MACHINE LEARNING MODELING OF ADVANCED AIR MOBILITY ROTOR INFLOW FROM A MID-FIDELITY AERODYNAMIC SOLVER

Alaaddin Furkan Uzunkaya
Politecnico di Milano
Milano, Italy

Pierangelo Masarati
Politecnico di Milano
Milano, Italy

ABSTRACT

In this study, two different rotor configurations were analyzed using DUST, a mid-fidelity aerodynamic solver, to investigate accurate rotor inflow prediction for advanced air mobility applications. A feedforward machine learning model was trained on the computed inflow fields and then integrated into a blade element momentum theory (BEMT) solver to evaluate its predictive capability. The trained model achieved an R^2 of about 0.95 and a range-normalized root mean square error (RN-RMSE) of around 0.03 when predicting the airflow field under given test conditions. Furthermore, when these airflows were fed into the BEMT solver, the resulting rotor thrust and torque predictions reached an R^2 of approximately 0.99 and an NRMSE of about 0.04, closely matching the reference analysis data. These results suggest that machine learning-based surrogate inflow models can significantly enhance the accuracy of real-time or near-real-time aerodynamic load predictions, and potentially aid certification processes.

INTRODUCTION

Advanced Air Mobility (AAM) vehicles, including eVTOL and multirotor configurations, require aerodynamic models that combine high fidelity with real-time computational performance to support design, simulation, and certification processes. Among widely used methods, Blade Element Momentum Theory (BEMT) offers computational efficiency but relies on simplified assumptions about rotor inflow, which can limit accuracy, particularly under varying operating conditions. Accurate prediction of rotor inflow is critical, as it directly affects estimates of thrust, torque, and vehicle handling qualities. Prior research has demonstrated the potential of surrogate modeling to enhance aerodynamic predictions. In (Ref. 1), a surrogate model for a fixed-pitch, speed-controlled coaxial rotor system achieved performance prediction accuracy within 5-10% using a hybrid BEMT-URANS solver. In (Ref. 2), low-fidelity neural network models were fused with high-fidelity CFD data to refine aerodynamic predictions, while (Ref. 3) developed a parameterized surrogate database for airfoil performance using OVERFLOW CFD, reaching a maximum error of 2.1% for untrained geometries. These studies highlight the effectiveness of data-driven approaches in improving low-cost aerodynamic models. Building on this foundation, the present work develops a machine-learning-based surrogate model of rotor inflow fields, trained on mid-fidelity aerodynamic data generated by DUST (Ref. 4), an aerodynamic solver developed at Politecnico di Milano. The trained inflow model is then integrated into a BEMT solver to evaluate its impact on predicted rotor thrust and torque. By combining the speed of BEMT with a data-driven representation of inflow dynamics, this approach aims to improve the accuracy of real-time or near-real-time aerodynamic load predictions, supporting advanced rotorcraft design and certification workflows.

THEORETICAL BACKGROUND

Blade Element Momentum Theory (BEMT)

Blade Element Momentum Theory (BEMT) combines momentum theory and blade element theory to estimate the aerodynamic forces on rotors by discretizing each blade into several elements and applying local flow assumptions. The local thrust and torque are computed based on sectional aerodynamic properties, such as airfoil lift and drag coefficients, and integrated over the blade span. BEMT offers high computational efficiency, making it widely used for preliminary design, performance estimation, and simulation of rotorcraft and propeller systems (Ref. 5). However, BEMT assumes a simplified, steady inflow model and often uses uniform inflow, which can limit accuracy.

Consider the blade section in Fig. 1; θ is the built-in twist of the blade since collective and cyclic pitch is not defined for the AAM, as it is controlled by the RPM variation instead of cyclic control. u_T is the experienced tangential velocity of the air relative to the blade in the rotor plane, u_P is the perpendicular component of experienced air velocity, u_R is the radial velocity component. Resultant velocity $U = \sqrt{u_T^2 + u_P^2 + u_R^2}$ and inflow angle is given as $\phi = \tan^{-1} u_P/u_T$. α is the angle of attack experienced by the blade section, and it is defined as $\alpha = \theta - \phi$

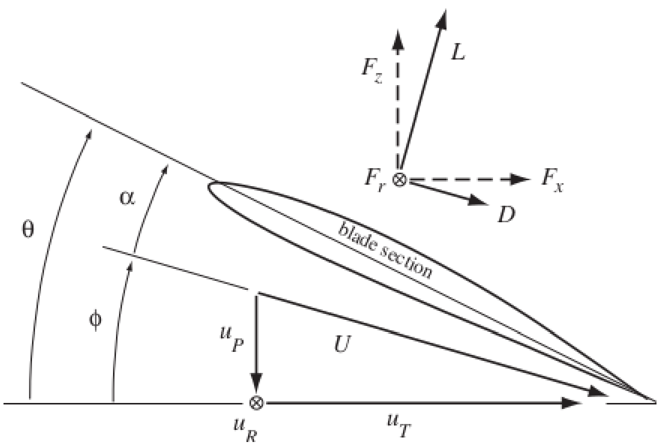

Figure 1. Blade section aerodynamics (Ref. 5).

In the figure, blade aerodynamic lift and drag forces L and D are represented as well. They are normal and parallel to the resultant velocity U, and are given in Eqs. (1) and (2), respectively.

$$L = \frac{1}{2}\rho U^2 c c_l \tag{1}$$

$$D = \frac{1}{2}\rho U^2 c c_d \tag{2}$$

where ρ is the air density, c is the chord of the blade section, c_l is the lift coefficient, and c_d is the drag coefficient of the airfoil of the blade section angle of attack, Reynolds' number, and Mach number. If these forces are decomposed with reference to the rotor plane, normal and in-plane forces are obtained, respectively, such that:

$$F_z = L\cos\phi - D\sin\phi \tag{3}$$

$$F_x = L\sin\phi + D\cos\phi \tag{4}$$

If the normal force of blade sections and the in-plane forces multiplied with the radial distance of the blade sections are summed and averaged over one revolution by the operator $\frac{1}{2\pi}\int_0^{2\pi}(...)d\psi$, it is obtained steady rotor thrust (Eq. (5)) and steady rotor torque with respect to the rotor hub (Eq. (6)), respectively.

$$T = \frac{1}{2\pi}\int_0^{2\pi}(N_b\int_0^R F_z dr)d\psi \tag{5}$$

$$Q = \frac{1}{2\pi}\int_0^{2\pi}(N_b\int_0^R rF_x dr)d\psi \tag{6}$$

Lifting Line Theory

Lifting Line Theory (LLT) represents a rotor blade or wing as a slender one-dimensional lifting line, where at each spanwise location r, the sectional aerodynamic coefficients of lift, drag, and moment – $\{c_l, c_d, c_m\}$ – are functions of the local angle of attack α, Reynolds number Re, and Mach number M:

$$\{c_l, c_d, c_m\} = f(\alpha, Re, M; \boldsymbol{r}) \tag{7}$$

The circulation distribution Γ_L(r,t) along the lifting line is then determined by solving a nonlinear problem that balances these aerodynamic coefficients with the induced velocities produced by trailing vortices. This method captures the spanwise lift distribution and induced drag more accurately than local blade element methods by including the effect of downwash. While LLT offers improved fidelity for steady, attached flow, it still relies on airfoil data and steady-flow assumptions, making it less suited to capture fully unsteady wake dynamics compared to higher fidelity vortex or CFD methods (Ref. 4).

Vortex Lattice Method

Vortex Lattice Method (VLM) extends the lifting line concept by discretizing the blade or wing surface into a mesh of lattice panels. Each panel has a bound vortex segment and sheds trailing vortices downstream. By applying the Biot-Savart law, the induced velocities from these vortices are computed, and the vortex strengths are solved to enforce flow tangency (no normal flow) at control points. VLM captures the spanwise variation of circulation and induced drag more accurately, and unlike LLT, it can represent the effect of blade planform shape and sweep. However, VLM is typically steady and linear; it assumes small angles of attack and attached flow (Ref. 6).

SCOPE AND OBJECTIVE

This study focuses on improving BEMT load predictions using machine-learned surrogate inflow models trained on mid-fidelity aerodynamic simulations. DUST was used to generate inflow velocity fields for a range of operating conditions of the rotors in Table 1.

Table 1. Rotors Data

	Rotor-A	Rotor-B
Radius (R)	0.85 m	1.1 m
Chord	0.1 m	0.2 m–0.14 m
# of Blades	3	5
Airfoil	NACA0012	NACA4406
Tip Mach # ($\Omega R/c$)	0.4–0.6	0.33–0.55
Root Cut-out	0.1 m	0.25 m
Blade Twist	15°–6°	20°–10°

Mid-Fidelity Solver DUST

DUST is a versatile aerodynamic solver built on an incompressible potential-vorticity framework that combines classical boundary-value methods with modern Lagrangian vortex-particle techniques. At its core, it decomposes the

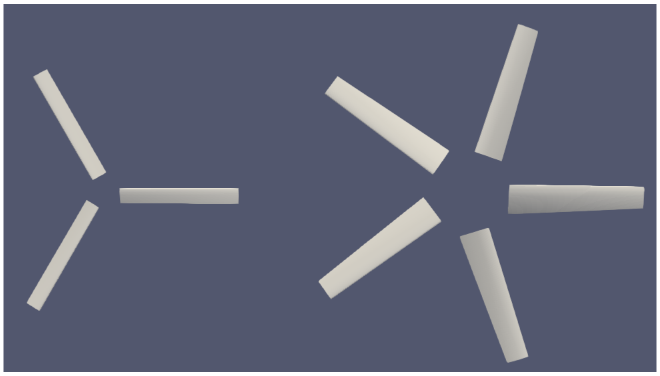

Figure 2. Rotor-A (left) and Rotor-B (right).

flow field via Helmholtz's theorem into an irrotational component—handled by surface panels and vortex-lattice elements—and a rotational component—modeled with free-vortex particles that are convected in a stable, mixed wake. Although tailored for incompressible flow, DUST can account for compressibility through the Prandtl-Glauert correction on surface panels and Mach-dependent tabulated data for lifting-line elements. Its modular architecture extends seamlessly from high-fidelity unsteady simulations to rapid steady-state computations, and an integrated post-processor offers extensive tools for flow visualization, force and moment extraction, and time-resolved analysis, making DUST a powerful platform for both research and practical aerodynamic design (Ref. 4).

Data Generation

To create the dataset required for the machine learning process, a total of 104 simulations were performed for Rotor-A: 75 simulations were allocated for training, and the remaining 29 were used for testing the model predictions. Similarly, for Rotor-B, 104 simulations were conducted, with 75 cases for training and 29 for testing. All training and testing cases are detailed in Table 2.

Table 2. Training and Test Conditions

Training Conds.	Rotor-A	Rotor-B
V_∞	0, 10, 20 m/s	0, 5, 10 m/s
V_∞ angle[a]	0°:15°:90°	0°:15°:90°
Tip Mach #	0.4:0.05:0.6	0.33:0.055:0.55
Test Conds.		
V_∞	0, 5, 15 m/s	2.5, 7.5 m/s
V_∞ angle[a]	22.5°, 52.5°, 82.5°	15°, 45°, 75°
Tip Mach #	0.425:0.05: 0.575	0.3575:0.055:0.5225

[a] In hover, angles of attack other than 0° were not analyzed.

As shown in Figure 4, a steady-state condition appears to be reached after the 7^{th} rotor revolution in terms of thrust and torque. However, to ensure convergence for all aerodynamic forces and especially for moments each simulation was extended to 10 rotor revolutions to achieve a steady-state solution confidently. The simulations were discretized such that

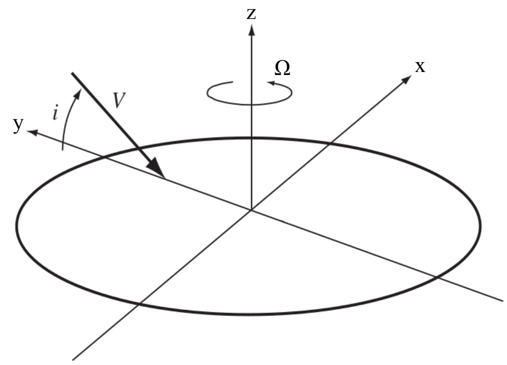

Figure 3. Angle of attack and rotor plane coordinate system.

each time step corresponded to a 2.5° degree azimuthal advance of the rotor, resulting in 144 steps per revolution. Since the rotor speed varies between the conditions, the simulation time step was adjusted in seconds accordingly. For the blade discretization, 20 uniformly distributed lifting line elements were used along the span. At each element's aerodynamic center, probe points were placed to record the local airflow velocities in the X, Y, and Z directions. The aerodynamic forces on the blades were calculated and stored for every blade azimuth angle ψ. The overall steady rotor thrust and torque were obtained by averaging the values over the final revolution.

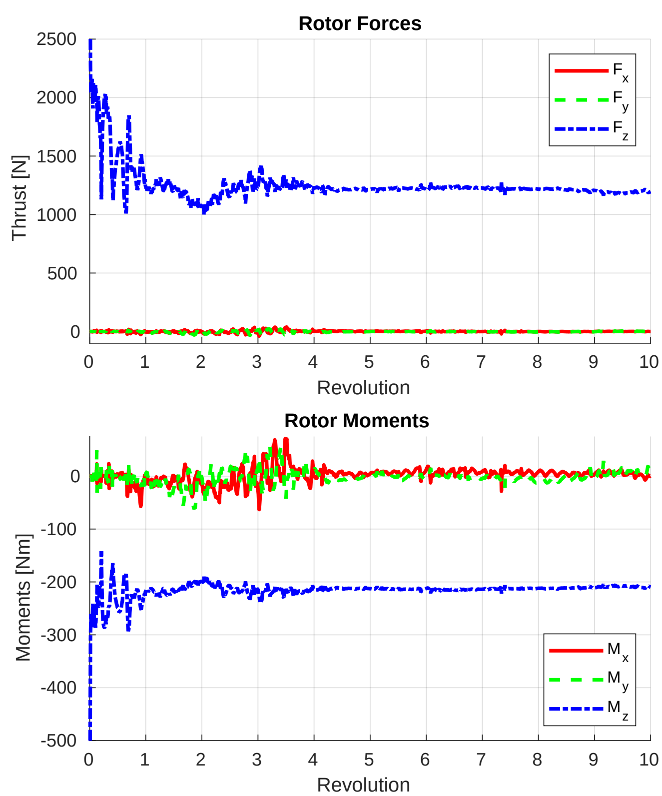

Figure 4. Rotor-B thrust and torque w.r.t. revolution.

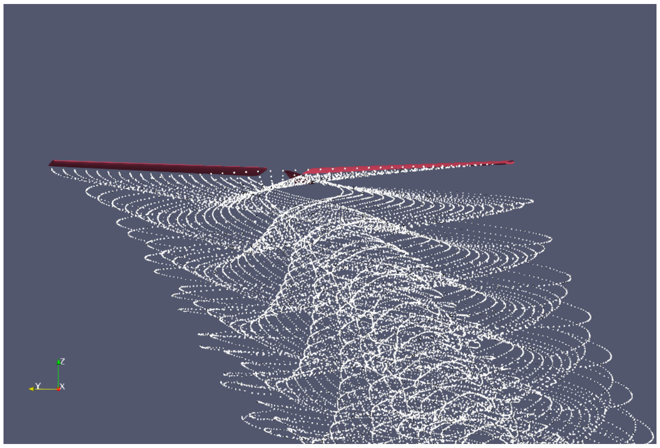

Figure 5. Rotor-A and particles @10$^{\text{th}}$ rev. for simulation case: V$_\infty$ = 10 m/s, V$_\infty$ angle = 45°, Tip Mach # = 0.45.

Machine Learning

After completing all simulations, the collected data were assembled into a structured matrix: the first column represents the free-stream velocity, the second is the angle of attack, the third is the nominal tip Mach number, and the fourth and fifth columns are the X and Y coordinates of the probe points in the rotor plane. The final three columns store the corresponding X, Y, and Z velocity components measured at the probe points. Since the X and Y velocity components contribute to the local dynamic pressure and influence force calculations, they were included as well as separate outputs to be trained and later used in the BEMT analysis.

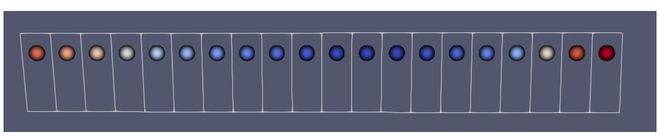

Figure 6. Rotor-A lifting line elements and probe points.

A feedforward neural network architecture with five inputs, three hidden layers with 64:32:32 fully-connected neurons, and one output was employed for the training. The five inputs are: magnitude of the free-stream velocity, angle of attack, nominal tip Mach number, and the X and Y positions of each probe point. The output is the predicted velocity component at the probe point (X, Y, or Z). To capture the full inflow and wake field, three separate models were trained, each corresponding to one of the velocity components. These trained models were then integrated into the BEMT solver to provide improved local flow predictions and enhance the fidelity of aerodynamic load calculations.

BEMT Implementation

The velocity field computed by the mid-fidelity solver DUST is incorporated into the BEMT code to enhance prediction accuracy while retaining computational efficiency suitable for real-time or near-real-time applications. Velocities are extracted at probe points positioned at the aerodynamic center (c/4 point) of each airfoil section, located at the mid-span of the lifting line and blade elements. For each time step during a rotor revolution, only those probe points that coincide with the instantaneous position of the aerodynamic center are considered; the rest are discarded. This ensures that the inflow velocities used in the BEMT calculation directly correspond to the actual airstream encountered by the blade sections over the rotation.

The Z-component of the extracted velocity field is directly used to replace the axial induced velocity term u_P. The X and Y components are incorporated by adding them to the tangential u_T and radial u_R velocity components. Using these enriched local velocity components, together with fixed lift and drag coefficients taken from (Ref. 7) airfoil dataset, the BEMT solver calculates the sectional lift and drag forces. These forces are then integrated to obtain the rotor thrust, in-plane forces, and torque. This method effectively combines detailed wake and inflow information from DUST with the computational speed of BEMT, resulting in improved aerodynamic load predictions without sacrificing runtime performance.

Tip Loss

Tip loss is a well-known aerodynamic phenomenon in rotorcraft, referring to the reduction in lift near the blade tips compared to the inboard sections. This effect is primarily caused by the formation of strong vortices shed from the blade tips, which induce a local spanwise flow that reduces the effective angle of attack and circulation at the outermost portion of the blade. The strength of these tip vortices leads to a notable decrease in lift generation over the outer 10-15% of the rotor span. Accurately accounting for tip loss is critical in rotorcraft analysis because it directly affects estimates of rotor thrust, torque, power consumption, and overall aerodynamic efficiency. Neglecting tip loss typically leads to overprediction of total lift and underestimation of required power, which can significantly impact both design calculations and flight simulation fidelity.

In traditional rotorcraft aerodynamics references such as (Ref. 5) and (Ref. 8), the Prandtl-Glauert tip loss factor is commonly applied to correct the lift distribution by reducing the lift coefficient near the blade tips. This factor gradually decreases toward the tip, capturing the reduction in circulation due to three-dimensional effects. However, in this study, direct comparison with mid-fidelity DUST results revealed that the standard tip loss formulation tends to overpredict the lift. It was observed that the onset of tip loss effects begins around the 17$^{\text{th}}$ lifting line element, corresponding to approximately 85% of the blade radius. Although some sources (Ref. 9) and (Ref. 10) recommended adjusting the tip loss factor empirically, it was found that applying synthetic, fixed tip loss factors provided a good agreement with DUST results. Specifically, for Rotor-A, the factors were set to 0.95 for the 17$^{\text{th}}$ element, 0.90 for the 18$^{\text{th}}$, 0.80 for the

19th, and 0.58 for the 20th (blade tip) element. For Rotor-B, the corresponding values were 0.88, 0.82, 0.73, and 0.54, respectively. These modified values were then used directly in the lift term calculations, improving the agreement between BEMT-predicted lift and mid-fidelity analysis data.

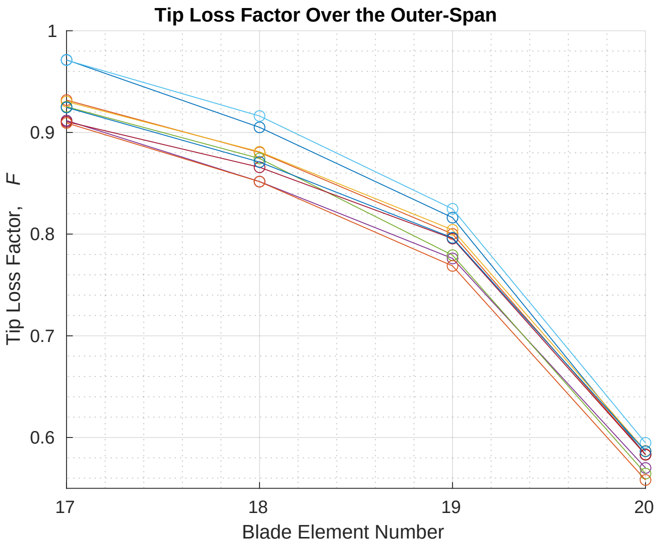

Figure 7. Synthetic tip loss factor for Rotor-A in hover state for different nominal tip Mach numbers.

Model Evaluation

In machine learning applications, especially those involving large and complex datasets, systematic and quantitative evaluation of predictive accuracy is essential. Visual inspection of prediction trends can be informative for small datasets; however, for large-scale aerodynamic simulations and surrogate modeling, standardized statistical metrics provide an objective and reproducible assessment. Since in this study the inflow field is predicted over 144 azimuthal positions (with a time step set so that the rotor blades advance by 2.5° per step) for 75 training cases and 29 test cases, this results in an enormous amount of data to evaluate. Therefore, rather than relying solely on visual inspection, consistent and quantitative metrics are necessary to objectively assess prediction quality.

In this study, two commonly used evaluation metrics are employed: the coefficient of determination (R^2) and the range-normalized root mean square error (RN-RMSE). The coefficient of determination (R^2) measures the proportion of variance in the reference (baseline) data that is captured by the model predictions. It is defined as (Ref. 11):

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2} \quad (8)$$

where y_i are the true values, $\hat{y}_i$ are the predicted values, and $\bar{y}$ is the mean of the true values. An R^2 value of 1 indicates perfect agreement between predictions and true values, while an R^2 of 0 implies that the model performs no better than simply predicting the mean. Negative values can also occur, signifying that the model's residual variance exceeds the variance of the reference data indicating that the model underperforms compared to the baseline mean.

The root mean square error (RMSE) quantifies the standard deviation of the prediction errors and is defined as:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2} \quad (9)$$

RMSE directly measures the average magnitude of prediction errors; however, since it retains the same scale as the target variable, its interpretation can be influenced by the dataset's range and units. To overcome this and enable comparison across datasets with different scales, the range-normalized RMSE (RN-RMSE) is used (Ref. 12). It normalizes RMSE by the range of the target variable:

$$\mathrm{RN-RMSE} = \frac{\mathrm{RMSE}}{y_{\max} - y_{\min}} \quad (10)$$

where y_{max} and y_{min} are the maximum and minimum values of the true data, respectively. A lower RN-RMSE value indicates better predictive accuracy and provides a scale-independent assessment, making it especially useful in surrogate modeling of aerodynamic phenomena where data ranges can vary significantly.

Together, R^2 and RN-RMSE provide complementary insights: the former reflecting the model's explanatory power relative to the data variance, and the latter quantifying the average prediction error normalized by the data range.

DISCUSSIONS

The trained machine learning model is evaluated for inflow velocities under hover conditions for both training and test cases, as shown in Figure 8. A very good agreement is observed between the DUST simulations and the predicted inflow for the training cases, as expected. For the test cases, the overall agreement remains strong; however, as the Mach number increases, discrepancies between the predicted and simulated inflow fields become more noticeable, particularly near the blade tip. In these regions, the machine learning model tends to overpredict the inflow. Due to the large number of training and testing conditions for both Rotor-A and Rotor-B, and because the inflow field becomes azimuthally dependent with increasing free-stream velocity, direct visual comparison of simulation and prediction results becomes challenging. Therefore, the accuracy of the predicted inflow, as well as the X and Y velocity components of the airflow, is assessed using the statistical methods discussed previously.

In Table 3, the coefficient of determination R^2 and RN-RMSE values of the predicted velocities are listed separately for each rotor and for each velocity component. From the table, training conditions have a good match for the analysis data for both rotors, whereas, for the test conditions, prediction quality drops a little bit, although the good agreement between predictions and the analysis continues.

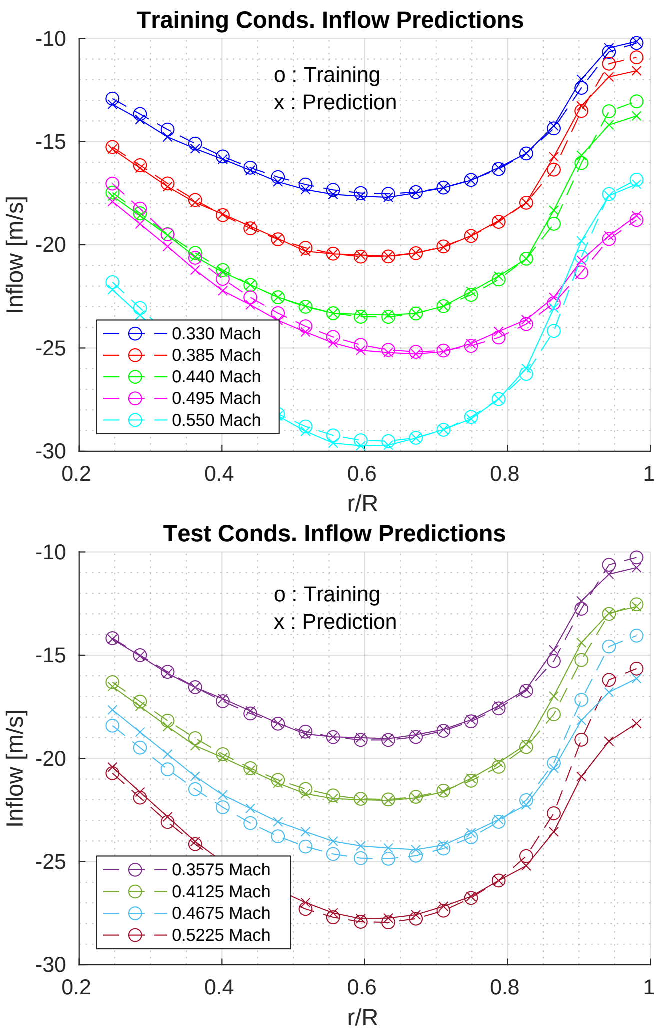

Figure 8. Predicted inflow comparisons with DUST in hover for Rotor-B

From Figure 9 to 16, Rotor-A and Rotor-B thrust and torque predictions obtained using BEMT with machine-learning-predicted inflow are compared with mid-fidelity simulation results for both training and testing conditions. In addition to the Lifting Line method, a subset of randomly selected cases has also been analyzed using the Vortex Lattice Method (VLM) to provide a cross-validation of agreement between lifting-line-based predictions and BEMT results. In each figure, percentage errors computed using Eq. 11 are also plotted to illustrate the relative deviation between the predicted and simulated results. A positive error indicates that the prediction overestimates the simulation output, while a negative error reflects underestimation.

$$\%\mathrm{Err} = (\mathrm{Prediction} - \mathrm{Analysis})/\mathrm{Analysis} \cdot 100 \qquad (11)$$

It should be noted that, due to the density of the data and for the sake of visual clarity, the rotor angles of attack for the training and testing conditions are not explicitly represented in the figures. Nonetheless, a general trend can still be interpreted: as the angle of attack of the rotor increases, the predicted rotor thrust tends to decrease. This behavior can be

Table 3. Evaluations of airflow predictions

Rotor-A	R^2	RN-RMSE
Vel-X Training	0.992	0.009
Vel-X Test	0.965	0.023
Vel-Y Training	0.998	0.007
Vel-Y Test	0.949	0.041
Vel-Z Training	0.989	0.012
Vel-Z Test	0.960	0.037
Rotor-B		
Vel-X Training	0.998	0.007
Vel-X Test	0.951	0.035
Vel-Y Training	0.998	0.007
Vel-Y Test	0.980	0.021
Vel-Z Training	0.994	0.009
Vel-Z Test	0.964	0.030

attributed to the reduction in effective blade angle of attack caused by increased inflow, which lowers lift and thus thrust.

Figures 9 and 10 present the comparison of thrust predictions for Rotor-A under both training and testing conditions. In hover, which represents the simplest aerodynamic regime, the results obtained from the Lifting Line method are in good agreement with those of the BEMT model, exhibiting deviations within 4% for both training and test cases. The VLM also provides thrust estimates within the same error margin, demonstrating its accuracy under these low-complexity conditions.

Thrust values obtained using a uniform inflow assumption are also shown for comparison in Figures 9 and 10. In both training and test cases, the uniform inflow significantly overpredicts the thrust yielding values approximately 35% higher than those predicted by the LL and VLM models. This highlights that, while uniform inflow theory may be suitable for rough estimations, a more detailed inflow representation is essential for accurate and realistic rotor performance predictions.

For non-hover training conditions, the thrust predictions using BEMT remain mostly within about 10% error tolerance when compared to mid-fidelity simulations. However, several interesting behaviors are observed across different regimes. In axial flight cases, total thrust approaches near-zero values. This causes the relative percentage error to become disproportionately large, even if the absolute deviation remains small. Such cases highlight a limitation of percentage-based error metrics in low-thrust regimes, where even small absolute deviations can result in disproportionately large percentage errors. To address this, rotor thrust and torque predictions are additionally evaluated using the same statistical metrics applied to the airflow field. Table 4 presents the R^2 and RN-RMSE values for both Rotor-A and Rotor-B. As shown, the agreement between the BEMT predictions and DUST simulations is strong, with R^2 values around 0.99 and RN-RMSE values approximately 0.04, indicating excellent predictive performance across the evaluated cases.

Notably, BEMT does not exhibit a consistent bias toward

Table 4. Evaluation of thrust and torque predictions

Rotor-A	R^2	RN-RMSE
Thrust Training	0.999	0.009
Thrust Test	0.990	0.040
Torque Training	0.999	0.020
Torque Test	0.994	0.042
Rotor-B		
Thrust Training	0.999	0.021
Thrust Test	0.985	0.045
Torque Training	0.999	0.045
Torque Test	0.996	0.056

underpredicting or overpredicting thrust. Depending on the operating condition, it alternately overestimates or underestimates the simulation results. This variability suggests that BEMT's assumptions about inflow and induced velocity are sensitive to flow regime and rotor orientation.

Although both the Lifting Line method and the VLM are based on inviscid, irrotational, and incompressible flow assumptions, it is observed that VLM tends to overpredict thrust compared to LL. One possible explanation lies in the use of linear potential theory in VLM, which assumes small perturbations and does not inherently capture nonlinear aerodynamic effects. In contrast, the LL method incorporates airfoil aerodynamic data via .c81 tables, which include lift, drag, and moment coefficients as functions of angle of attack, Reynolds number, and Mach number. These tabulated coefficients implicitly account for effects such as compressibility, mild flow separation, and viscous phenomena to some extent. As a result, while LL still shares some of the same limitations as VLM, particularly in its simplified wake modeling, it may yield more realistic aerodynamic predictions due to the embedded nonlinearities and empirical grounding provided by the .c81 data.

As expected, thrust predictions for training conditions align more closely with the simulation data than those for testing conditions. The deviation between BEMT and DUST increases with the rotor angle of attack. In these test cases, the predicted inflow field struggles to capture the combined effects of free-stream and induced velocities. Specifically, in cases with strong axial free-stream components, the ML model tends to overpredict the inflow, which reduces the effective angle of attack experienced by the blade elements and, consequently, underestimates the generated thrust.

Another important observation is that as the rotor's rotational speed increases, the relative influence of the free-stream velocity diminishes, bringing the rotor closer to hover-like conditions. In these high-speed cases, the flow becomes increasingly dominated by the rotor-induced velocities, and the ML-inferred inflow field better matches the actual conditions. This leads to improved agreement between BEMT and simulation results and highlights the robustness of the surrogate modeling approach in hover and near-hover regimes.

In general, it can be concluded that for Rotor-A, the predicted thrust results show good agreement with DUST simulations across a wide range of conditions. This supports the applicability of the machine-learning-augmented BEMT framework as a fast and reasonably accurate tool for rotor performance estimation.

In Figures 11 and 12, the torque predictions for Rotor-A exhibit trends that are generally consistent with the thrust results, showing good agreement between the BEMT-based predictions and the simulation data for both training and test conditions. In hover conditions, the torque prediction errors remain within approximately 4%, indicating a reliable match. On the other hand, VLM results tend to underpredict the torque, which can be attributed to its systematic underestimation of drag forces due to the assumptions of inviscid flow and lack of detailed boundary layer modeling.

Rotor torque predictions based on the uniform inflow assumption are also presented for comparison. The uniform inflow approach overpredicts the rotor torque by approximately 20%. This overestimation stems from the simplified nature of the uniform inflow model, which neglects the spanwise variation of the induced velocity. Such simplifications lead to an inaccurate estimation of the local angle of attack and dynamic pressure across the rotor disk, ultimately affecting the torque calculation. Although the torque error is less pronounced than in thrust, this result still emphasizes the need for incorporating a more realistic inflow model, especially in applications where accurate power estimations and performance assessments are critical.

As the free-stream velocity increases and axial flow effects become more prominent, the discrepancy between predicted and simulated torque values grows, reaching up to about 20% in certain cases. However, these larger percentage errors are primarily associated with conditions where the absolute torque values are small. Therefore, complementary statistical metrics such as R^2 and range-normalized RMSE are also used to ensure a robust evaluation of model accuracy.

Overall, the torque results further support the capability of the proposed approach, demonstrating strong agreement with the mid-fidelity DUST simulations across a wide range of operating conditions.

For Rotor-B, a key observation is that the nominal percentage errors are significantly reduced compared to Rotor-A. This improvement is attributed to the fact that both thrust and torque values across the training and test cases are sufficiently large, avoiding the amplification of relative errors due to division by small reference values. As a result, all errors for both thrust and torque remain below 10%.

The second important observation is that the analyzed free-stream velocities for Rotor-B are lower than those for Rotor-A. Consequently, the influence of the free-stream flow is minimal, and the flow conditions can largely be considered as near-hover. Moreover, as the rotor speed increases, the relative impact of the free-stream velocity diminishes even further.

Similar to the Rotor-A case, the VLM continues to overpredict thrust and underpredict torque. This systematic trend suggests a consistent behavior of the VLM-based inflow predictions across different rotor configurations.

Overall, BEMT predictions for Rotor-B demonstrate high accuracy. Both thrust and torque errors for training and test cases remain below 10%. As can be seen in Table 4, BEMT results are in strong agreement with DUST simulations, confirming the effectiveness of the machine-learning-enhanced inflow model.

FUTURE WORK

This study establishes a promising foundation for enhancing aerodynamic modeling in Advanced Air Mobility (AAM) rotor systems by enabling accurate, mid-fidelity inflow field prediction suitable for integration with computationally efficient methods such as BEMT. However, several directions for future work are identified to extend the applicability and robustness of the proposed framework.

First, the methodology developed here holds significant potential for rotor-rotor interaction studies, which are critical in AAM configurations featuring closely spaced rotors. The ability to predict inflow fields in near-real-time can enable increased fidelity, computationally efficient modeling of interference effects, making this approach valuable for preliminary design, certification simulation environments, and flight control applications.

Second, the current training dataset consists of a relatively limited number of operating conditions and rotor configurations. To expand the generalization capability and validity region of the trained model, future studies should incorporate a broader range of flight conditions, including off-design and extreme scenarios. This would improve the model's reliability across diverse operational envelopes, especially in safety-critical AAM missions.

Moreover, while this study relies on a mid-fidelity aerodynamic solver (DUST) for generating training and testing data, future iterations could benefit from incorporating high-fidelity computational fluid dynamics (CFD) data or detailed experimental measurements, where available. These datasets would enhance the accuracy and confidence of both the training process and the final integrated BEMT predictions, particularly for complex flow features such as tip vortices and unsteady effects.

Additionally, this framework may be extended to support aeroelastic analyses of rotors in AAM applications. By coupling the predicted inflow fields with structural dynamic models, it would be possible to evaluate structural loads and deformations in real-time, paving the way toward fully integrated flight dynamic simulations and structural health monitoring systems.

Another promising direction is the inclusion of gust responses within the training dataset. Modeling and predicting rotor response under gusty conditions is essential for urban operations, where airflows are often highly unsteady. Training the model to account for such perturbations could enable accurate real-time predictions of unsteady aerodynamic loads, enhancing flight safety and control robustness.

Finally, the present work assumes quasi-steady inflow conditions, neglecting the time-dependent dynamics of the wake. Future work may focus on incorporating temporal sequences into the model, enabling the prediction of dynamic inflow behavior, which enhances the model's ability to capture transient phenomena and improves its fidelity in dynamic maneuvers or time-critical control applications.

CONCLUSIONS

In this study, a methodology was presented to enhance the aerodynamic modeling fidelity of rotor systems, particularly for Advanced Air Mobility (AAM) applications, by integrating machine learning-based airflow predictions into the BEMT framework. The following conclusions can be drawn:

1. The accuracy of Blade Element Momentum Theory (BEMT) was significantly improved by prescribing inflow velocity fields predicted by a trained machine learning model.
2. The machine learning model demonstrated much faster prediction times compared to mid-fidelity aerodynamic simulations, making it a promising approach for rapid estimations.
3. Although faster, the current machine learning implementation still requires optimization in terms of computational load and latency to be suitable for real-time applications where timing constraints are critical.
4. The methodology is not limited to mid-fidelity simulation data; it can be extended using high-fidelity CFD results or rotor test data, provided the full velocity field is available.
5. The proposed approach lays the groundwork for studying rotor-rotor interactions in AAM platforms by enabling real-time or near-real-time aerodynamic load predictions.

Authors contact:
Alaaddin Furkan Uzunkaya
alaaddinfurkan.uzunkaya@polimi.it

Pierangelo Masarati
pierangelo.masarati@polimi.it

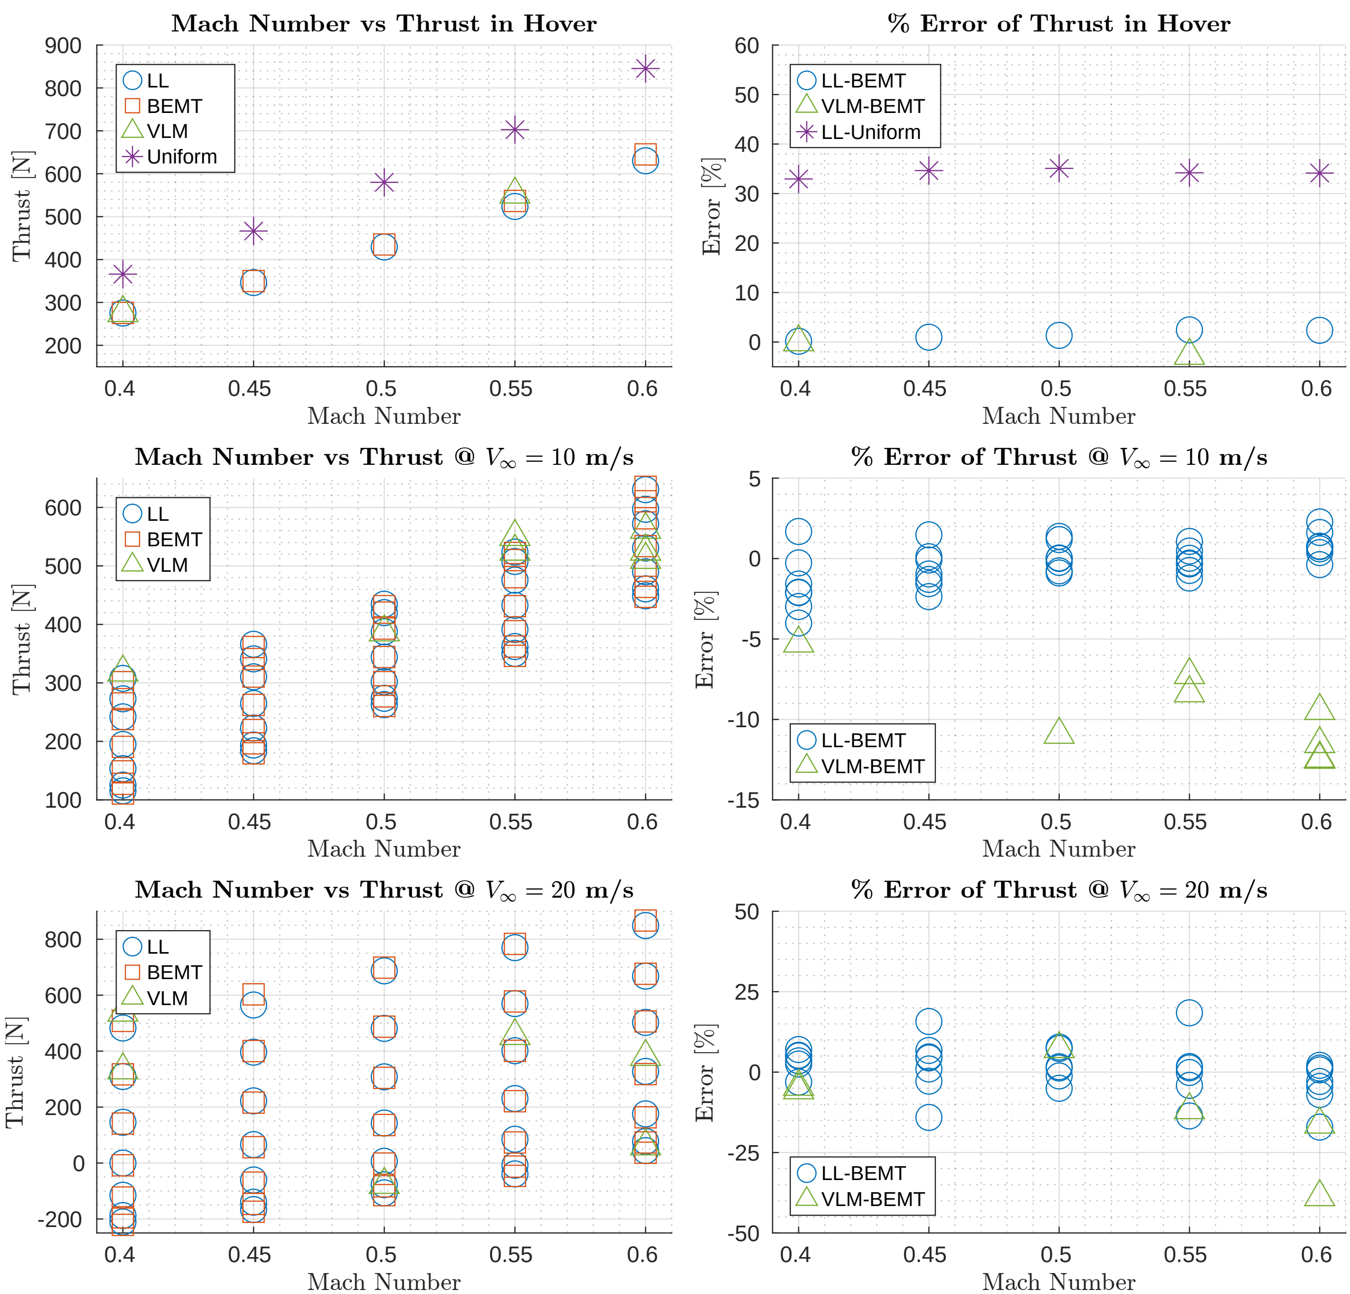

Figure 9. Rotor-A training conditions thrust comparisons and % error.[1]

[1]Due to near zero thrusts, excess % errors not shown

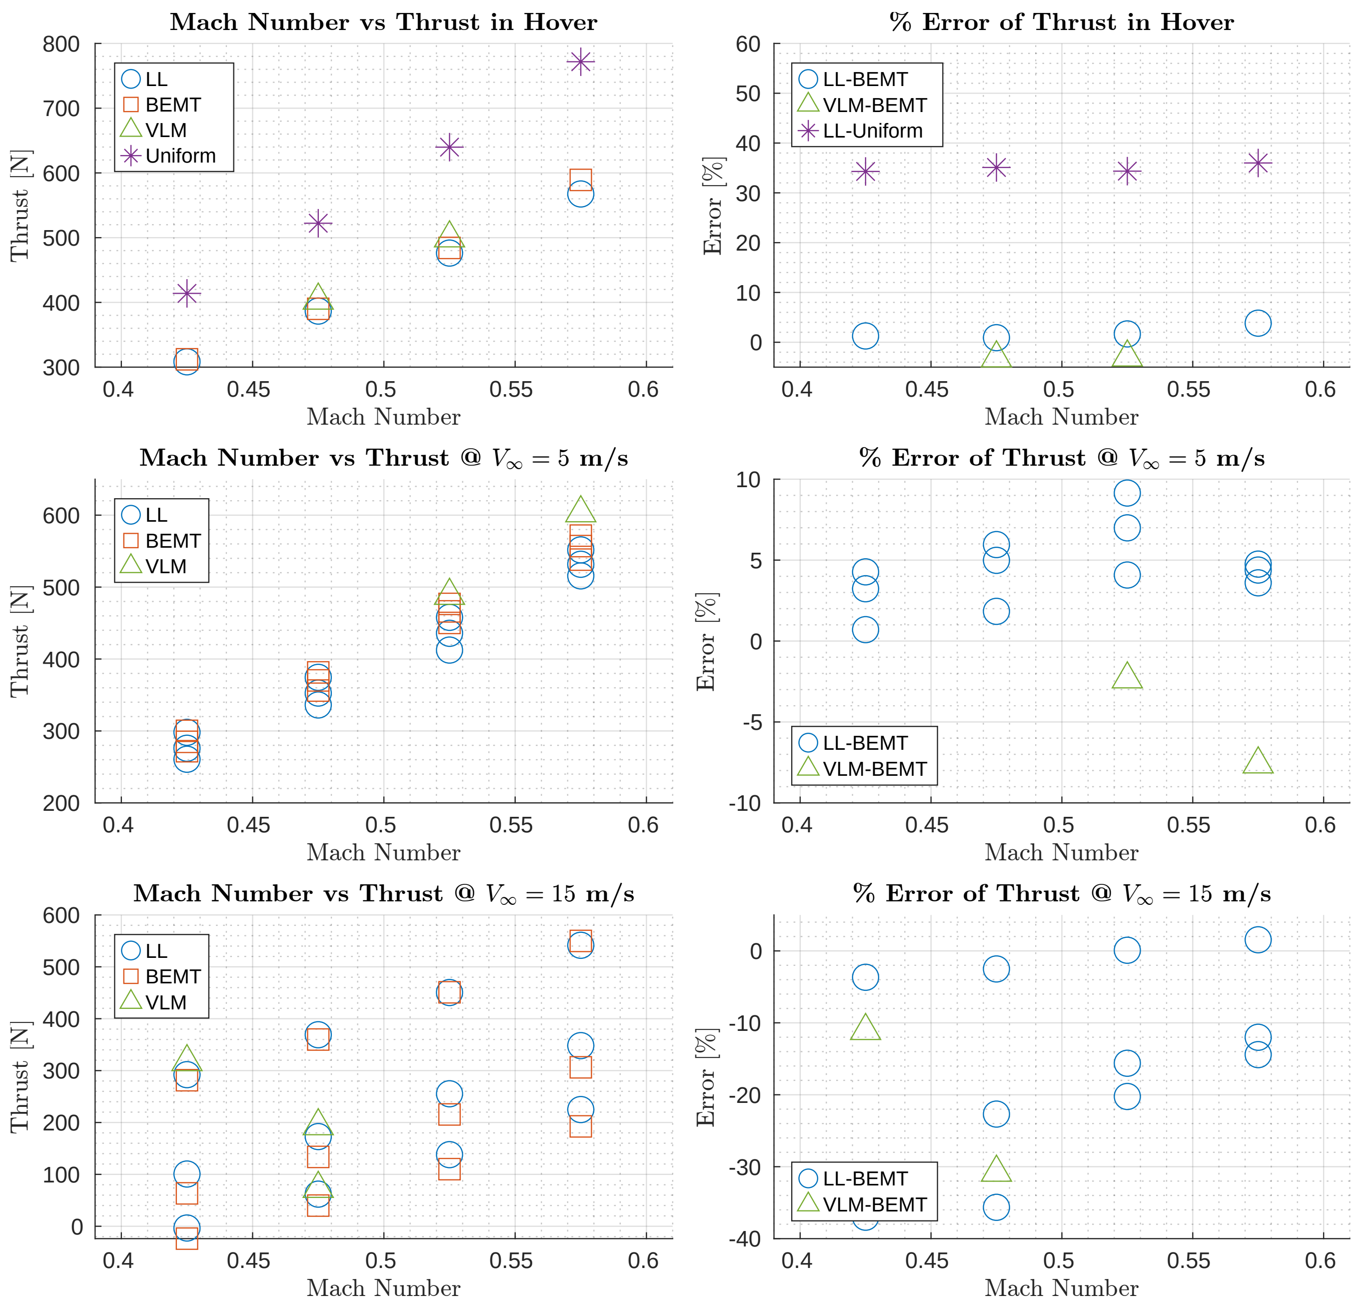

Figure 10. Rotor-A test conditions thrust comparisons and % error.[1]

[1]Due to near zero thrusts, excess % errors are not shown

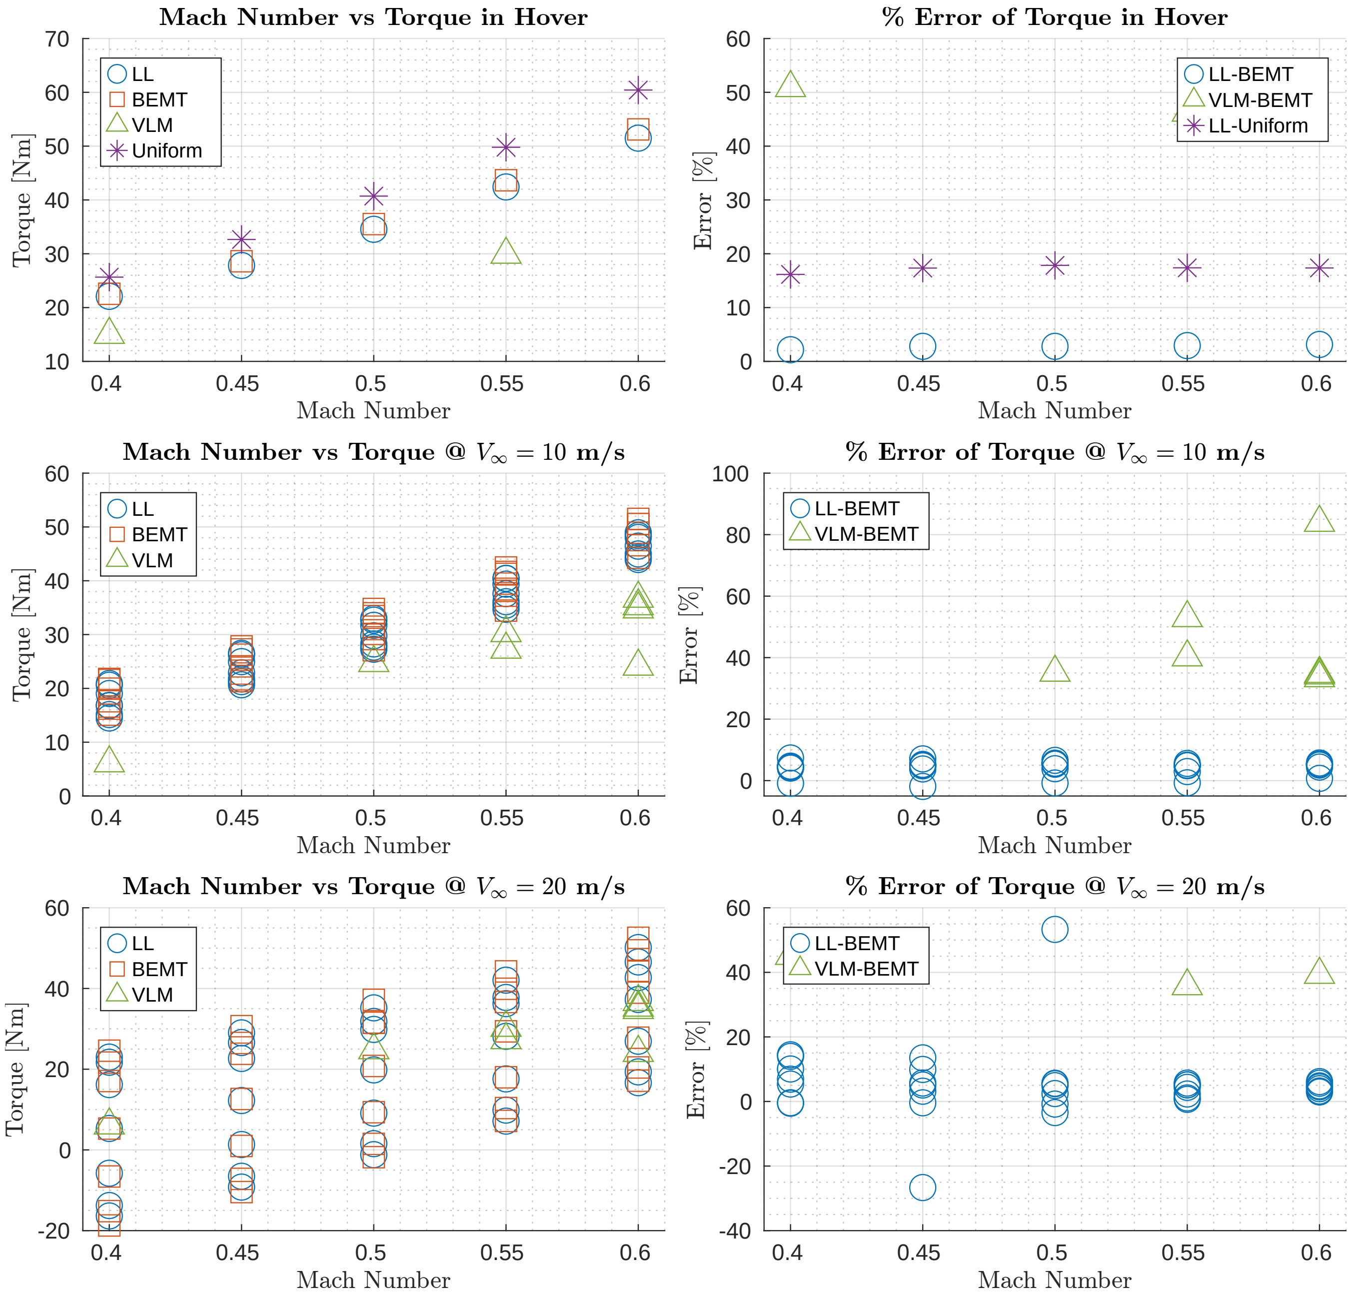

Figure 11. Rotor-A training conditions torque comparisons and % error.[1]

[1]Due to near zero torques, excess % errors are not shown

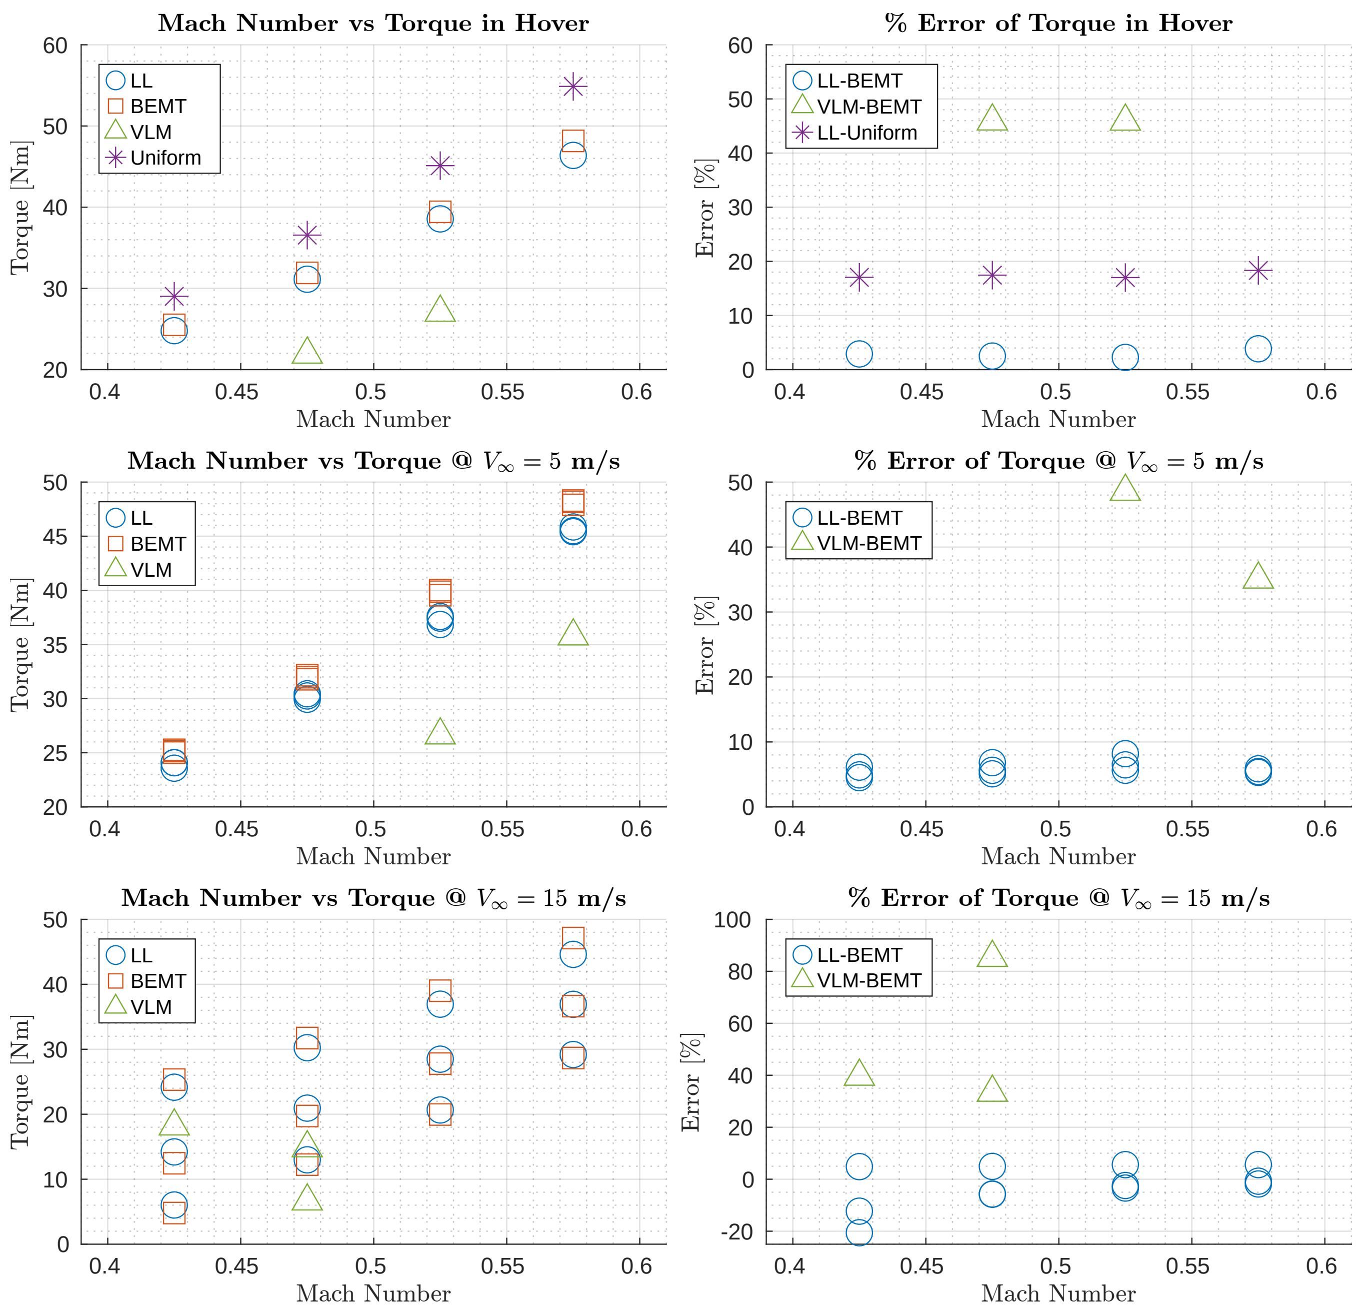

Figure 12. Rotor-A test conditions torque comparisons and % error.[1]

[1]Due to near zero torques, excess % errors are not shown

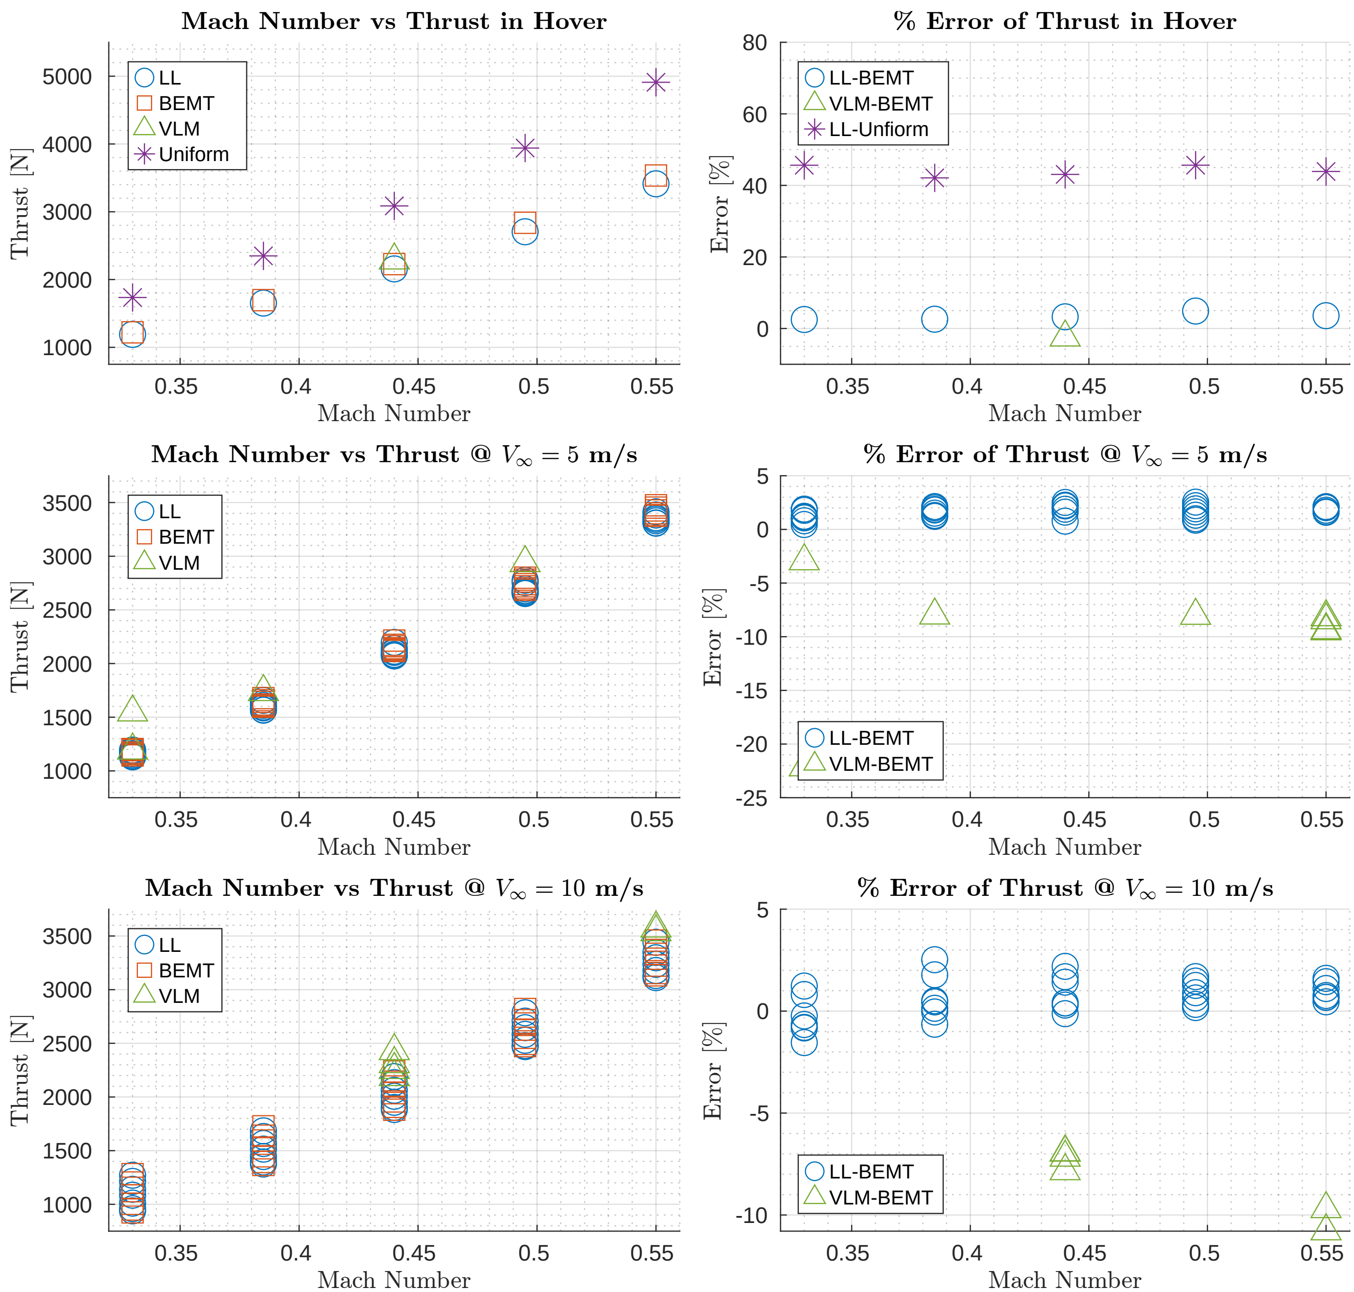

Figure 13. Rotor-B training conditions thrust comparisons and % error.

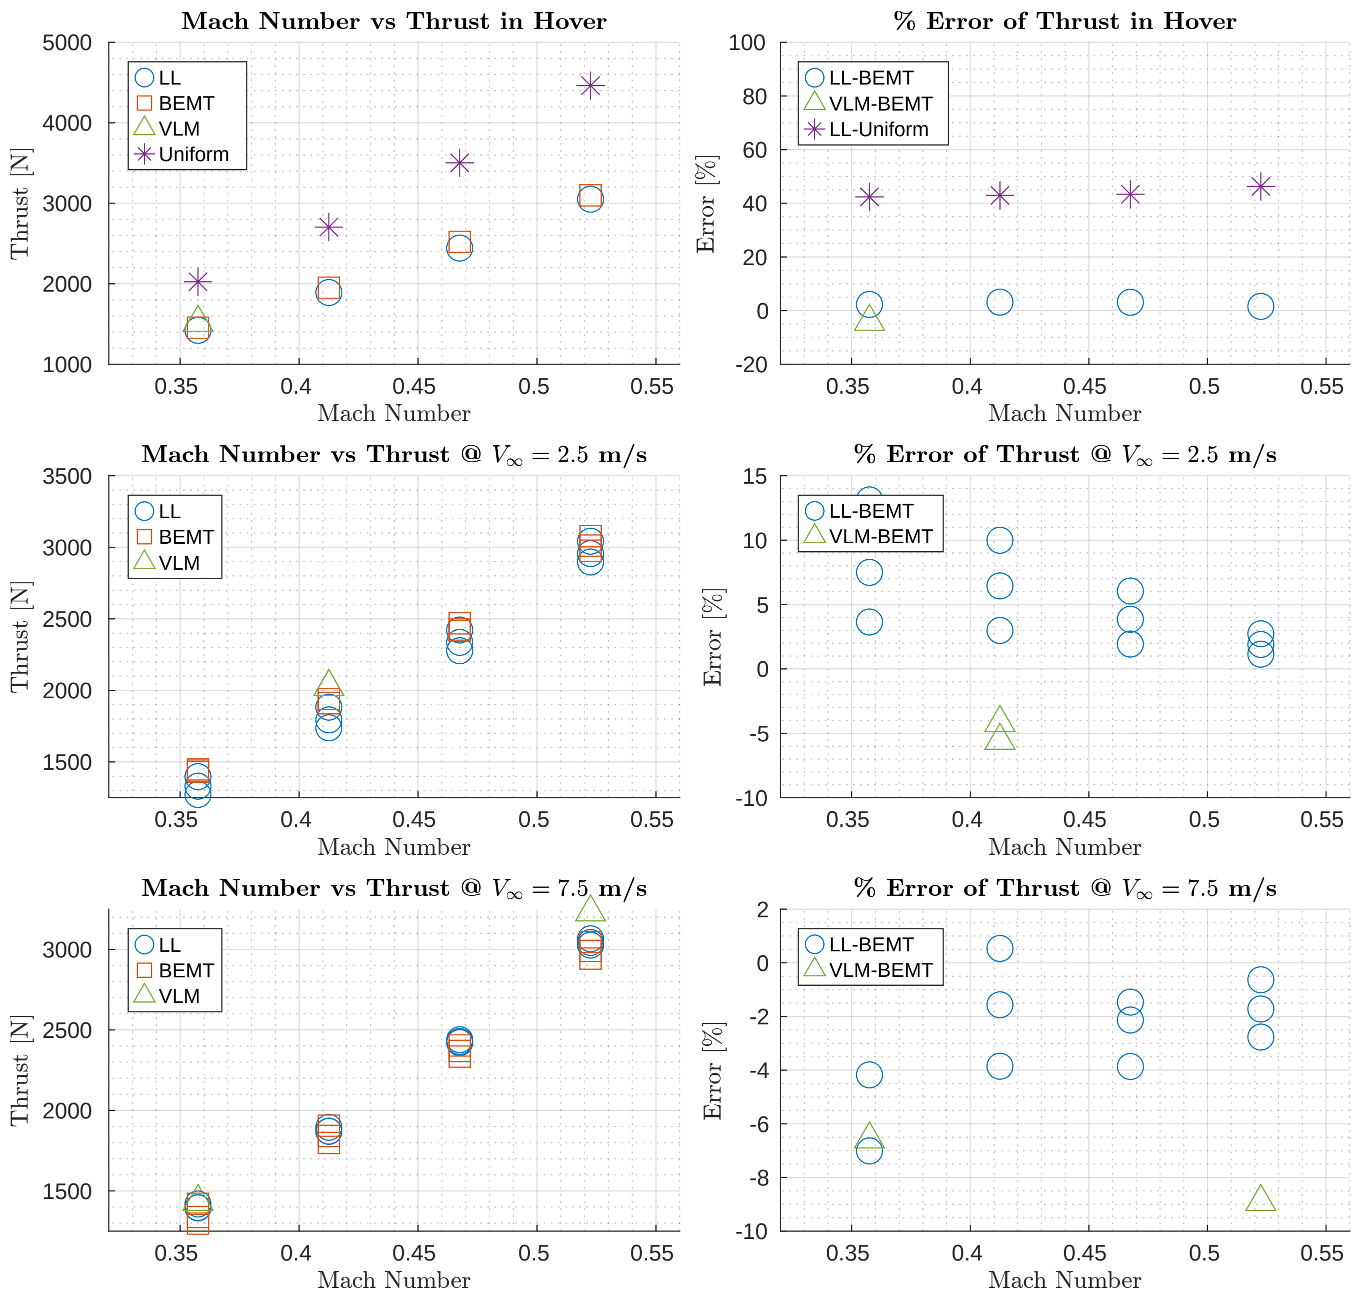

Figure 14. Rotor-B test conditions thrust comparisons and % error.

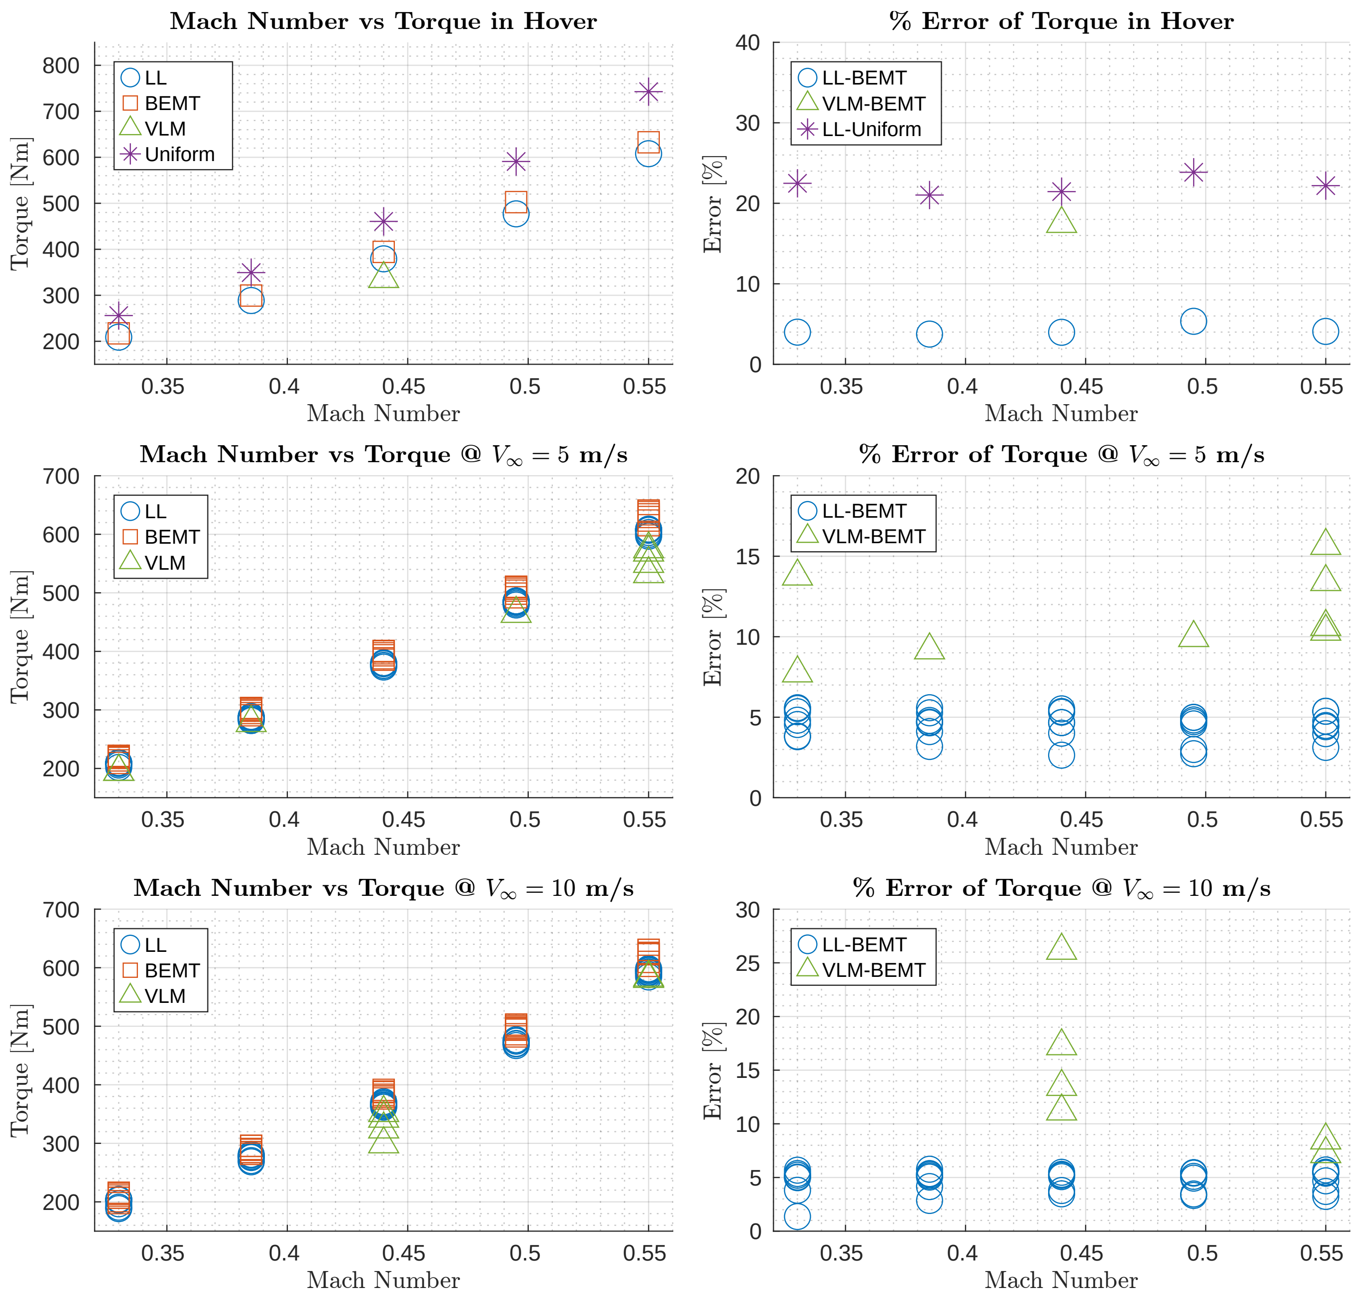

Figure 15. Rotor-B training conditions torque comparisons and % error.

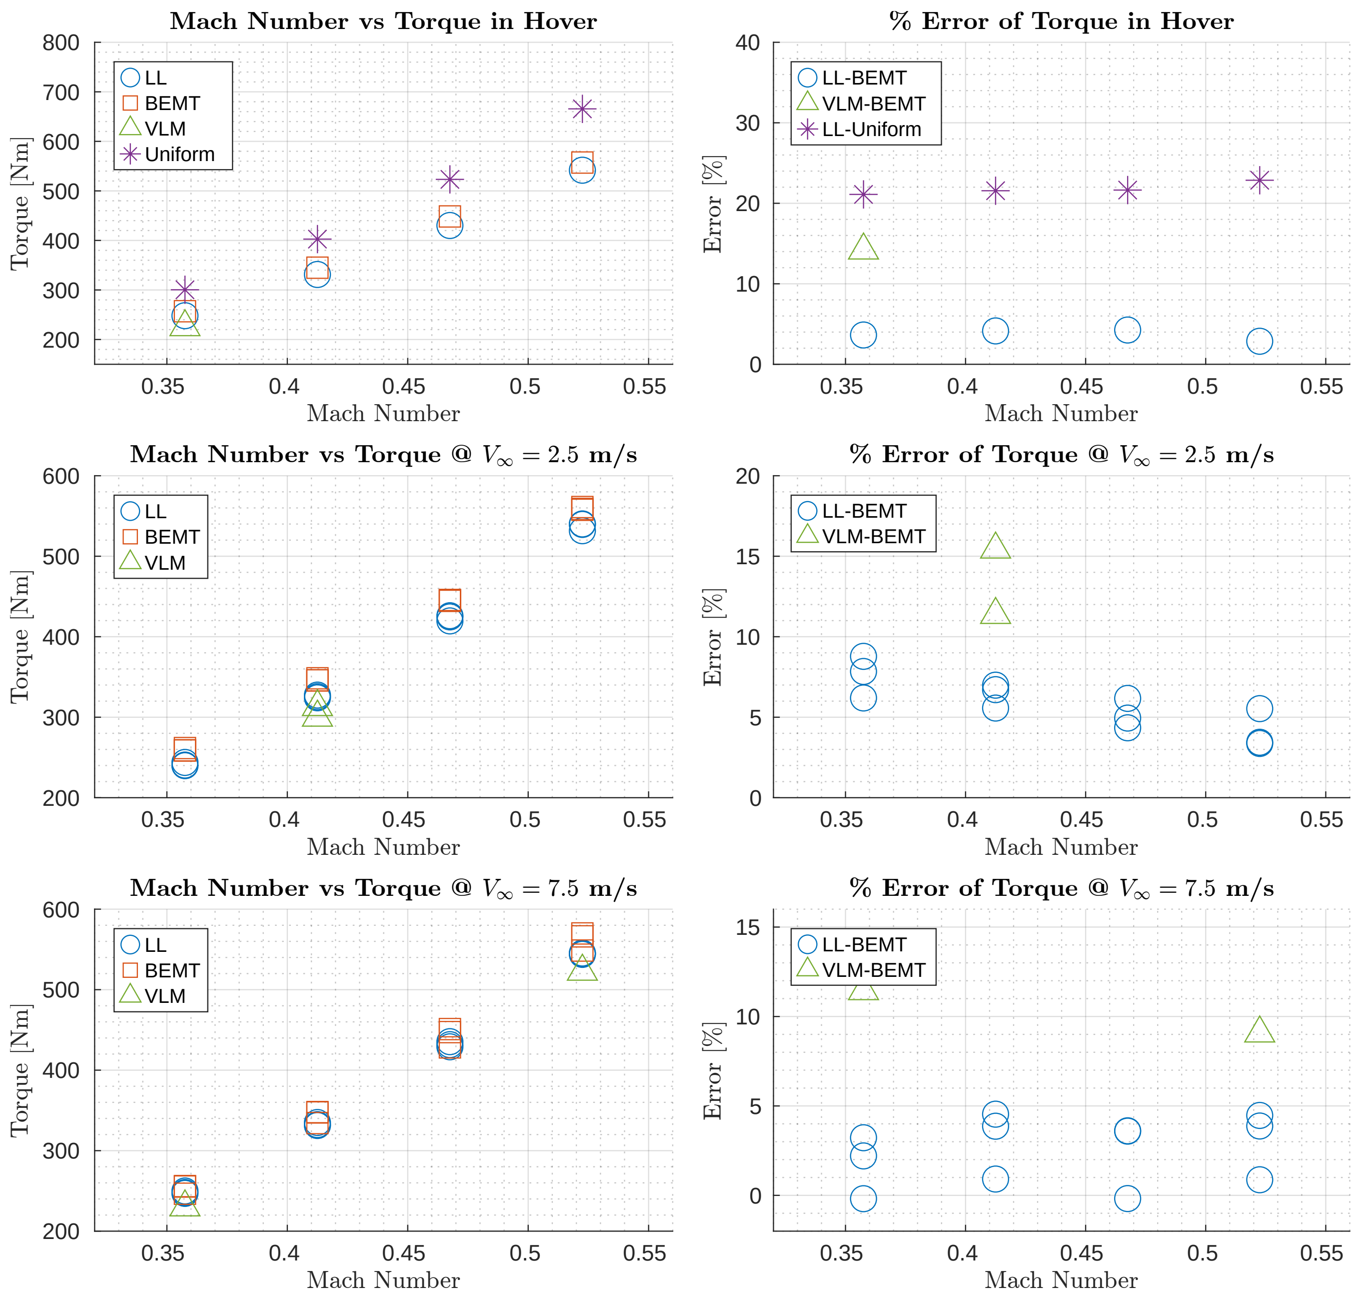

Figure 16. Rotor-B test conditions torque comparisons and % error.

REFERENCES

1. Cornelius, J., and Schmitz, S. "Coaxial Rotor CFD Validation and ML Surrogate Model Generation," 2024 *Transformative Vertical Flight*, February 2024

2. Pedrioli, A., Capone, P., Righi, M., Garcia-Sanchez, E., Pinsard, L., and Gomes, J. V., "MODEL-SI: Modeling and Simulation - Multi-Fidelity Surrogate Model of an eVTOL for Certification," AIAA 2024-1624. *AIAA SCITECH 2024 Forum*, January 2024. https://doi.org/10.2514/6.2024-1624

3. Cornelius, J., Peters, N., Aagren, T. S., and Lugo, D. N., "PALMO: An OVERFLOW Machine Learning Airfoil Performance Database", AIAA 2025-0038. *AIAA SCITECH 2025 Forum* January 2025. https://doi.org/10.2514/6.2025-0038

4. Tugnoli, M., Montagnani, D., Syal, M., Droandi, G., and Zanotti, A., "Mid-fidelity Approach to Aerodynamic Simulations of Unconventional VTOL Aircraft Configurations," *Aerospace Science and Technology, page 106804, 2021*, doi.org/10.1016/j.ast.2021.106804

5. Johnson, W., *Rotorcraft Aeromechanics*, Cambridge University Press, Cambridge, 2013, pp. 39–91

6. Katz, J., Plotkin, A., *Low-Speed Aerodynamics*, Vol. 13, Cambridge University Press, 2001.

7. Cornelius, J., "PALMO: An OVERFLOW Machine Learning Airfoil Performance Database Version 1.0 NACA 4-Series", NASA/TM-20240014546, 2024.

8. Leishman, J. G., *Principles of Helicopter Aerodynamics*, Cambridge University Press, New York, NY, 2008, Chapter 10.

9. Shen, W.Z., Mikkelsen, R., Sørensen, J.N., and Bak, C., "Tip loss corrections for wind turbine computations" *Wind Energy*, Vol. 8, 2005, pp. 457–475. https://doi.org/10.1002/we.153

10. Shen, W. Z., Sørensen, J. N., and Mikkelsen, R., "Tip Loss Correction for Actuator/Navier-Stokes Computations" *ASME. Journal of Solar Energy Engineering*, 127(2), 2005, pp. 209–213. https://doi.org/10.1115/1.1850488

11. Chicco, D., Warrens, M.J., Jurman, G., "The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation," *PeerJ Computer Science*, 2021, pp. 7:e623. https://doi.org/10.7717/peerj-cs.623

12. Chakraborty, D., Elzarka, H., "Performance testing of energy models: are we using the right statistical metrics?" *Journal of Building Performance Simulation*, 2018, 11:4, pp. 433–448, https://doi.org/10.1080/19401493.2017.1387607

CLASSIFICATION & COPYRIGHT