Andrew Vallent (andrewcv24@gmail.com), Alyson Zhang (alyson.lq.zhang@gmail.com), Jess Irvin Peleo (jessirvino.peleo@gmail.com), Benedict Martinez (benedictryanmartinez@gmail.com), John Carlo Figueras (jfigueras719@gmail.com)

Advised by Dr. Muhammad Hajj

Stevens Institute of Technology

Neural Network Informed Hydrofoil Optimization Pipeline

ABSTRACT

Designing hydrofoils that maintain high performance across varying operating conditions remains a complex and time-intensive process. A pivotal step in hydrofoil design is shape optimization. The goal of which is reducing hydrodynamic drag while maintaining sufficient lift, and can significantly lower propulsion power requirements. However, traditional design approaches rely on computational fluid dynamics (CFD) simulations and physical model testing, both of which require significant computational effort, time, and experimental resources. Even modest improvements in lift-to-drag ratio can translate into substantial fuel or energy savings over extended operating periods, thereby reducing operating costs and associated greenhouse gas emissions. Currently, engineers lack a rapid, automated framework capable of efficiently generating optimized foil geometries while reliably evaluating their performance.

In this work, an artificial intelligence–based design framework for rapid hydrofoil shape optimization and performance assessment was developed. The proposed workflow integrates multiple neural networks into a unified automated framework. A convolutional neural network first compresses the foil geometry into a reduced set of six design parameters. These parameters are subsequently reconstructed into candidate foil shapes using a fully connected neural network. A neural optimization machine is then employed to iteratively adjust the foil geometry to maximize hydrodynamic efficiency. Lift and drag coefficients are evaluated using the available physics-informed neural network NeuralFoil, and the outputs from the network inform us of the shape efficiency.

To validate the proposed framework, selected optimized foil designs were fabricated using additive manufacturing and experimentally tested in the Davidson Laboratory towing tank. Experiments were conducted at Reynolds numbers between 150,000 to 450,000 and angles of attack ranging from 0° to 8°, representing realistic operating conditions for small hydrofoil craft models. Performance of the optimized foils was compared with standard reference foils, such as the NACA 0012, under identical testing conditions. The complete optimization process—from initial geometry input to finalized foil design—required only a few minutes of computational time.

Overall, the proposed optimization framework enables rapid hydrofoil prototyping and testing of different configurations for performance evaluation at substantially lower computational cost than conventional simulation-based methods. Although the current framework does not yet account for free-surface interactions, three-dimensional effects, structural constraints, or manufacturing complexity, the results demonstrate the strong potential of artificial intelligence as a fast preliminary design tool. With expanded training datasets and further development, this approach could evolve into a

comprehensive system for optimizing hydrofoils in practical marine applications, supporting the development of more efficient and lower-emission high-speed vessels.

INTRODUCTION

High-speed vessels such as ferries, research craft, and seaplanes are becoming increasingly dependent on efficient hydrodynamic performance to meet demands for reduced fuel consumption, improved speed, and lower environmental impact. Hydrofoils help achieve these goals by lifting a vessel's hull out of the water to reduce drag, but the design process for hydrofoil shape optimization remains slow and computationally expensive. Traditional workflows rely heavily on iterative computational fluid dynamics (CFD) simulations or towing-tank testing, both of which limit rapid design exploration.

TalarAI was created to address these challenges by applying machine learning to hydrofoil optimization. The system integrates a convolutional neural network (CNN) encoder, a fully connected decoder, and NeuralFoil, a physics-informed neural network (PINN) trained to predict lift and drag coefficients. This automated workflow enables the rapid evaluation and refinement of two-dimensional hydrofoil geometries.

BACKGROUND

To evaluate optimized foils meaningfully, it is necessary to understand the current baseline hydrofoil performance and physical testing environment.

The project uses the Ski-Cat craft, a high-speed catamaran with hydrofoils, as a reference craft. This vessel has a forward hydroski, as well as a swept, fully submerged main hydrofoil at the longitudinal center of gravity (LCG), and damping plates at the stern. Available data reported on a 1/3 scale manned vessel, while a 1/16 model was available at the Davidson Laboratory.

According to experimental data and prior reports, the foil system behaves as follows:

- The main foil provides approximately 90% of the lift during high-speed operation.
- Forward hydroskis contribute about 4% of the lift and help stabilize low-speed motion.
- Aft damping plates contribute roughly 6% of the lift and supports pitch stability.
- The main foil, a 6% thickness Ogive section, is submerged one chord length below the free surface.
- At 1/3 scale, flaps are used for lift control; at 1/16 scale, angle of attack (AoA) is adjusted instead.

Reynolds numbers for the 1/16 scale Ski-Cat model range from roughly $1.50 \times 10^5$to 4.50×10^5, using Froude scaling as per testing requirements. To simulate Ski-Cat conditions, Reynold's scaling will be used in testing to directly compare lift and drag coefficients to vessel requirements.

DESIGN METHODOLOGY

TalarAI's design workflow integrates machine learning models with hydrodynamic prediction tools to create a repeatable optimization cycle.

Training Data Development

To train the model, a set of training data was created from known airfoil shapes. Approximately 1600 airfoils from the UIUC Airfoil Database (Selig n.d.) were each interpolated into 80 standardized points (40 upper, 40 lower). From the points, 256 x 256 pixel PNG images were generated for CNN training. For ease of training models, foils outside the normalized [0,1] chord domain were removed.

Figure 1: Training set example

CNN Encoder Development

A convolutional neural network encoder was built to compress each foil image into six latent parameters. Convolutional layers extracted geometric features from the 256 x 256 foil images. The encoder outputs a 6-parameter latent vector for each foil. Training used mean squared error (MSE) loss, Adam optimizer (learning rate 5×10^{-4}), with batch size 32, and 80 epochs.

MSE loss directly measures how close the reconstructed foil is to the original 80-point geometry. It penalizes even small deviations, which is important when small geometric errors can lead to large aerodynamic differences. The Adam optimizer converges faster and more reliably than standard gradient descent, especially in problems with many parameters and nonlinear relationships like airfoil shapes. The learning rate of provided a good balance between stable training and speed, larger values led to unstable updates, while smaller ones slowed convergence. A batch size of 32 ensures that each update step reflects a representative sample of the dataset while still fitting comfortably into GPU memory. 80 epochs were required for the model to converge to a low reconstruction error, as shown in the training log (Figure 2) where validation loss continuously decreased before flattening. This training setup consistently produced smooth convergence and minimized overfitting, which is why it was used across all training runs.

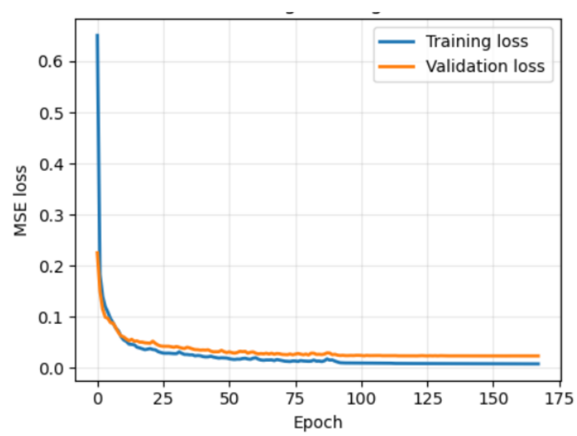

Figure 2: Encoder Training Log

Six latent parameters were found to minimize reconstruction loss effectively and provide a compact yet expressive encoding of foil geometry. The number of latent parameters controls how much geometric variation the model can represent. Multiple latent dimensions were tested, and their reconstruction losses were compared (see Figure 3).

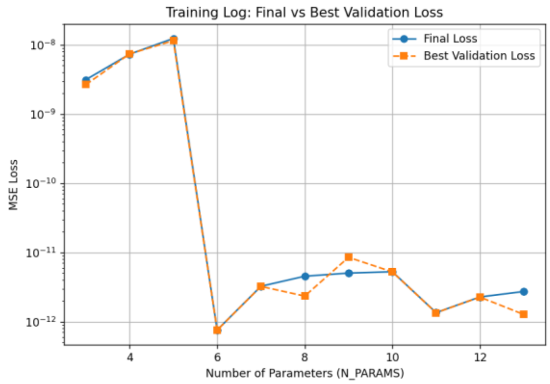

Figure 3: Loss vs. number of parameters

For this neural network configuration, the minimum-loss point occurs at 6 parameters. The performance drop after 6 indicates that adding more parameters gives the model too much capacity, causing it to memorize rather than generalize. Therefore, 6 latent parameters were chosen because:

1. They achieve the lowest reconstruction loss out of all tested values
2. They represent the optimal balance between accuracy and model complexity
3. They avoid the overfitting observed when using 8-10 parameters
4. They align with the recommendation to pick the optimal point before degradation begins

Decoder Development

The fully connected decoder reconstructs foil shapes from the latent parameters.

The network architecture takes 6 input nodes to 100 hidden nodes to 1000 hidden nodes, then finally 80 output nodes. Decoder outputs are the y-coordinates for 80 foil points; x-coordinates follow the standardized distribution used in the dataset. Inputs are normalized during training using a standard scaler for numerical stability.

The decoder successfully recreates foil shapes with low geometric error, enabling a closed-loop mapping from latent parameters back to physical geometry. Figure 4 shows a sample comparison between hydrodynamic lift and drag at Re = 150,000 for the original foil geometry from the database and the reconstructed foil.

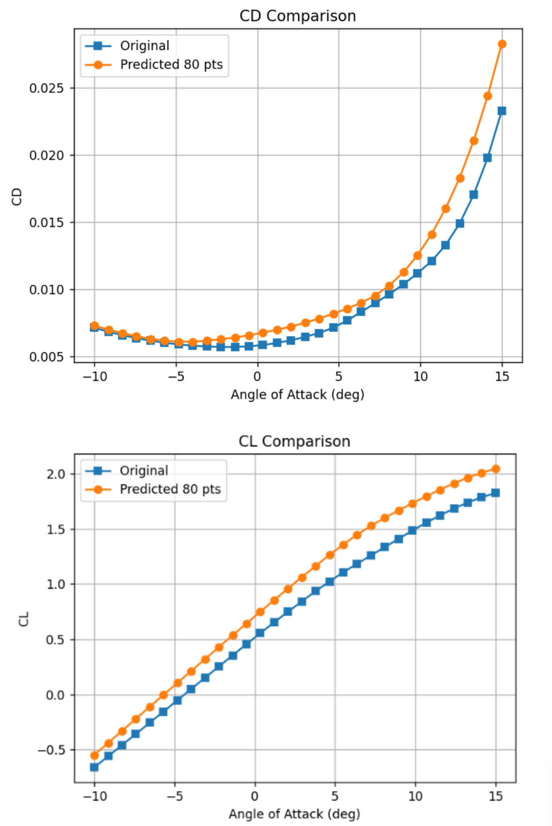

Figure 4: CL/CD vs AoA plot.

Hydrodynamic Evaluation With Neuralfoil

NeuralFoil is a physics-informed neural network trained to mimic XFoil predictions while improving robustness. It was selected over XFoil as an evaluation tool because it eliminates many of the non-convergence issues that occur with XFoil for complex shapes.

To validate the usage of NeuralFoil for the evaluation step in the pipeline, NeuralFoil evaluations of the E61 foil were compared against XFLR 5 predictions and tabular data. Neural foil evaluations resulted in negligible error compared to the reference data. Comparisons of 2D lift and drag predictions at

Re = 100,000 and Re = 500,000 can be seen below. These conditions were chosen based on the availability of data from airfoiltools.com.

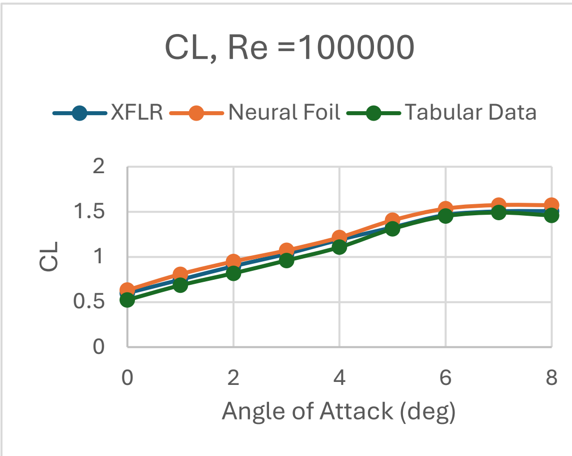

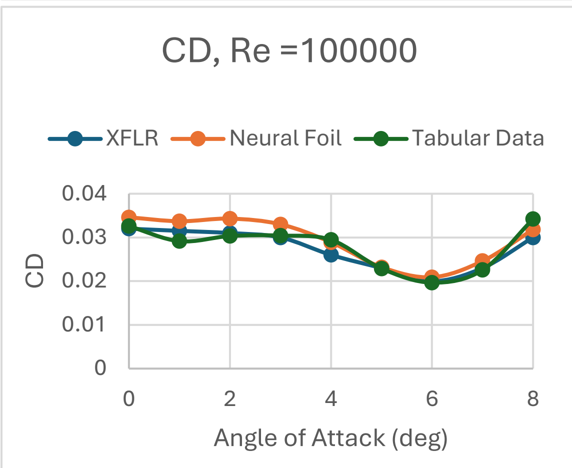

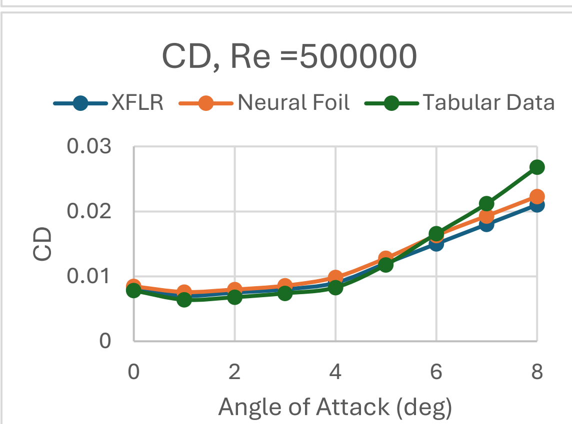

Figure 5: NeuralFoil Validation

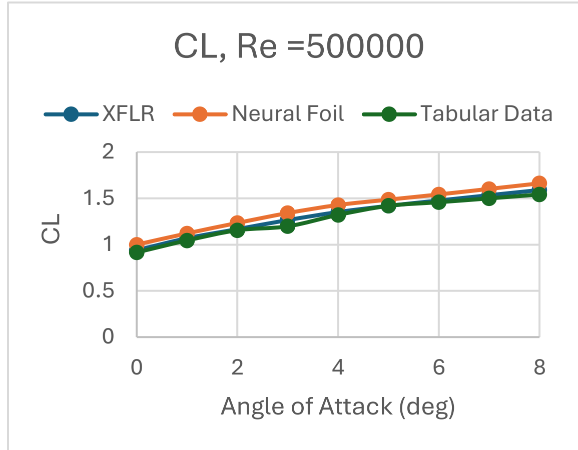

Optimization Workflow

Preliminary optimization was performed using random search across the latent parameter space. A more robust optimization machine was required for complete foil optimization within the design space.

The Neural Optimization Machine is the core optimization driver of the TalarAI pipeline. Building on the preliminary random search framework, the NOM coordinates the encoder, decoder, and NeuralFoil evaluator into a closed-loop optimization cycle targeting maximum hydrodynamic efficiency. The NOM minimizes the drag-to-lift ratio (CD/CL) across the operating angle of attack range, subject to the following constraints:

- Angle of attack bounded within 0 degrees $< \alpha <$ 10 degrees
- Upper foil surface y_u must remain above lower surface y_1 at all chordwise positions
- Minimum thickness constraint applied for structural manufacturability
- Latent parameter vectors bounded within the training distribution to ensure physically realistic geometries

The fitness score for each candidate foil is computed as the mean C_D/C_L ratio evaluated by NeuralFoil across angles of attack from 0 degrees to 8 degrees at the target Reynolds

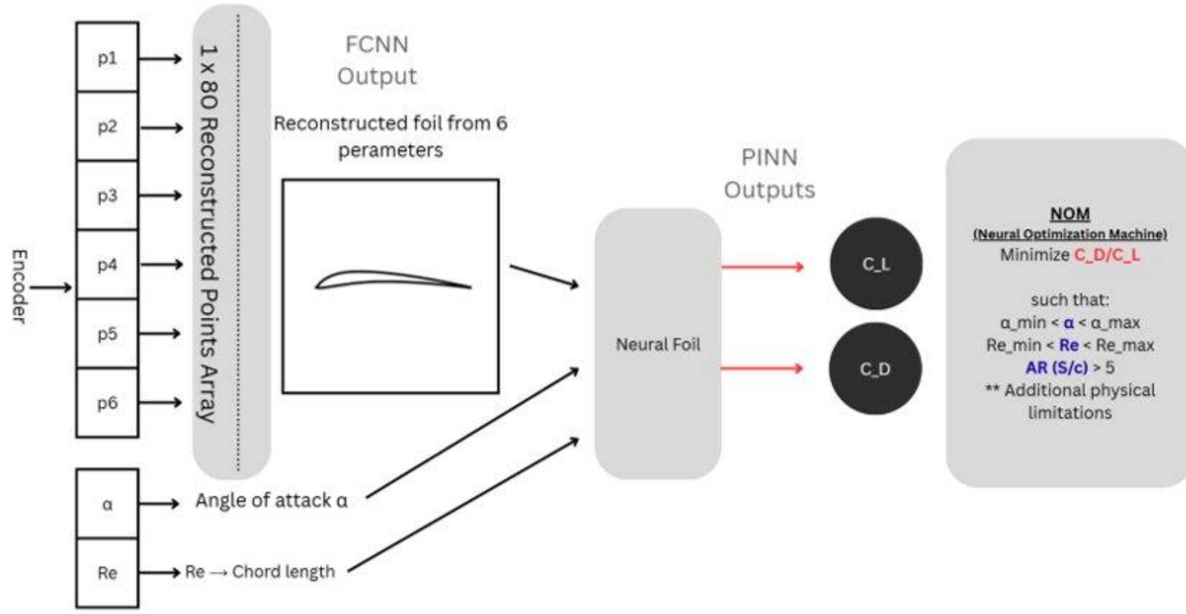

Figure 6: TalarAI pipeline diagram

number. Lower fitness scores correspond to higher aerodynamic efficiency.

Pipeline Execution

The complete framework of the pipeline may be seen in Figure 6. Each optimization pass executes the following steps:

1. Sample a candidate 6-parameter latent vector within bounded ranges
2. Pass the vector through the fully connected decoder to reconstruct an 80-point foil geometry
3. Apply geometric constraint checks; reject any foil violating physical bounds
4. Evaluate the foil with NeuralFoil across the operating AoA sweep
5. Computer the fitness score and store the candidate if it improves on the current best
6. Repeat across the full candidate pool and return the optimal foil geometry

Pipeline Speed Advantage

A key advantage of the TalarAI pipeline over traditional design workflows is evaluation speed. A single NeuralFoil evaluation of one foil geometry across the full AoA sweep completes in approximately 0.1 to 0.5 seconds. By comparison, a comparable XFoil or CFD simulation for one geometry typically requires 10 to 30 minutes of setup, meshing, and computation time. This means the TalarAI pipeline can evaluate hundreds of candidate geometries in the time it would take traditional CFD to evaluate a single foil, enabling a level of design space exploration that would otherwise be computationally prohibitive.

MODEL VALIDATION

To further evaluate the quality of optimized foils over baseline foils, physical model testing was used in the Davidson Laboratory at Stevens Institute of Technology. Foil models were constructed using carbon fiber 3D printing with PET-C as the material. Two aluminum rods were added along the span of the foil for both added rigidity and ease of mounting to endplates. For ease of fabrication, Reynolds scaling was used do adjust foil geometry and speed range such that the chord length of model foils was twice that of the original foils.

Testing Parameters

Testing was conducted in the Davidson Laboratory high-speed towing tank. For closest simulation to 2-dimensional conditions, surface piercing end plates were fixed to either end of the foil. In between the end plates and test foils, plastic mounting plates were fixed to the steel foil rods. These mounting plates allowed for adjustment of angle of attack 0 degrees through +8 degrees. See Figure 7 for a visual representation of angle of attack adjustment.

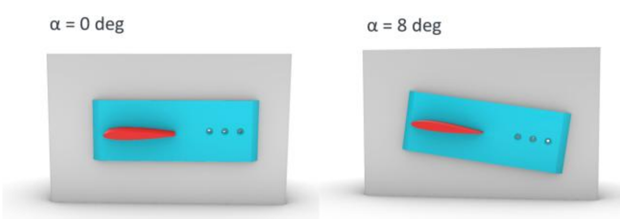

Figure 7: Angle of Attack adjustment on end plates

The aluminum endplates were fixed to an aluminum frame above the water line, to which an inclinometer and 6-pin drag balance were fixed. The inclinometer was used to collect pitch data to ensure the drag moment did not cause considerable pitch deflection. The 6-pin drag balance was selected for its ability to measure both lift and drag forces. The sensor used measured up to a maximum of 250 in-lbs of moment in both the pitch and roll directions. Initially, the maximum allowable roll moment limited the speed at which the foil could be tested, but this was fixed after a pitch correction. The connection between the pitch and roll moment is unknown. After pitch correction, the pitch moment, caused by drag, became the limiting factor to testing. See Figure 8 for representations of the model setup.

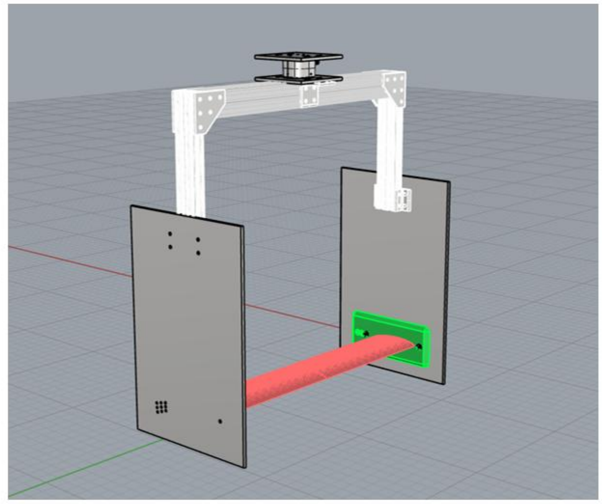

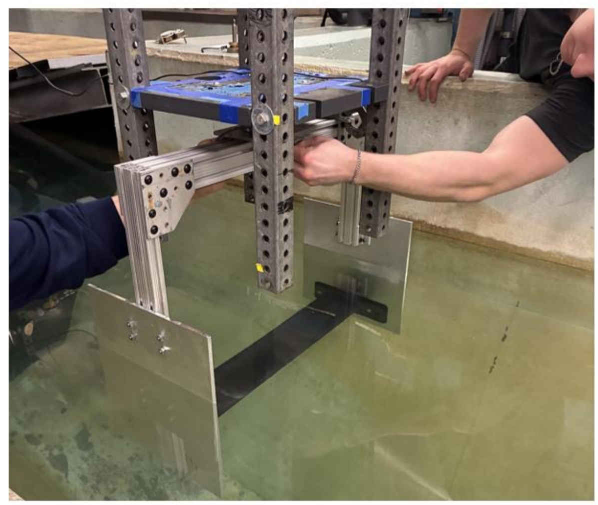

Figure 8: Picture (bottom) and 3D Model (top) of model testing configuration

Testing occurred in 3 stages: endplate testing, E61 foil testing, and optimized geometry testing. First, the endplates and plastic mounting plates without an attached foil were tested at the range of speeds from 4.32 ft/s to 12.96 ft/s. This range corresponds to the associated Reynolds numbers over which the foil was optimized. After testing the endplates, both the E61 and optimized foils were tested at the same range of speeds, with variation in angle of attack from 0 to 8 degrees.

Testing Results

End plate data resulted in a quadratic correlation between speed and drag force. The results from experimentation were plotted and a trendline applied. The equation of this trendline was used in calculation of the drag caused by end plates during physical testing. During testing, it was determined that adjusting the angle of attack of the mounting plate alone had an insignificant effect on drag and lift. Additionally, lift force produced by the end plates was negligible in all conditions.

The end plates were tested at both 2-chord length immersion and 1 chord length immersion. Values at the same speed at these 2 difference immersions were used to linearly forecast the drag value at the operating immersion of the end plates. Plots for drag force vs. speed can be seen below in Figure 9. Using the equation, these drag values were subtracted from the drag values measured during foil testing.

After completion of the end plate testing, the E61 and optimized foils were tested. The foils were tested from a range of 4.32 ft/s to 12.96 ft/s. These values correspond to Reynolds numbers of approximately 150,000 and 450,000 respectively. Comparison between expected and experimental results at the design condition of Re = 150,000 and alpha = 2 deg. can be seen in figures below.

It can be seen from Figure 12 that experimental lift values increase in difference from theoretical 2D predictions at higher speeds. This is likely caused by higher bending in the foil, altering the direction of the lift vector and decreasing measured lift.

Lift-to-drag ratio at both design and maximum speed can be seen plotted below in Figure 14. In comparison to the reference E61 experimental values, the optimized foil achieved a 17.67% lower L/D ratio at design speed and an 11.98% lower L/D ratio at maximum lift condition.

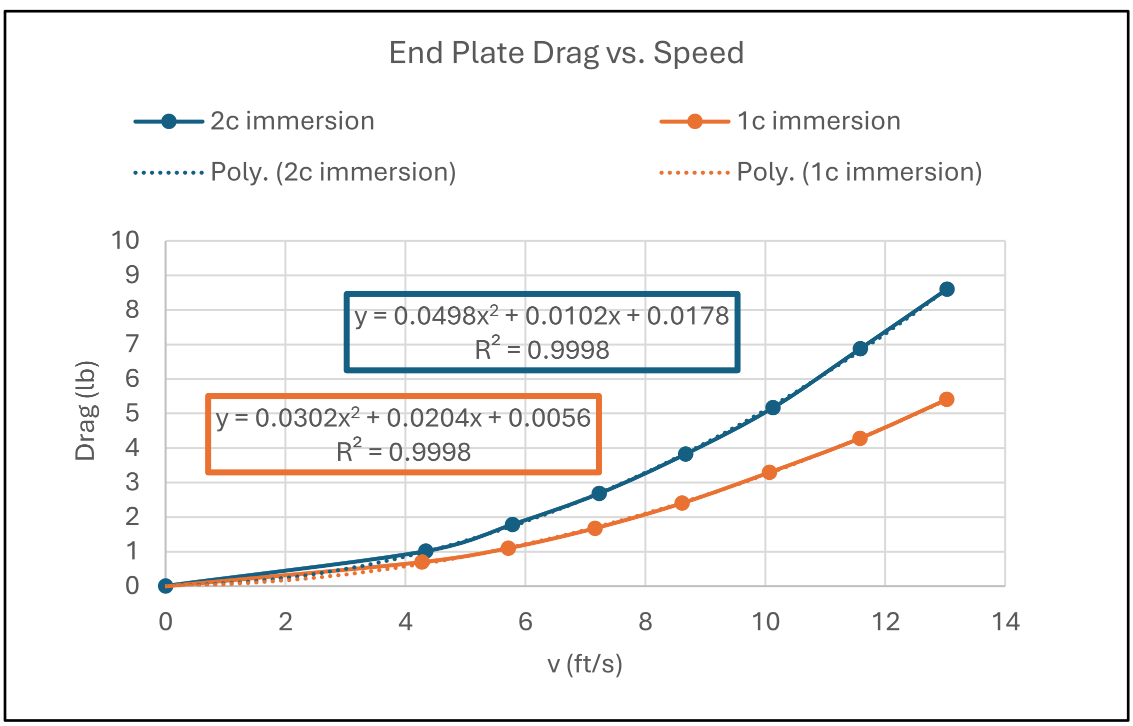

Figure 9: Endplate drag vs. speed for differing immersion

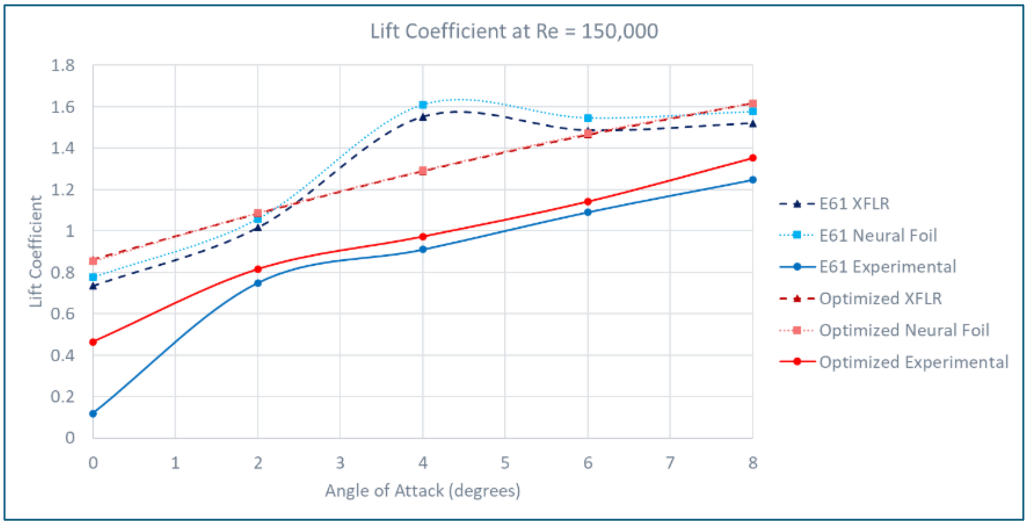

Figure 10: CL vs. AoA at design speed

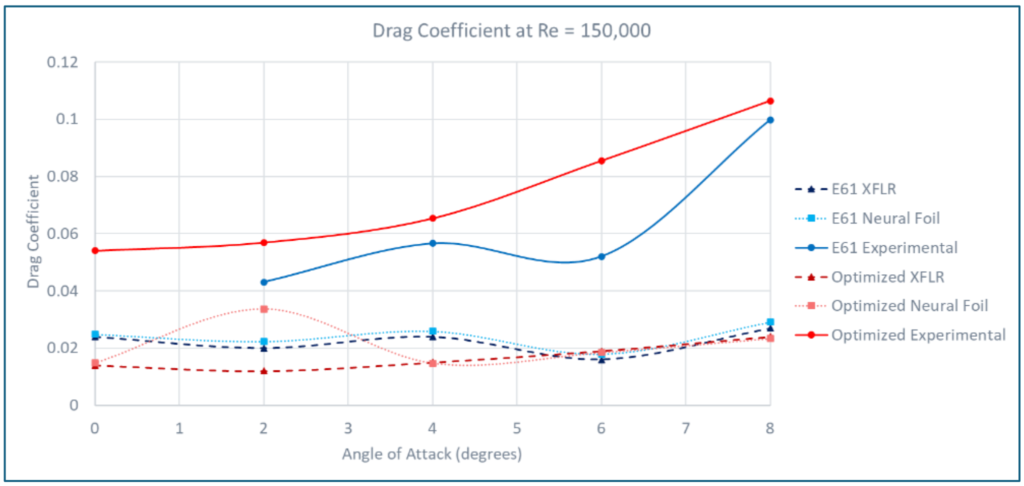

Figure 11: CD vs. AoA at design speed

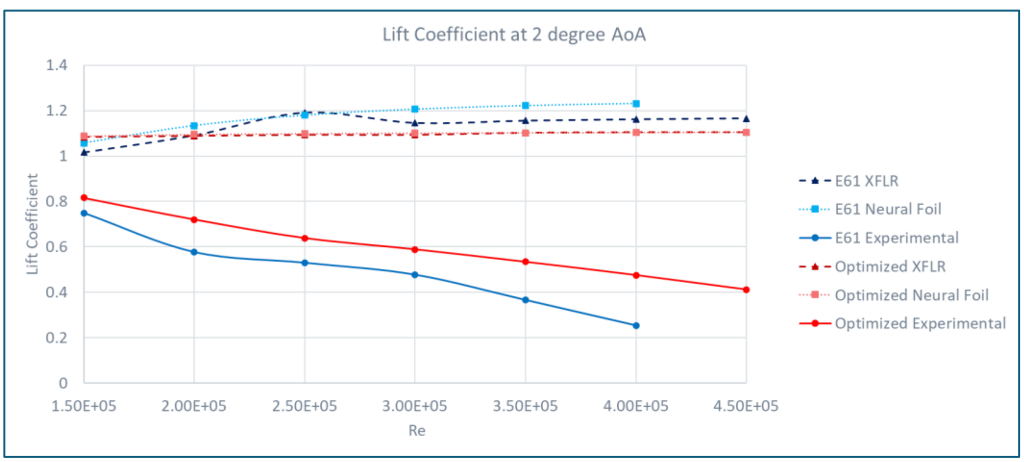

Figure 12: CL vs. speed at design AoA

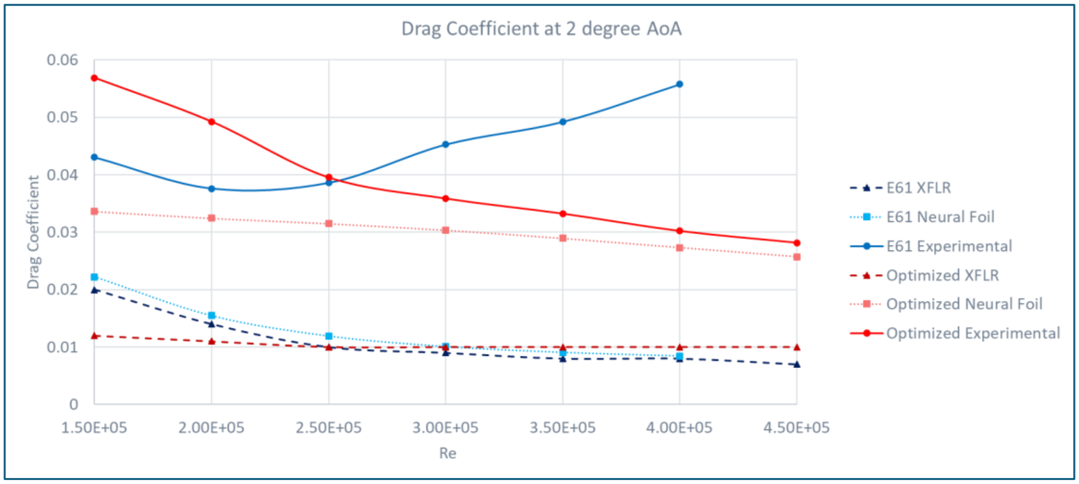

Figure 13: CD vs. speed for design AoA

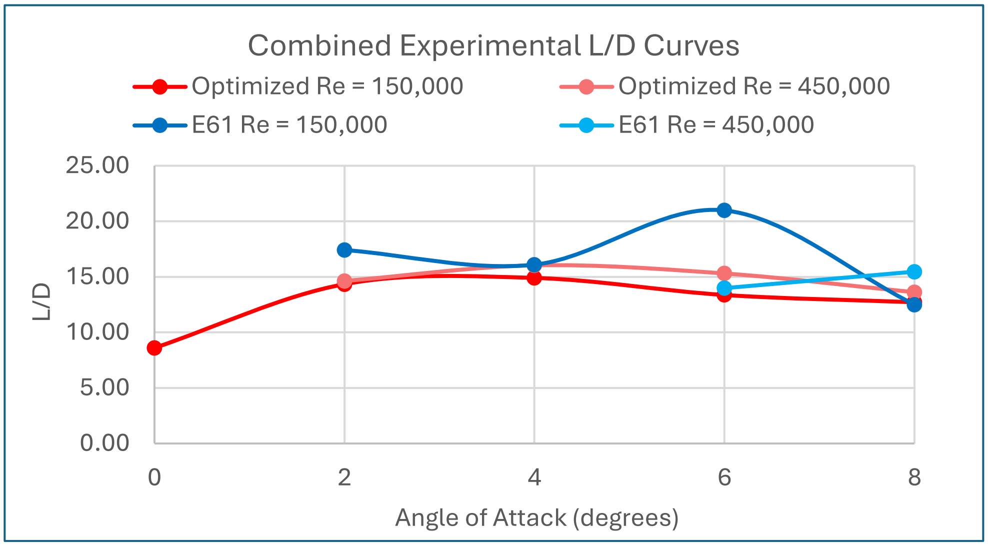

Figure 14: Combined experimental L/D curves

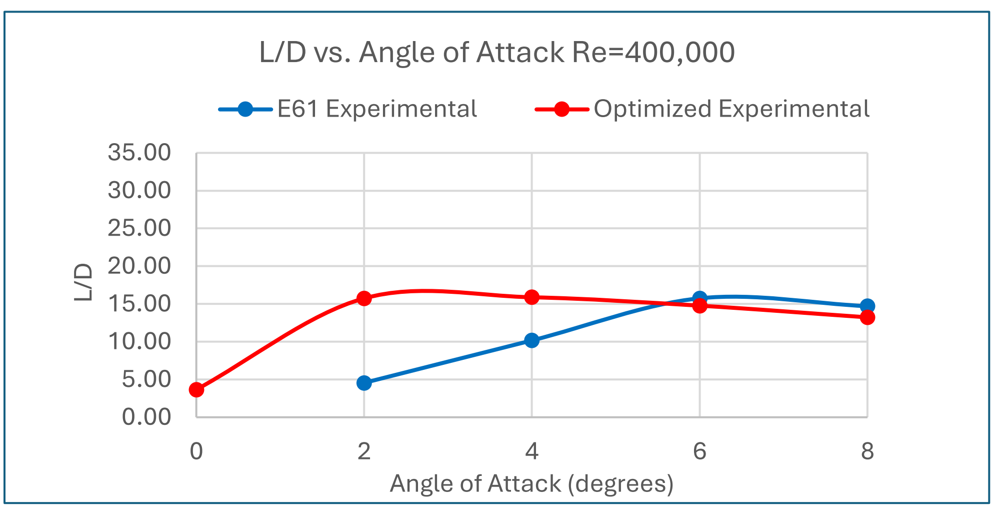

Figure 15:Experimental L/D curves at max speed

Overall, lift-to-drag value did not improve from the E61 control foil at the design condition. At some conditions however, the optimized foil produced noticeably better results than the control foil. At 2 degrees Re = 400,000 condition, the optimized foil saw a L/D ratio triple that of the E61 foil Figure 15.

Differences in experimental data from predicted values are likely caused by imperfections experienced during experimental testing. Drag consistently increased during testing relative to predicted values. This is likely caused by surface roughness of the 3-D printed foils. Additionally, interference from the end plates may induce additional drag over the foil, as well as decrease the lift force over the foil from the induced turbulence. As discussed previously, bending of the foil may have led to additional losses in lift.

CONCLUSION

Based on the results achieved by this iteration of the pipeline, the following recommendations for future work are proposed:

- Apply correction for friction drag on foils from carbon surface, explore foil construction less prone to bending.
- Verify 2D prediction trends by simulating fully three-dimensional foils, capturing spanwise flow and tip effects.
- Prepare 5-10 optimized foil shapes for towing-tank testing to establish statistically meaningful comparisons.
- Train neural networks on more data, allowing for more robust optimization.

The TalarAI system successfully demonstrated its ability to encode, decode, evaluate, and optimize hydrofoil geometries using artificial intelligence and machine learning. During Fall 2025, the team developed a CNN-based encoder-decoder framework, integrated NeuralFoil for hydrodynamic evaluation, and conducted preliminary optimization trials.

With these building blocks in place, an optimized foil was evaluated against a control foil (E61) at design conditions matching that of the Ski-Cat. At the design condition and maximum condition, the optimized foil failed to outperform the control foil. However, significant improvements were achieved in other tested conditions. Refined and expanded testing is required to quantify improvements achieved in the case of the Ski-Cat foil. While experimental results did not show improvement in the foil, errors in experimental procedure and real-world conditions make validation of 2D data against 3D experimental data more complex than could be accomplished during this round of testing. Both XFLR and Neuralfoil predicted improvements in performance in these foils after entering the optimization pipeline.

After incorporating more training data and more advanced models, the hydrofoil optimization tool could prove to be an invaluable tool for naval architects for hydrofoil design. The pipeline optimized foils on the order of minutes in comparison to the days computational fluid dynamics would take. This saves upstream engineering time, allows for more rapid prototyping, and more design iterations. Overall, the TalarAI optimization pipeline is a promising stepping stone in the construction of an AI-enabled hydrofoil design tool that could both decrease required engineering time in hydrofoil design and improve vessel performance.

REFERENCES

Abbott, I. H. 1959. *Theory of Wing Sections - Including a Summary of Airfoil data.* Dover Publications.

Jagupilla. 2025. *NE 423 Course Notes.* Stevens Institute of Technology.

Savitksy, Roper, Mercier. n.d. "Final Report on Design and Construction of 1/3-Scale

Manned Model of Ski-Cat High-Speed Catamaran." Davidson Laboratory.

Savitsky, Daniel. 1964. *Hydrodynamics of Hull.* Hoboken: Stevens Institute of Technology under ONR Contract.

Selig, M.S. n.d. *Airfoil Coordinate Database.* https://m-selig.ae.illinois.edu/ads/coord_database.html.

Y. Y. Liu, J. X. Shen, P. P. Yang, X. W. Yang. 2024. "A CNN-PINN-DRL driven method for shape optimization of airfoils."

AKNOWLEDGEMENTS

The authors would like to recognize Dr. Muhammad Hajj of Stevens Institute of Technology for his mentorship and advising throughout the completion of this project. Additionally, we are grateful for the support provided by the faculty of the Stevens Institute of Technology Department of Civil, Environmental, and Ocean Engineering and Davidson Laboratory including Dr. Raju Datla, Dr. Mahmoud Ayyad, Dr. Ruben Paredes Alvarado, Dr. Hossam Aboalela, David Plaza, and Uihoon Chung. Finally, the team would like to recognize Robert J. Walling of Leidos Inc. for his support throughout the project.