelity databases [4]. While the described medium-fidelity simulations provide rapid and inexpensive predictions of rotorcraft aerodynamics, the generation of their required high-quality lookup table databases requires strenuous management of complex workflows and large computational cost. Researchers can, however, mitigate immense computational expenses by relying on linear interpolation between lookup table data points at the cost of the solution's degradation. This, however, is not possible in the context of aerodynamic design optimization, where geometry optimization processes nullify the usability of previously generated lookup tables which are each respectively associated with a unique geometry.

Surrogate-models find their utility in contexts where computational resources are a limiting factor, and the required lookup tables for medium-fidelity rotorcraft simulation span a large domain of flow boundary conditions and geometries. Recently, several machine learning techniques have been used to accomplish this task in rotorcraft surrogate modeling. Sridharan and Sinsay [5] used 70021 CFD-generated data points to map a response surface with Gaussian Process Regression (GPR). This response surface was capable of the prediction of aerodynamic coefficients over various geometries, α values (-5°~15°), and M values (0.3~0.9). Similarly, Cornelius and Schmitz [6] used approximately 62000 high-fidelity data points to train a Neural Network (NN) capable of the prediction of aerodynamic coefficients over various geometries, α values (-20°~20°), M values (0.25~0.9), and Re values (75k~8M). GPR and NNs were shown to successfully map these studies' respective solution domains and allowed for an immense reduction of computational cost linked to rotorcraft concept design optimization. The development of these surrogate models, however, required very large datasets represented by full-factorial Design of Experiments (DoE).

These studies' databases, however, span small dimensional spaces and also make use of a full-factorial DoE definition. The resulting extreme data density used in the development of the described surrogate models do not leverage the full interpolative capabilities of Machine Learning (ML) algorithms. While full-factorial DoE might be feasible for problems with few dimensions, its complexity grows exponentially with the number of dimensions. For example, a full-factorial experiment with a constant number of levels l per dimension N will have a complexity of $O(l^N)$, making it unfeasible for complex simulation cases with a great number of dimensions.

The current work aims to build upon previous attempts at surrogate modeling for rotorcraft. This study broadens the range of flow conditions to attempt to account for the entire spectrum of conditions a rotor might experience. Rotorcraft lookup tables often span across α ranges of -180°~180°. Furthermore, this study extends to larger Reynolds numbers, exceeding $Re > 20\times10^6$ in extreme cases near rotor tips. The presented model is trained using Deep Learning (DL), a method that has been successfully applied in the development of NN models with Multilayer Perceptron (MLP) architectures for aerodynamic coefficient prediction [7]. The generated CFD-database used in DL is composed of sparse data points to avoid incurring massive computational expenses. Latin Hypercube Sampling (LHS), a form of DoE that generates sparse and uniform sampling intervals, has been shown to be effective in cost-efficient airfoil surrogate modeling [8], and is used in the current work.

2. Problem Domain

The model is trained on sparse data spanning greatly varying geometries and large ranges for flow boundary conditions as described in Table 1. These ranges are meant to represent the full state-space of possible geometries, and standard, reversed, deep stall, and transonic flow conditions that might be present for an earthbound rotorcraft's blade-sections. The α range accounts for situations of hovering, forward flight, and reversed flow, while the range for M is meant to represent the range of possible Mach numbers resulting from the combination of wind and rotor velocity vectors. The angle of attack will vary with the resulting Mach number combination. Furthermore, Re is calculated as a function of M, temperature (T), and density (ρ), where Sutherland's law, the barometric formula, and the ideal gas law are used to close the presented multi-variable system. Hence, realistic values for T and ρ are selected based on a sweep of low-earth atmosphere. By coupling M with α and constraining Re to values of T and ρ found in low-earth atmosphere, the problem domain is restricted to realistic combinations of variables, consequently eliminating the generation of unnecessary data points and saving computational resources during CFD database generation and ML training. Sets of flow condition variables, α, M, and Re, are sampled with an LHS DoE and are each paired with a unique geometry.

Table 1. Rotor geometry & boundary flow condition domain.

Variable	Range
Geometry	16541 combinations of NURBS controls points
Angle of attack	$\alpha \in$ [-180, 180] °
Mach number	$M \in [0.2, 0.9]$, as a function of α
Reynolds number	$Re \in [2.5, 20.0] \times 10^6$, as a function of M

The geometrical space is covered by 16541 geometries that are generated with NURBS using unique combinations of 9 control points. The control points' weights are fixed for better control of the resulting generated geometries, while their Cartesian coordinates are varied according to an LHS DoE. Three sampled geometries are illustrated in Figure 1.

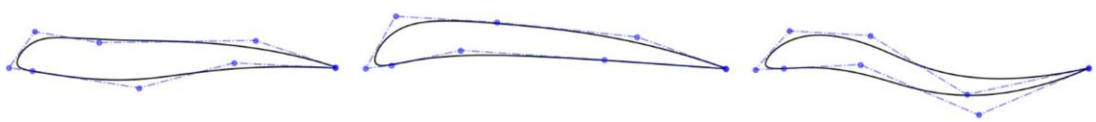

Figure 1. Examples of airfoils generated with NURBS control points.

3. Methodology

The methodology for the development of the ML surrogate model is illustrated in Figure 2. The workflow necessary in the development of the ML surrogate model is highly complex, involving automated meshing, flow solver input generation, CFD simulations, bifurcation analysis, heuristic ML training, and hyperparameter meta-optimization.

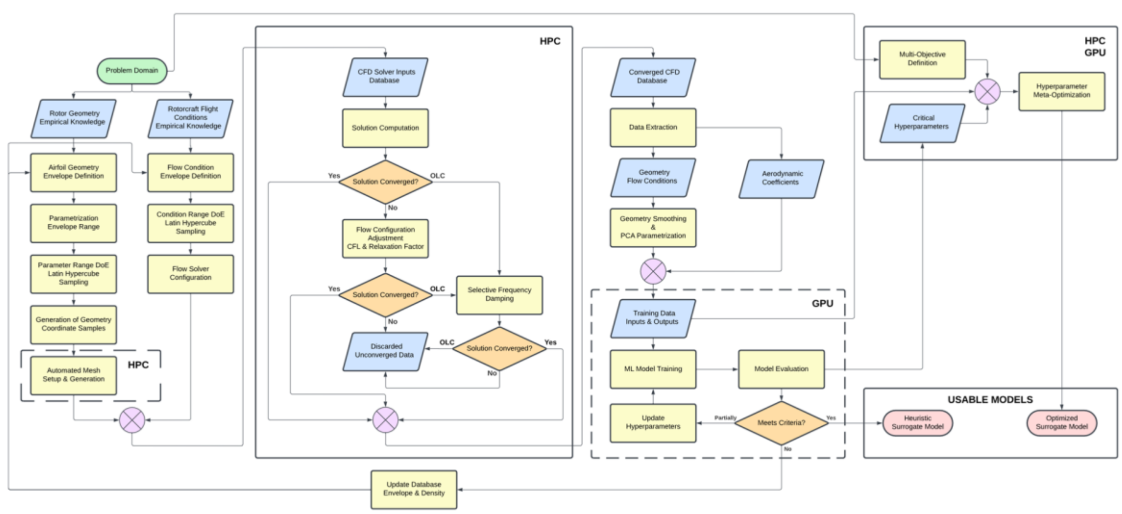

Figure 2. Machine learning surrogate model development workflow (OLC = Outer Limit Cycle).

3.1. *CFD Workflow*

Once the design space is defined, LHS is used to split the variable spaces into a uniform number of intervals to permit efficient sampling. The inputs of 16541 CFD simulation cases are subsequently generated through an automated process consisting of mesh setup, mesh generation, and flow solver input configuration. The meshes are each composed of 240×120 cells to provide an appropriate balance between grid resolution and computation time. The meshes share the same minimum wall distance of $\Delta s = 1\times10^{-6}$ chord, permitting a minimum nondimensionalized wall distance $y^+ < 1$ throughout the entire domain. Accordingly, the finite volume CFD solver CHAMPS [9] is configured conservatively to allow for convergence in the most extreme of cases present in the problem domain. The flow solver is paired with the Spalart-Allmaras [10] RANS turbulence model. Bifurcation procedures can be utilized in cases where the implemented aerodynamic force Cauchy-windowing convergence criterion is not met. The solver's Courant-Friedrich-Lewy (CFL) number and relaxation factors are adjusted in divergent simulations. In the case of detected residual Outer Limit Cycles (OLC), Selective Frequency Damping (SFD) [11] is added to permit convergence. Simulations that do not respect the Cauchy-windowing criterion and an additional ρ-residual < 10^{-2} criterion are rejected. 734 of the initial 16541 cases were pruned, while the remainder are used in ML development. Similarly, 71 of 1098 cases were pruned in a supplementary testing dataset composed of the NACA 0012, NACA 2416, and NACA 4608 airfoils.

3.2. Machine Learning Development

The machine learning development workflow begins with the raw converged CFD simulations data. Parameterization is required as the number of degrees of freedom associated with the geometry Cartesian coordinates is too large. The combination of a Cartesian coordinate smoothing for airfoil geometry generalization and Principal Component Analysis (PCA) for dimensionality reduction is found to be an effective solution. The smoothing procedure consists of fitting the geometries' initial coordinates to cubic splines that are subsequently used in the generation of updated geometries that follow a consistent distribution and number of coordinates across all cases. The CFD data is then split into a training, validation, and testing datasets. The smoothing operation is used in addition to the PCA to ensure generality as the PCA cannot accurately represent geometries with differing coordinate distributions and number of coordinates.

The surrogate model is developed to predict C_L, C_D, and C_M all at once. The developed surrogate model makes use of MLP architecture with fully connected dense linear layers and Exponential Linear Unit (ELU) activation functions. *Z*-score and min-max normalization are respectively used on model inputs and outputs. During training, several training parameters are optimized heuristically as a function of validation performance. Heuristic optimization permits the identification of key hyperparameters. These key hyperparameters are further optimized with black-box meta-optimization tools using the Tree-structured Parzen Estimator (TPE) algorithm [12].

4. Results

The meta-optimization procedure allows for the determination of the Pareto front between the number of model inputs, the model's number of neural connections represented by a scaled complexity objective, and the model's validation loss represented by a scaled loss objective. The TPE algorithm minimizes the weighted sum of the loss and complexity objectives, with equal weights assigned to both in this instance. The resulting front, composed of the models with the lowest combined objectives, is shown in figure 3.

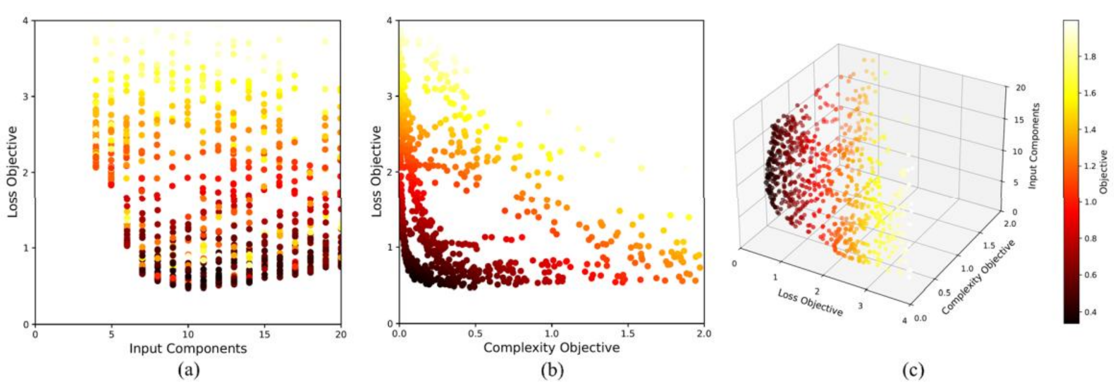

Figure 3. Multi-objective Pareto front between validation loss objective and (a) input components, (b) complexity objective, (c) input components & complexity objective.

The optimized model configured with 10 input components, 3 hidden layers, and 256 neurons per hidden layers can now be evaluated on testing datasets. The predictive performance of this model configuration is evaluated on two test datasets as described in Table 2. The first dataset is composed of a 10% random sampling split from the global dataset produced from the previously described workflow. The second dataset is composed NACA 0012, NACA 2416, and NACA 4608 geometries across the entire range of boundary flow conditions. The second dataset is used to evaluate generalization performance as no NACA geometries are used in the model's training or validation. The model's generalization is tested in this instance as the training dataset's geometries and NACA geometries are generated very differently.

Table 2. Evaluation of predictive performance on random split & NACA testing datasets.

	C_L R^2, MAE, σ	C_D R^2, MAE, σ	C_M R^2, MAE, σ
10% Random Split from Global Dataset	0.996, 0.0302, 0.0479	0.999, 0.0115, 0.0163	0.996, 0.0128, 0.0211
NACA 0012, 2416, 4608 Testing Dataset	0.993, 0.0420, 0.0615	0.998, 0.0177, 0.0218	0.990, 0.0233, 0.0305

Unlike the first dataset, the NACA dataset contains multiple α for each flow condition. Aerodynamic coefficient curves can then be visualized. Figure 4 depicts a typical example of coefficient curves. The optimized ML model follows the tendencies found in the CHAMPS solver's CFD simulation data but does not predict the same stall phenomena near max C_L. However, there is also a large uncertainty classically associated with CFD simulations in this highly non-linear region near the onset of stall.

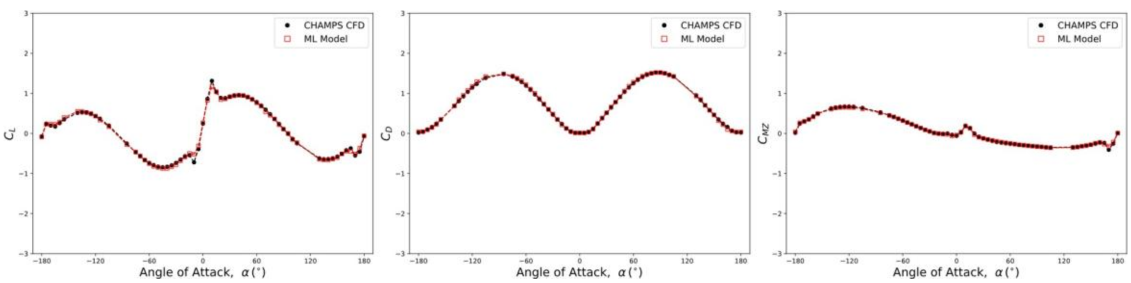

Figure 4. Typical example of CFD and ML model predictions for NACA 2416 aerodynamic performance coefficient curves at $M = 0.5$ and $Re = 11.6\times10^6$ for (a) C_L, (b) C_D, and (c) C_M.

The developed surrogate models represent an immense prediction acceleration with respect to conventional CFD simulations. The acceleration factors are shown in Figure 5, where for example, the optimized model that has 3 hidden layers and 256 neurons per layer will accelerate predictions by a factor of 0.54×10^9. In this case, 16.8 and 2.3×10^{-8} core-years were required to predict the entire training database using conventional CFD and the developed surrogate ML model, respectively.

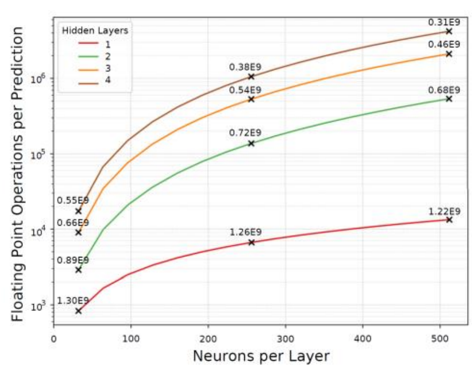

Figure 5. ML acceleration factor with respect to conventional CFD at different model complexities.

5. Conclusion

The presented study demonstrates the successful development of a machine learning surrogate model capable of predicting rotorcraft aerodynamic coefficients across nearly the entire range of possible flow conditions a rotor might experience. The limitations of previous surrogate modeling efforts were addressed by leveraging the interpolative capabilities of deep learning algorithms across complex domains decomposed with Latin Hypercube Sampling. Rigorously developed workflows permit the generation of a full state-space CFD simulation database. Cartesian coordinate smoothing and PCA are data processing techniques that were respectively found to be useful in generalization and dimensionality reduction. Hyperparameter meta-optimization was also employed to find the Pareto front between the optimal number of model inputs, model complexity, and model validation loss. The predictive performance was further evaluated on a model with an optimized set of hyperparameters, being 10 input components, 3 hidden layers, and 256 neurons per layer. The optimized model achieves significant accuracy while generalizing to unseen datasets, as evidenced by its performance on both randomly sampled and NACA testing datasets. However, challenges are seen in regions near the onset of stall. The model's ability to deliver rapid and cost-effective predictions highlights its potential utility in early-stage rotorcraft design optimization and flight simulators. This work sets the stage for further refinement of surrogate modeling in rotorcraft, where various architectures, geometry parameterization techniques and differing levels of data density have yet to be analysed on a problem domain of this scale.

6. Acknowledgements

This work was supported by CAE Inc., NSERC, and the CRIAQ *PARTENAR-IA* program. Compute resources were provided by the Digital Research Alliance of Canada and Calcul Québec.

7. References

[1] M. Parenteau, É. Laurendeau, and K. Sermeus, 'VLM coupled with 2.5d RANS sectional data for high-lift design', in *2018 AIAA Aerospace Sciences Meeting*, 2018. doi: 10.2514/6.2018-1049.
[2] S. Düssler and R. Palacios, 'Enhanced Unsteady Vortex Lattice Aerodynamics for Nonlinear Flexible Aircraft Dynamic Simulation', *AIAA Journal*, vol. 62, no. 3, pp. 1179–1194, Mar. 2024, doi: 10.2514/1.J063174.
[3] A. Cocco, A. Savino, A. Colli, P. Masarati, and A. Zanotti, 'A non-linear unsteady vortex-lattice method for rotorcraft applications', *Aeronaut. j.*, vol. 128, no. 1328, pp. 2308–2330, Oct. 2024, doi: 10.1017/aer.2024.55.
[4] V. Proulx-Cabana, G. Michon, and E. Laurendeau, 'Parametrization Effects of the Non-Linear Unsteady Vortex Method with Vortex Particle Method for Small Rotor Aerodynamics', *Fluids*, 2024, doi: 10.3390/fluids9010024.
[5] A. Sridharan and J. D. Sinsay, 'Accelerating Aerodynamic Design of Rotors using a Multi-Fidelity Approach in TORC: Tool for Optimization of Rotorcraft Concepts', in *AIAA AVIATION 2023 Forum*, San Diego, CA and Online: American Institute of Aeronautics and Astronautics, Jun. 2023. doi: 10.2514/6.2023-4305.
[6] J. Cornelius and S. Schmitz, 'Dragonfly Rotor Optimization using Machine Learning Applied to an OVERFLOW Generated Airfoil Database', in *Proceedings of the Vertical Flight Society 80th Annual Forum*, Montreal, Canada: The Vertical Flight Society, May 2024, pp. 1–20. doi: 10.4050/F-0080-2024-1316.
[7] X. Du, P. He, and J. R. R. A. Martins, 'Rapid airfoil design optimization via neural networks-based parameterization and surrogate modeling', *Aerospace Science and Technology*, vol. 113, p. 106701, Jun. 2021, doi: 10.1016/j.ast.2021.106701.
[8] K. Marepally, N. Paternostro, B. Lee, and J. D. Baeder, 'Data Puncturing and Training Strategies for Cost-Efficient Surrogate Modeling of Rotorcraft Airfoil Aerodynamics', in *AIAA SCITECH 2023 Forum*, National Harbor, MD & Online: American Institute of Aeronautics and Astronautics, Jan. 2023. doi: 10.2514/6.2023-2042.
[9] M. Parenteau, S. Bourgault-Cote, F. Plante, and E. Laurendeau, 'Development of Parallel CFD Applications on Distributed Memory with Chapel', in *2020 IEEE International Parallel and Distributed Processing Symposium Workshops*, New Orleans, LA, USA: IEEE, May 2020, pp. 651–658. doi: 10.1109/IPDPSW50202.2020.00110.
[10] P. R. Spalart and S. R. Allmaras, 'One-equation turbulence model for aerodynamic flows', *Recherche aerospatiale*, no. 1, 1994, doi: 10.2514/6.1992-439.
[11] V. Liguori, F. Plante, and É. Laurendeau, 'Implementation of an efficient selective frequency damping method in a rans solver', in *AIAA Scitech 2021 Forum*, 2021. doi: 10.2514/6.2021-0359.
[12] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, 'Algorithms for hyper-parameter optimization', in *Advances in Neural Information Processing 24 (NIPS 2011)*, 2011.