Abstract

Active aerodynamics is an emerging technology that alters vehicle subsystems in real time to optimize as track conditions change. Repeated solver or computational fluid dynamics (CFD) evaluations inside a control loop are too slow for continuous optimization, often forcing fixed geometries and limiting performance. This paper proposes a physics-informed reduced-order model (PI-ROM) that replaces inefficient aerodynamic evaluations with fast, accurate coefficient predictions for active aero rear wing control. A training/testing dataset was generated using XFOIL CFD sweeps for NACA airfoils. The PI-ROM combines thin-airfoil and drag-polar physics as baselines with a neural network residual to improve accuracy for non-ideal conditions. The framework underwent three iterations: calibrating weights and preventing negative drag; improving drag prediction by leveraging the predicted lift within the residual; and revising the controller from curvature-only changes to an objective-driven strategy that prioritizes drag minimization on straights and downforce maximization in corners. The PI-ROM informed a closed-loop controller which was evaluated against a fixed-wing baseline in a vehicle dynamics track simulation described by ordinary differential equations. The PI-ROM achieved a 26.58x speedup relative to XFOIL prediction time while maintaining accuracy (Mean Absolute Error: 0.0298 for CL, 0.0034 for CD; Normalized Mean Absolute Percentage Error: 1.37% for CL and 8.73% for CD). In simulation, active control reduced lap time by 1.2 seconds and saved 5.61 kJ more energy versus a fixed wing on the same track. Future work would involve 3D wing implementation and hyper-realistic vehicle dynamics to match real-world motorsports.

Keywords: reduced-order model; physics-informed model; active aerodynamics; computational fluid dynamics; neural network residual; aerodynamic coefficient prediction; drag polar; thin airfoil theory; vehicle dynamics simulation; lap time optimization; XFOIL; NACA airfoil; wing control; downforce; drag minimization

Ethics Statement

This study is a computational and simulation-based investigation. No human subjects, animal subjects, biological materials, microorganisms, or vertebrate animals were involved at any stage of the research. No experiments requiring institutional ethical approval were conducted. The work was carried out in accordance with applicable academic integrity and safety guidelines. The sole safety consideration was extended computer use, mitigated by taking breaks every 30–45 minutes to prevent eye strain.

Reduced-Order Model for Active Aerodynamics Prediction and Wing Control

Anuj Subramanian[1]
[1] *Adlai E. Stevenson High School, 1 Stevenson Drive, Lincolnshire, IL 60069, USA*
anujsubra@gmail.com
Corresponding Author: Anuj Subramanian
Sponsor: Dr. Susan Brontman | sbrontman@d125.org
Academic Year: 2025/2026

Table of Contents

Acknowledgements

I would like to acknowledge this project's sponsor, Dr. Brontman for her consistent support and suggestions for bettering the project. It is through her that the project became more structured and logical. Additionally I would like to thank Mrs. Palffy and Mrs. Juneja for also supporting the project and allowing it to develop accordingly. I would finally like to thank my friends, families, and previous mentors for providing me with guidance and knowledge.

Problem or Need

The problem this design intends to address is the lack of a quicker, less computationally expensive alternative to Computational Fluid Dynamics (CFD) for predicting aerodynamic coefficients of drag and downforce. Could a physics-informed reduced-order model (PI-ROM) enable real-time wing control and improve lap time/energy consumption compared to baseline solvers and fixed wings? This would allow for real-time wing control which provides an innovative and safe feature for motor sports and vehicles in general. Without a reduced-order model like this, active aero wings, wings that change their geometries in real time, would not be able to change in real time and instead rely on the time/cost expensive CFD or other plausibly ineffective models for optimization.

Current active aerodynamic wings have fixed or non-continous parameters and do not take advantage of reduced-order models in the most effective way they could. Potentially implementing reduced-order models in each aerodynamic device could greatly improve efficiency in all aspects of a motor vehicle when performing tasks such as racing. Being able to optimize and simulate lap times is an essential need in the new era of active aerodynamics for motor sports, and so, within a dynamic system, these aerodynamic properties of a rear wing will

be manipulated based on the reduced-order model to improve performance and safety.

Determining aerodynamic coefficients from a car's airfoil/wing through a design like this will severely cut down time and computational costs which would address this issue successfully. It allows for control loops in active aerodynamics to function quicker and more efficiently which is crucial in safety and performance for motorsports. The control loop will also capitalize on active aerodynamics and demonstrate how active aerodynamics, the manipulation of vehicle subsystems in real time such as tires or the wing, can be beneficial for real motor sports by reducing energy and lap times.

Future machine learning (ML) and real-time control of aerodynamic systems can be achieved through quicker prediction and manipulation of the wing to alter lift and drag (key aerodynamic components of a car) with this design. This application can save numerous lives and even extend to airplanes and other aerodynamic situations. The control algorithm and reduced-order model design will align perfectly with the introduction of active aerodynamics in professional motor sports, proving its relevance and potential if optimized.

Performance Criteria

- The multi-output machine learning (ML) model should and is expected to achieve range-normalized mean absolute percentage error (MAPE) percent values of under 5% for coefficient of lift and under 10%. Mean absolute error (MAE) for coefficient of drag since drag does not vary as much as lift so similar MAEs will result in larger MAPEs. Effectiveness of design will be measured through the MAPE loss values between the reduced-order ML model and CFD dataset. MAE/MAPE values will be used since it is easier to interpret and outliers should not be weighed as much since data is coming from a CFD which will have occasional outliers. For general criteria, the percent values for

MAE are taken with respect to the test-range which means percentages will not be extreme if values for coefficient of drag and coefficient of lift are too small. These values should be obtained when performing on unseen data.

- R2 score to test regression accuracy, should and is expected to be greater than 0.70 for both coefficients of drag and lift to signify a strong correlation. R2 scores are common for input-output evaluation as it determines the strength of the correlation between the two.
- To be applicable for real time prediction, the prediction time should and is expected to be around 1 milliseconds.
- Lap time reduction should and is expected to be around 0.5-1.0 seconds from fixed wing.
- Energy lost should and is expected to be no greater than around 80 kJ

The reduced-order model and control system shouldn't be computationally expensive or time expensive to use. The ordinary differential equations (ODEs) describing the vehicle dynamics should capture all the important information to make it the most reflective of the real world and simulate the vehicle's behavior properly for accurate lap performance metrics. Another thing to be especially cautious about is that the collected CFD data should be consistent and accurate from XFoil. It can easily be overlooked, and contradictory data will result in poor training of the model, so the data cleaning has to be thorough.

Review of Literature

Solving the problem of Computational Fluid Dynamics' inefficiency in calculating coefficients of drag and lift will benefit the real world, and is of interest, because it will allow for active aerodynamic optimization, which is much more efficient than static aerodynamic systems due to its constantly adaptive nature. The motorsports world is shifting from drag reduction

systems, which had a big presence this year, to active aerodynamics where things like a car's wings or other subsystems such as the tires are actively changing depending on the stretch of the track that the car is on (Noble, 2024). For example, in straights, a race car minimizes drag to optimize the energy saved and improve overtaking ability, which is why drag reduction systems were initially implemented (Motorsport, 2023). However, the shift to active aerodynamics allows for this and more, such as having downforce during curves, which is crucial for safety and minimizing drag on straights which saves on energy consumption. Greater aerodynamic adaptability and control which is provided with active aerodynamic configurations reduces lap times immensely due to the modulation and allocation of downforce (Rijns et. al, 2025).

Aerodynamics: The Theory and Certain Tradeoffs

Firstly, the underlying aerodynamics and relationships between functioning systems in a vehicle must be understood. The thin airfoil theory is a method of simplifying aerodynamic behavior with certain assumptions and is used as a representation of the physics for an ideal airfoil (Cadence CFD Solutions, 2023). Airflows can be compressible (density can change) or incompressible and the Navier-Stokes equations describe fluid flow in laminar and turbulent flow conditions with two different forms depending on compressibility.

Figure 2:

Model Diagram of Wing with Acting Forces and Features (Cadence CFD Solutions, 2023)

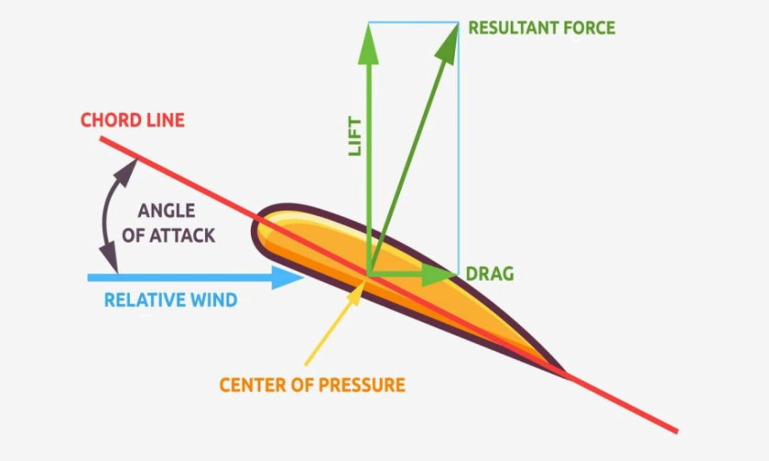

Figure 2 describes how an ideal wing looks with the aerodynamic forces and properties acting on it and where they are located.

The Navier-Stokes equations are used to calculate aerodynamic forces for traditional solvers and so they can be the reason the solvers are more complicated than necessary for the rear wing. While the Navier-Stokes equations are a very important differential equation in fluids, this project will be using the thin airfoil theory in order to neglect the compressibility and reduce complexity to make prediction simpler and more effective. This makes the thickness of the wing small and wingspan long so that the fluid flow is incompressible and inviscid (no viscosity/thickness, doesn't resist force) (Cadence CFD Solutions, 2023). If thin airfoil theory is used, then the coefficient of lift (CL) is 2 times π times angle of attack as shown in the equations below.

Equation 1:

Coefficient of Lift Relation to Angle of Attack (Cadence CFD Solutions, 2023)

$$C_L = 2\pi((a - a_0)(180/\pi))$$

Leveraging the thin airfoil theory will provide basic physics underlying the relationships between aerodynamic coefficients and airfoil properties. This will allow the design to be able to

compare the coefficients of lift and drag in association with the geometric configuration of the wing such as its camber to make it possible to generate sample data points. Below are the equations for the forces of drag and lift which are to be minimized or maximized:

Equations 2 and 3:

Equations for Forces of Drag and Lift (NASA Glenn Research Center, 2024)

$$F_{drag} = \frac{1}{2} \cdot \rho \cdot C_d \cdot A \cdot v_x^2$$

$$F_{lift} = \frac{1}{2} \cdot \rho \cdot C_l \cdot A \cdot v_x^2$$

where:

ρ: Air density;

C_d and C_l: Drag and lift coefficients;

A: Vehicle frontal area.

Another equation similar to the thin airfoil equation but for drag is the quadratic drag polar relationship that relates the coefficient of lift (CL) to the coefficient of drag (CD) (De Marco, 2025). By utilizing this connection, the prediction of CD can also be rooted in physics. It is shown in Equation 4.

Equation 4:

Coefficient of Drag Relation to Coefficient of Lift, Drag Polar Relationship (De Marco, 2025)

$$C_D = C_{D0} + KC_L^2$$

However, with these minimizations and maximizations comes certain tradeoffs on the track. Higher downforce (inverse of lift) is beneficial for vehicle control, corner performance, and braking, but causes an increase in both drag and straightline speed. Increasing angle of attack or

camber (wing curvature) can increase lift and the magnitude of downforce thus increasing drag as well which needs to be minimized for energy consumption. This complicates things in terms of optimal setup. In Mustafa Cakir's 2012 study using CFDs on a rear wing, he lists the various forces acting on a car such as drag, lift, and gravity (Cakir, 2012). He also details the equations for the coefficients of drag and lift for rear wings by manipulating equations 2 and 3 above to solve for CD and CL.

Coefficients of drag and lift are preferred since they are unitless and are therefore universal regardless of metric system. Cakir also details the benefits of CFDs and how they are commonly preferred over wind tunnel testing since it is cost and time effective comparatively. A company that spends $100,000 on wind tunnel testing would only spend $20,000 for CFD testing, highlighting a five time decrease in cost. Similarly, in regards to time, there are improvements up to 30% for CFD in comparison to wind tunnel testing (Engineering Daily, 2025).
However, a reduced-order model would be able to specialize itself in a certain aspect of aerodynamics and be able to cut down time and computational cost even more which is what this project intends to accomplish for the rear wing subsystem.

Figure 3:

CFDs as the Intersection of Various Fields (Cakir, 2012)

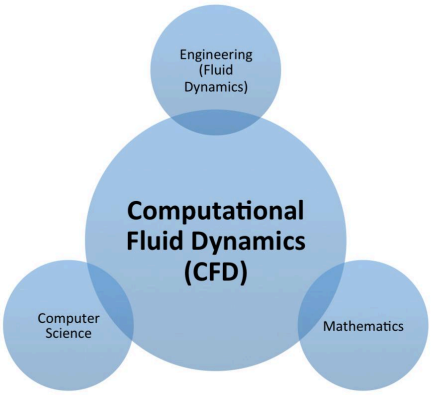

Figure 3 signifies that CFDs integrate aspects from computer science, mathematics, and engineering with fluid dynamics.

Cakir's results show that the CFD was able to numerically analyze different aerodynamic properties of a spoiler on a car. Every time a wing generates downforce, the inverse of lift, it generates drag as well and this is a primary trade off (Cakir, 2012). Passenger cars would benefit from greater downforce as it improves traction but since it generates drag, it slows down the car which high speed racing cars do not benefit from. The use of active aerodynamics will try to omit the cons of balancing different aerodynamic coefficients by manipulating the wing at different times. Regulation of aerodynamics yields a balanced output, not favoring drag nor downforce at a time and allowing for alternate settings of a wing thus overstepping many tradeoffs.

Methods for Modeling (Surrogates/ROM)

Throughout existing literature, there is a prevalence of reduced-order modeling used as quick and accurate alternatives to traditional CFD or wind tunnel methodologies. "Reduced-Order Modeling (ROM) Approach to Blunt-Body Aerodynamic Modeling" by Dean et. al for AeroSpace Research Central introduces reduced-order modeling for modeling complex dynamic systems found with blunt-body vehicles for example (Dean et. al, 2025). The scientists attempt to bypass calculating dynamic stability coefficients with CFD simulations by using cost-effective ROMs as their main purpose. It is very similar to what this project intends to do with active rear wings. There is great promise in ROMs for modeling the dynamic aerodynamics in the results. They not only maintain accuracy for each configuration but also drastically improve the time efficiency (Ghoreyshi et. al, 2014). Proper orthogonal decomposition (POD) is a common way to reduce dimensionality using proper and approximate the aerodynamic

coefficients for unseen trajectories based on low MAE values during evaluation. POD works by taking the most significant modes (patterns) of a data to still retain information of the original model and uses singular value decomposition, a linear algebra technique, to capture the most amount of variance in the data with a smaller number of dimensions (Dean et. al, 2025). This project will use a similar methodology and will base the model (design) off CFD data and use it for calculating aerodynamic coefficients to be used in dynamic systems. This project's ability to capture unseen trajectories for systems as complex as blunt-body aerodynamics shows strong validation towards ROMs.

Furthermore, an experiment by Li et. al states that ROMs fit under 2 categories: system-identification-based methods and projection-based methods (Li et. al, 2021). The former utilizes an input-output mechanism that is extracted and measured by the source of a system that is of interest. It described the more complicated, high-fidelity, full order system through a mathematical equation, making it a strong choice for this project's design, since this project will be extracting data from the track and want to get quality results at a low cost and short time. Additionally the model will be basing its predictions off an input-output mechanism between wing parameters and aerodynamic coefficients. According to Li et. al's findings, some common models for prediction for aerodynamic reduced-order models include neural networks, support vector machines (SVM) and Volterra series model (Li et. al, 2021). These approaches are possible design options for this project's ROM.

The latter category of ROM approaches the task by breaking down the high-dimensional flow data into many smaller low-order parts through Proper Orthogonal Decomposition (POD) similar to Dean et. al's study (Dean et. al, 2025). The unsteady nature of airflow has proven to be

difficult to work with, so using one of these models can prove to be largely beneficial, especially with the ample amounts of AI reference models available.

In searching literature for surrogate models with similar inputs/outputs, a study by Andres-Perez et. al was found (Andres-Perez et. al, 2021). The paper modeled 3 aeronautical configurations:

1. NACA0012 airfoil
2. RAE2822 airfoil
3. 3D DPW wing

Of the three, the NACA0012 airfoil stands out as being a common and simple baseline airfoil with 0% max camber. Linear regression models are simple functions of the following structure (Figure 4), where w and x are features from the data set and can be used as a simple mathematical model for predicting outputs given a set of inputs:

Figure 4:

Baseline Linear Regression Model as Example of Mathematical Models (Andres-Perez et. al, 2021)

$$\hat{y}\left(\vec{w}, \vec{x}\right) = w_0 + w_1 x_1 + w_2 x_2 + \cdots + w_n x_n,$$

Figure 4 is how the model class for a linear regression model is constructed. From this, parameters (weights/biases) are learned.

There are also other, more complicated vector related regression optimization algorithms that can be potentially utilized for this project's design if linear regression turns out to be inadequate, such as:

- Least Squares
- Ridge Regression
- Least Absolute Shrinkage and Selection Operator (LASSO)

- Least Angle Regression (LARS)
- Bayesian Ridge

Neural networks can be essential to this project as well, since they are an integral and easy architecture for modern machine learning. This can be achieved with the help of Radial Basis Function Neural Networks (RBFNNs) (Andres-Perez et. al, 2021). These are neural networks that implement radial basis functions during activation which produce an output from a given distance away from a center point. These distances and output values are trained during the training process of the RBFNNs and can be much faster for training than backpropagation and more easy to generalize since it is similar to clustering techniques. A more generalizable and faster model is valuable for this design. These examples of common ROMs and dimensionality reduction techniques are essential for the design of the project and training an adequate ML model to replace CFD for real-time aerodynamic coefficient prediction.

Potential and Previous Uses of Machine Learning in the Control of Active Aero

Machine learning is a common and successful supplement to aerodynamic control and active aerodynamics (Renn & Gharib, 2022). Model-free reinforcement learning is a common method to allow for aerodynamic control of unmanned air vehicles like a UAV or wind turbines in turbulent wind conditions. Wind conditions are determined by Reynolds's number which determines if the air flow is turbulent or laminar. In a previous study, the scientists used two models (TD3 and LSTM-TD3) with LSTM-TD3 performing better (Renn & Gharib, 2022). LSTM-TD3 performed better due to the fact that it was a recurrent neural network with sequential steps and since the full flow of the dynamics weren't defined, TD3 wasn't able to infer and learn the state/physics of the system (Renn & Gharib, 2022). This means the design for this project would benefit from possibly being a recurrent neural network and being trained using

reinforcement learning for the controller. However, there is a big downside in time and data which are key components for the project's success criteria. They can effectively learn the dynamics of nonlinear and chaotic systems which this project's dynamic system may exhibit; however, it is computationally expensive. Control-based techniques similar to the ones used in Renn's paper could be applied to the simulated rear wing in this project and the system will be less chaotic since it is on a race track, not turbulent wind conditions.

Various sources also prove the benefits of active aerodynamic wings over other wings like single or dual elements (Occam's Racer, 2023). The experiment simulates lap times with different wing curvatures that have active aero zones in which the wing changes configurations. Occam's Racer's experiment could be used for possible lap time simulation options and performance evaluations but this project would expand by creating a design for a machine learning alternate model and having automated controllers instead of fixed which was used in this source (Occam's Racer, 2023). Wings designed for low drag were used during their experiment. Dual element wings, which are common in motor sports now, have two fixed positions: one for downforce and one for low drag (Occam's Racer, 2023). This project will have a range of options instead of just these two fixed configurations, creating a continuous geometric parameterization. There are plenty of active aero wings from Carbon Miata, NINTE, TRK1 that could benefit from being less computationally expensive. Occam Racer's experiment was simulated using OptimumLap simulations and this project could also do the same in addition to dynamic systems simulation. Dual element active aero wings were found to perform faster than fixed wings but only by seconds. The aforementioned papers demonstrate the practical benefits of machine learning in aerodynamic control and active aerodynamics as well. They support this design's purpose of pairing the two for real-time control for the rear wing.

Interactions With Vehicles and Optimal Lap Times

Vehicle dynamics and physical modelling also plays a significant role in the solutions of common aerodynamics problems such as the Minimum Lap Time Problem. These dynamics will serve as evaluation of the design. “A Survey of Vehicle Dynamics Modeling Methods for Autonomous Racing: Theoretical Models, Physical/Virtual Platforms, and Perspectives” by Zhang et. al (Zhang et. al, 2024) provides existing vehicle dynamic models but highlights their lack of inclusion of challenging driving conditions with high speeds common in racing. It serves as guidance for what dynamic model to choose depending on the driver's situation and provides this project with key examples of relationships between state variables and which ones to implement. The paper discusses control algorithms which will be helpful for this project. It also summarizes the current state of specific models. It further emphasizes that control algorithms should be aware of vehicle dynamics in order to prepare for things like tire slips (Zhang et. al, 2024). Current vehicle dynamic models may be applicable to autonomous vehicles, but due to their computational inefficiency are inadequate for autonomous racing. This is the gap that this design intends to address.

Lap time simulation and optimization is essential for racing and motor sports. It is crucial for design evaluation. Advancements like active aerodynamics further require the need of lap time optimizations. The paper by de Buck & Martins addresses gaps in minimum lap time methodologies (de Buck & Martins, 2022). It details the system dynamics (state variables are position or distance along centerline, heading angle, and distance normal from the centerline of a racetrack) and control problems with steering wheel angle being controlled. The ODEs for each positional state variable are:

Figure 5:

ODEs for position, heading angle, and distance from centerline used in de Buck & Martins' experiment (de Buck & Martins, 2022)

$$\dot{s} = \frac{V\cos(\alpha - \lambda)}{1 - n\kappa}$$
$$\dot{n} = V\sin(\alpha - \lambda)$$
$$\dot{\alpha} = \Omega - \frac{\kappa V\cos(\alpha - \lambda)}{1 - n\kappa}$$

These equations shown in Figure 5 demonstrate existing ODEs for common state variables (change in distance along centerline, change in distance from centerline, and change in heading angle) and many of the existing ODEs similar to these can be integrated and paired to create the ODEs for each state variable in the dynamics system evaluation of the design. These states specifically govern the vehicle's position along a track (de Buck and Martins, 2022).

For example, the change in position along the centerline of a vehicle (ṡ) is dependent on the current heading angle (**α**) and current distance from centerline (n) and curvature (κ) among others. Additionally it is shown that the change in heading angle of a vehicle is proportional to the curvature of the track, kappa (de Buck and Martins, 2022). These current values determine how the position will change based on the given equation and defining this relationship is the purpose of an ODE. Solving an ODE allows for the tracking of different variables across long timeframes.

In a paper by Elbal et. al, the experiments prove that simulation techniques are effective in testing and optimizing vehicle parameters especially for the Minimum Lap Time problem (Elbal et. al, 2024). The methodology used in this paper features vehicle control, vehicle dynamics, and active aerodynamics to optimize certain parameters for the Minimum Lap Time problem. The Minimum Lap Time problem is a very common problem statement that models are trained to be enhanced for and it serves as a baseline framework for motorsports optimization. The results of Elbal et. al's study demonstrates improvements in energy reduction and other aerodynamics areas and proves the effectiveness of simulation techniques for vehicle performance optimization. To summarize, this paper reinforces the necessity of a simulation of something like the Minimum Lap Time problem for evaluating the design this project creates. It is clear evidence that continuous active aerodynamics,governed by a control system, is optimal for the Minimum Lap Time problem. A Minimum Lap Time simulation framework is necessary for assessing the validity of a design and it is underscored by the use of several common ODEs for states like position and heading angle (Elbal et. al, 2024). This paired with a reduced-order model can optimize for real time control.

De Buck & Martins, in their experiment, assume a linear relationship with aerodynamic coefficients, therefore this project improves by using a machine learning model to determine the aerodynamic coefficients more accurately. De Buck and Martins also highlight how rear wing states improve performance when in the control loop. Active aerodynamics was found to significantly reduce lap times and the relationships between vehicle subsystems is important for new complex technologies such as active aerodynamics (de Buck and Martins, 2022). Trajectory optimization was important in developing optimal control systems for both papers. By replacing slow aerodynamic evaluations, the reduced-order model would improve on the results found by

de Buck and Martins. The reduced-order model should be able to optimize accordingly and perform similarly to that of the results found by Elbal and his group in the experiment (Elbal et. al, 2024). The simulations will be able to leverage different configurations and tracks to determine the performance of various control system/reduced-order model designs for this project. Both de Buck and Elbal et. al's methodologies include vehicle dynamics and control thus providing a basis for the design which will expand upon and also model dynamically instead of with quasi-steady-state assumptions (Elbal et. al, 2024; de Buck and Martins, 2022).

Common minimum lap time simulation techniques benefit from vehicle dynamics modelling over quasi-steady-state modelling (Massaro & Limebeer, 2021). Many current simulations use quasi-steady state models (not dynamic so state variables are constant). This is a simplified approach, one this design will improve upon through the use of vehicle dynamics simulations. Dynamic simulations are too computationally expensive, albeit more accurate (Massaro and Limebeer, 2021). Thus, the reduced-order model design in tandem with the existing vehicle dynamics equations from Zhang et. al's survey can address the minimum lap time problem without the simplified quasi-steady state model approach (Zhang et. al, 2024). This is more accurate while also being computationally efficient. Additionally, it is conveyed that drivers are commonly looked at as the control system in simulations; for this design, however, the control system will be the active aerodynamic rear wing that leverages the reduced-order model to make informed optimization decisions.

Finally, active rear wings have been shown in the literature to be much more beneficial compared to passive wings in terms of energy efficiency (Georgiev et. al, 2023). There are changes in the active wing for downforce maximization or drag maximization, similar to the purposes of this design, and it is shown that the active wing outperforms the passive in terms of

braking distance reduction, regenerated energy increase, and deceleration. The relevance to this design is in how aerodynamic coefficients can be related as functions of angle of attack and these can be used for determining optimal features of different aerodynamic devices such as possibly a rear wing (Georgiev et. al, 2023). This validates the use of aerodynamic coefficients as input for the control loop with changes in the angle of attack and other parameters as outputs.

Overall the background research proves that active aerodynamics is beneficial for reducing and optimizing lap times and performance of cars. Furthermore, CFD simulations for active aerodynamics are very computationally expensive and so there would be very big benefits in creating reduced-order models for the same task. Additionally, control algorithms and vehicle dynamic models are common for optimizing and simulating performance for active aerodynamics, an emerging technology for motor sports. There are plenty of techniques for these algorithms and existing models detailed in the literature to support this project. Thus this project would be very beneficial for expanding upon the idea of creating a reduced-order model for CFD tasks. It has been consistently shown there is substantial potential in improving vehicle performance with active aerodynamic configurations and the aerodynamic control that underlies it (Rijns et. al, 2025; Occam's Racer, 2023; Georgiev et. al, 2023). This project expands on all these previous methodologies by integrating the ROM with a control system for the rear wing subsystem of a race vehicle. Finally, evaluating its performance with a dynamic system that is guided by the Minimum Lap Time problem statement framework simulation is a very feasible and effective method for truly determining the validity and extent to which the design is actually worth using over traditional CFDs and fixed wings.

Design Plan

Figure 6:

Design Flow Chart for Creation and Control Loop of the ROM/Controller

Figure 6 shows two flow charts depicting the procedural process for creating and iterating the design. It involves four subsystems with data collection, ROM implementation, control system design, and vehicle simulation. There are multiple interacting and moving parts.

The parameters refer to the shape of the wing and are continuous parameterizations instead of fixed, since active aerodynamic wings involve modifying/adjusting the wing in small increments rather than fixed steps like previous active aero wings. The parameters that will be modified are the wing features that the control system optimizes each loop based on the reduced-order model's predictions. The machine learning model itself will have its own parameters such as weights and biases while the physics-informed section of the model will have parameters that are to be fitted by least squares. It will also have hyperparameters such as batch size and learning rate that will be configured with hyperparameter searches accordingly. The wing features will directly impact aerodynamic coefficients like drag or lift or downforce and are as follows:

a. Wing Curvature (Camber Percentage)
b. Angle of Attack (Degrees)

The angle of attack determines the overall angle of the wing while camber is how curved the wing is according to a base line.

During lap time simulation there will be operating parameters for speed, Reynold's number (determines relative turbulence of airflow), steering angle, general air/wind, coefficient of friction, weight.

Both coefficients of lift and drag and Reynold's number are dimensionless while camber is a percentage so it is also dimensionless. Angle of attack is measured in degrees.

Constructing and Testing the Prototype

Materials

- Laptop/Computer for running XFOIL and model
- MIT's open source XFOIL as the CFD source (program for analysis and visualization of fluid related data on airfoils) for getting data and aerodynamic values https://web.mit.edu/drela/Public/web/xfoil/
- X11/XQuartz GUI System for running 2D CFD visualizations through XFoil
- Google Colab Notebook for running Python with necessary libraries
- Scikit-learn and PyTorch for creating the model
- NumPy, Pandas, Matplotlib for visualizing and handling data/results

Procedure

Data Collection in XFOIL

1. Download xfoil6.97.tar.gz through https://web.mit.edu/drela/Public/web/xfoil/
2. Download .dmg file for XQuartz in mac or whichever Operating System is being used through same link
3. Through the terminal use tar -xzf xfoil6.97.tar.gz to extract the downloaded xfoil file
4. Prepare the computer for XFoil usage by editing XFoil Makefiles. Follow this youtube tutorial to do so (https://youtu.be/XdaRlTHf7BE?si=WzKCSzHumTWbQKIl)
5. Open XFoil in the terminal by typing xfoil in PATH of terminal
6. Select NACA 2412 (2% camber) by typing NACA 2412
7. Enter operating mode by typing OPER
8. Enable viscous flow calculations by typing VISC

9. Set Reynolds number to 3*10^5 by typing 3E5
10. Initiate PACC mode by typing PACC into terminal (polar accumulation mode)
11. Enter polar.txt file name based on current wing/Reynolds number
12. Click enter again to avoid polar dump file
13. Perform ASEQ -5 15 0.5 (angle of attack [AoA] sweep from -5 degrees to 15 degrees with step of 0.5)
14. Exit PACC mode to save files
15. Enter OPER mode again
16. Use VISC and Re of 5*10^5, 5E5
17. Repeat steps 10-15
18. Use VISC for Re of 8*10^5, 8E5
19. Repeat steps 10-15
20. Use VISC for Re of 1*10^6, 1E6
21. Repeat steps 10-15 (Total of 4 polar files for NACA 2412)
22. Exit OPER, return to XFOIL using QUIT
23. Load NACA 1212 (1% camber) by typing NACA 1212
24. Repeat steps 7-22 with NACA 1212 for 4 additional polar files
25. Load NACA 4412 (4% camber) by typing NACA 4412
26. Repeat steps 7-22 with NACA 4412 for 4 additional polar files
27. Load NACA 0012 (0% camber) by typing NACA 0012
28. Repeat steps 7-22 with NACA 0012 for 4 final polar files (16 total)

Python Preprocessing

29. Create Python script to read each file using pandas and identify non-converging values (missing a row)
30. Same script to remove duplicate values, invalid data points (not a number)
31. Extract relevant columns (camber, angle of attack, coefficient of lift, coefficient of drag, Reynold's number)
32. Put preprocessed data with relevant columns into single CSV file separated by camber (based on airfoil), angle of attack (AoA) value, Reynolds number (Re) value, calculated coefficient of lift (CL), calculated coefficient of drag (CD)
33. Create training dataset by filtering for datapoints where Re is not 1*10^6 (around 500 datapoints)
34. Create testing dataset for datapoints where Re is 1*10^6 (around 180 datapoints)

Reduced Order Model Design (Physics + Neural Network)

35. Create initial physics-based CL predictor using parameterized thin-airfoil theory equation: $CL = CL_{\alpha} + (\alpha + \alpha_0)$ where $CL_{\alpha} = \alpha_0 + \alpha_1 * c$ and $\alpha_0 = \beta_0 + \beta_1 * c$. (Coefficient of Lift in terms of Angle of Attack [α], Camber [c], and Parameters [α_0, α_1, β_0, β_1])
36. Create initial physics-based CD predictor using parameterized drag-polar relationship: $CD = CD_0 + k * CL^2$ where $CD_0 = \gamma_0 + \gamma_1 * c + \gamma_2 * log(Re)$ and $k = k_0 + k_1 * c$. Parameterized Drag Polar Equation (Coefficient of Drag in terms of Reynolds Number [Re], Camber [c], Coefficient of Lift [CL], and Parameters [γ_0, γ_1, γ_2, k_0, k_1])

37. Fit the parameters in the parameterized models with least squares learning algorithm from numpy
38. Create linear sequential neural network design (input layer with neurons for each wing value, two hidden layers with hyperbolic tangent activations, 3 hidden linear layers, output layer with 2 neurons for the corresponding CL and CD residuals)
39. Use ADAM optimizer and Mean Squared Error (MSE) as loss function to train
40. Add custom layer for using physics-predicted lift as additional input to improve drag prediction
41. Enforce physical consistency checks within the neural network constraining drag to be positive
42. Save the calibrated model weights for future use (torch.save())
43. Update final CL and CD predictions to be calculated as physics predicted values + neural network residuals
44. Track MAE loss, R^2 scores, time to run for performance metrics on unseen data (testing set) using sci-kit learn to calculate

Vehicle Dynamics Setup

45. Define states (like position (x,y), velocity (v), and yaw) sufficient to evaluate lap time and energy under consistent track conditions.
46. Implement with ordinary differential equations (ODEs) for vehicle dynamics that take aerodynamic forces into account for longitudinal/lateral dynamics
47. Define track curvature as a function (return kappa values based on current position (x,y)
48. Set initial trajectories for the simulation

Controller Creation

49. Define a weighted objective to maximize downforce in corners and minimize drag on straights
50. Create an objective-driven controller that queries the physics model to predict current aerodynamic coefficients at each step
51. Using this information ensure the controller optimizes wing geometry continuously based on weighted objective
52. Implement rate limiters and position constraints on wing adjustments to ensure realistic actuator capabilities (maximum change per time step, physical angle limits)
53. Include limit constraints in both the controller and dynamics to avoid unrealistic changes/calculations
54. Integrate model predictions using its calls into ODE solver to update aerodynamics forces at each step

Testing and Reiteration

55. Simulate complete lap, tracking lap time, total energy loss, downforce, drag, speed over time
56. Run baseline simulation with fixed wing geometry (constant AoA and camber) with same track conditions
57. Simulate for multiple trials (around 5 runs) for each iteration/baseline
58. Calculate MAE, Mean Absolute Percentage Error (MAPE), R^2 scores and compare reduced order model predictions to XFoil predictions on test set
59. Compare lap times/energy loss between active aero controller and baseline fixed wing
60. Create data visualization plots: model prediction error heatmaps, lap time comparisons, speedup comparisons, speed over track position, drag/downforce over track position

61. Identify metrics that lack completion of criteria
62. Modify neural network, controller definition based on identified issues
63. Retrain/re-evaluate updated iteration
64. Compare against previous iteration to see improvement
65. Find best design after multiple iterations and report findings

Results and Discussion

Results of Testing and Redesign (Iteration 1)

The Physics-Informed Reduced-Order Model (PI-ROM) and controller design was tested over three iterations. The first iteration fitted the model's parameters and discovered weights for the neural network residual. The second iteration improved the performance of drag prediction by not only using coefficient of lift for the physics equation but also as an input for the neural network. The final iteration of the design focused on improving the vehicle dynamics simulation with tire friction and speed caps, it also changed the controller to be objective-driven in terms of drag and downforce rather than only modifying based on curvature of the track.

In the first iteration, The thin-airfoil theory equations used for the model's prediction of CL describe the closely linear relationship with lift and angle of attack. CL is important to bring for calculating downforce, the inverse of lift which is important in a motorsports environment. Drag prediction is based on the quadratic form drag-polar equation which meant it had more parameters and complexity in its parameterization. The parameterization of the thin-airfoil equation (Equation 1) is shown in Equation 5 and the parametrization of the drag-polar equation (Equation 4) is shown in Equation 6. The PI-ROM works by fitting the parameters in these equations using the least squares regression algorithm similar to Figure 4 as a baseline and then

tweaking it for more accuracy with a neural network as a residual. The residual is especially useful in non-ideal conditions which stray away from the regimes these equations model.

Equation 5:

Parameterized Thin Airfoil Equation (Coefficient of Lift in terms of Angle of Attack [α], Camber [c], and Parameters [α_0, α_1, β_0, β_1])

$$CL = CL_{\alpha} + (\alpha + \alpha_0),$$

where

$$CL_{\alpha} = \alpha_0 + \alpha_1 * c$$

and

$$\alpha_0 = \beta_0 + \beta_1 * c.$$

Equation 5 is the model class for the lift coefficient prediction and involves four parameters and two data inputs.

Equations 6:

Parameterized Drag Polar Equation (Coefficient of Drag in terms of Reynolds Number [Re], Camber [c], Coefficient of Lift [CL], and Parameters [γ_0, γ_1, γ_2, k_0, k_1])

$$CD = CD_0 + k * CL^2,$$

where

$$CD_0 = \gamma_0 + \gamma_1 * c + \gamma_2 * log(Re)$$

and

$$k = k_0 + k_1 * c.$$

Equation 6 is the model class for the drag coefficient prediction and involves five parameters with three data inputs.

Furthermore, drag contains more underlying complexity in its calculations due to numerous outside influences like friction being at hand that do not affect lift. This contributed to its poorer initial results and the need for a redesign in iteration 2. The prediction results on unseen data for the first iteration are shown in Table 1.

Table 1:

Aerodynamic Coefficient Prediction Metrics, First Iteration

Model Evaluation Metric	**Lift**	**Drag**
MAE	0.0298	0.0046
R2	0.9959	0.5583
Range-Normalized MAPE	1.370%	11.810%

Both lift and drag predictions were able to achieve low mean absolute error on the testing dataset (Reynold's number of 1,000,000 for all airfoils) of 0.0298 and 0.0046 for CL (Coefficient of Lift) and CD (Coefficient of Drag) respectively. However, when looking at the range-normalized MAPE (Mean Absolute Percentage Error) values which take into account the scale of the data, the results for drag prediction are not as accurate. Since Mean Absolute Error is based on the difference between true and predicted values, if these values are inherently small like they are for aerodynamic coefficients, the Mean Absolute Error may be deceivingly low and not a true indicator of accuracy. This is why a normalized percentage error is useful and while this value is low for lift, the percentage for drag prediction does not quite meet the criteria for the PI-ROM design. Table 1 also displays the R2 scores which indicates how well the model was

able to fit with the data. The closer these values are to +1 means a stronger relationship and for lift, the variability in the data is almost perfectly predicted by the model.

Figure 7:

Error Heatmap of PI-ROM Iteration 1 in CL Prediction

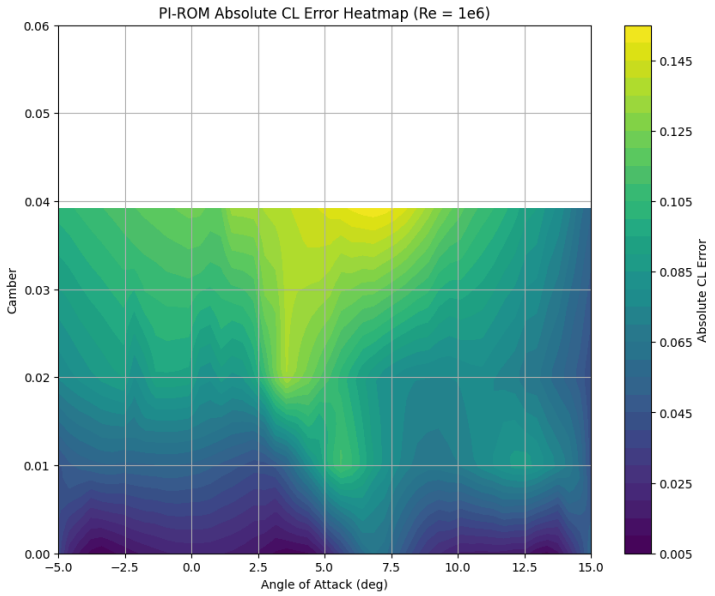

The error heatmap (Figure 7) demonstrates for CL prediction that as camber increases, error did too. Error heatmaps provide a better understanding of where the model falls short and where errors may arise in prediction which can be improved in future designs.

A valid reasoning behind the patterns shown in Figure 7 is that higher cambers move away from the ideal thin airfoil making the linear relationship for predicting CL less and less accurate. However, this error is not irrationally high as it could be due to the work done by the neural network residual.

Figure 8:

Error Heatmap of PI-ROM Iteration 1 for CD Prediction

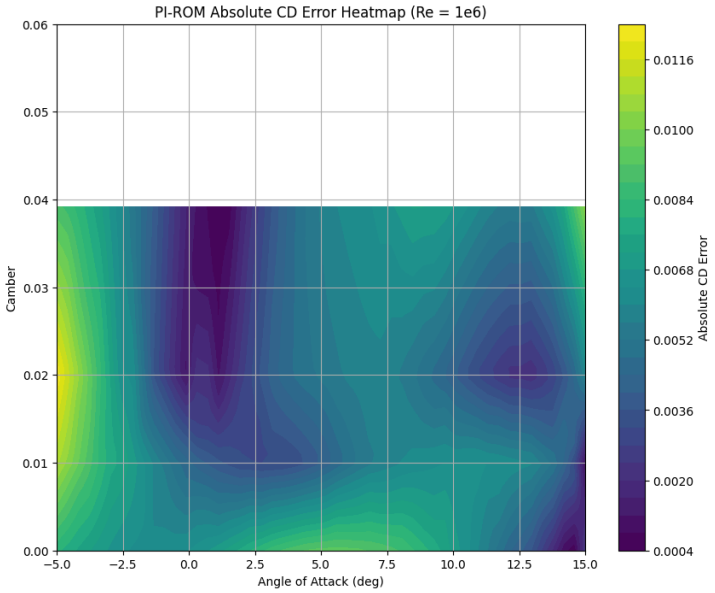

Figure 8 shows how the drag's error was highest (most concentrated on heatmap) near the left and bottom boundaries or was highest at an angle of attack of -5 degrees and camber of 0.00. This possibly demonstrates that the model struggled to predict the drag coefficient at boundary locations and a wider range of samples could alleviate this.

Figure 9:

PI-ROM Iteration 1 Lift Predictions over Angle of Attack and Camber Grid

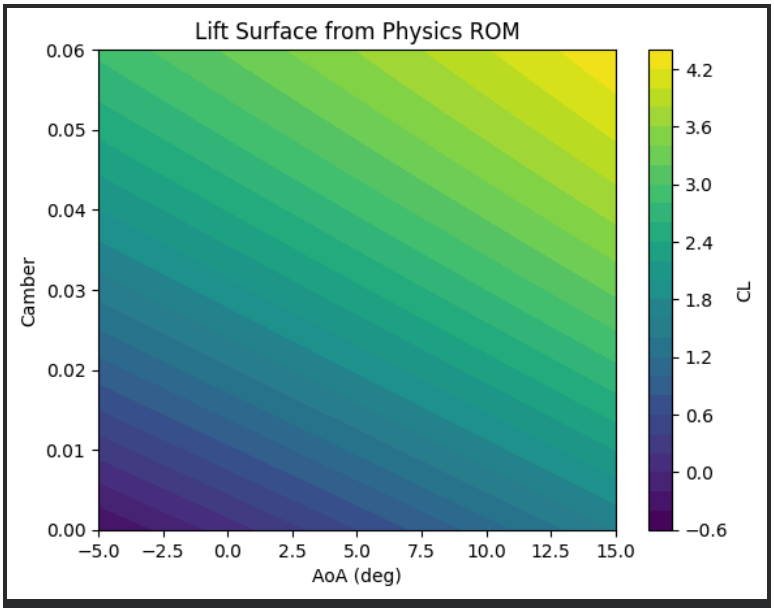

Figure 9 shows what the lift predicted was for a meshgrid of angle of attack values on the x-axis and camber on the y-axis.

The Lift Surface (Figure 9) also backs the validity of the model since according to literature and physics, as camber or angle of attack increases so should lift. This is shown in the heat map as the brighter (higher values of lift) correspond to the top right which means highest angle of attack and highest camber. Thus, since lift is shown to be the highest in the regions of highest angle of attack/camber the physics underlying the model check out and are not concerning.

Figure 10:

PI-ROM Iteration 1 Predictions vs. XFOIL Predictions of CL for Different AoAs

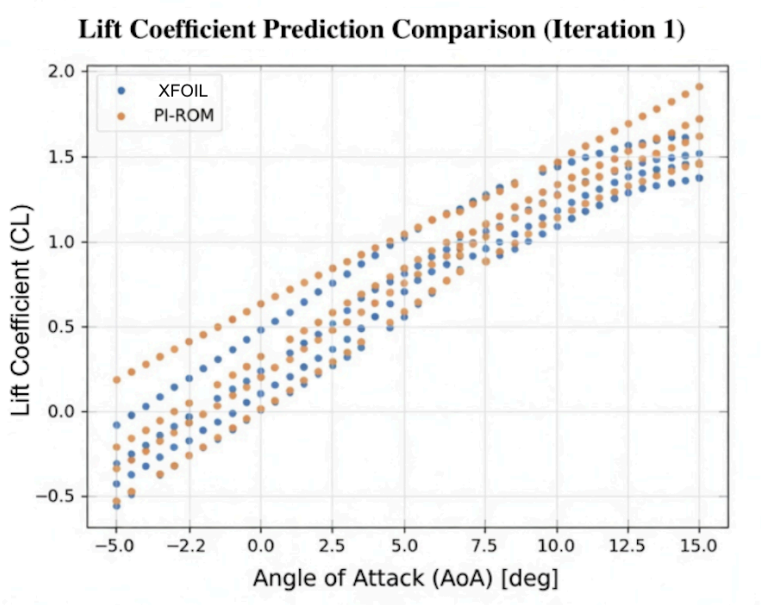

Figure 10 compares the predicted values of lift coefficient based on angle of attack for XFOIL and the design.

Figure 11:
PI-ROM Iteration 1 Lift Prediction True vs. Predicted Regression Scatter Parity Plot

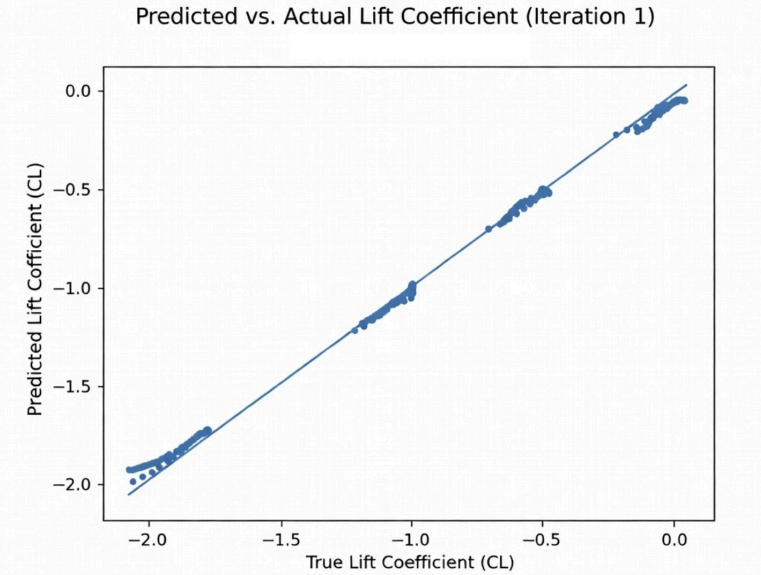

Figure 11 compares true vs predicted values for lift coefficient and the ideal line for this difference.

Figure 10 strongly demonstrates the linearly proportional thin airfoil relationship which both XFOIL and iteration 1 PI-ROM follow closely to predict the Lift Coefficient which is dimensionless. The scatter plot of true to predicted (Figure 11) shows how tightly centered around the line of best fit the predictions are. This signifies the potential to predict for airfoil configurations in the range as the model is able to explain variability and that the data obtained from XFOIL show clear patterns that the model is able to interpret. Meanwhile the accuracy of the PI-ROM is fortified in Figure 10 since the predictions overlap with XFOIL's predictions for almost all angles of attack. The PI-ROM only slightly mispredicts at the lowest and highest angle of attacks where the physics theory may not be as ideal. The goal of this model is to be faster but also similar to XFOIL in terms of accuracy and Figure 10 shows how aligned the values are while Table 2 compares prediction times.

Table 2:
XFoil and PI-ROM Prediction Times Comparison

XFoil Prediction Time	PI-ROM Prediction Time	Speedup Factor
2.4081 ms	0.0906 ms	26.58 times faster

Both the XFOIL and the PI-ROM were given the same conditions, a NACA 2412 airfoil (max camber of 0.02) with a Reynolds number of 1e6 and angle of attack of 5 degrees to calculate aerodynamic coefficients for. The average prediction time over 1000 such queries were calculated and are as shown in Table 2. The PI-ROM is thus shown to be much faster with a speedup (XFoil prediction time divided by PI-ROM prediction time) of 26.58 times faster and is an efficient replacement for XFOIL for an active aero wing. It successfully achieved the performance metric. 26.58 times is a factor that is monumental especially in terms of actual use cases where these calculations would be repeated. The small gains in time can add up to be drastic. This prediction time is consistent throughout iterations 2 and 3.

Figure 12:
Results of First Iteration Vehicle Dynamic Simulation and Wing Changes over Track

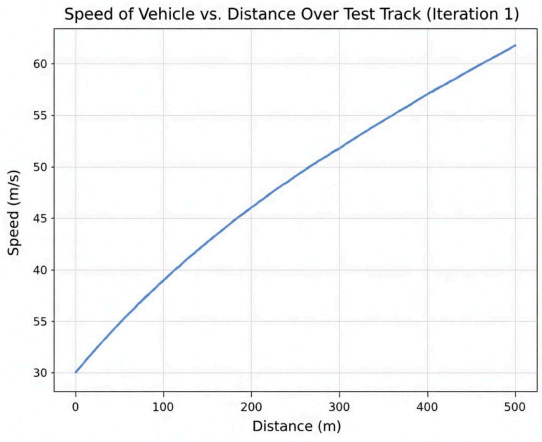

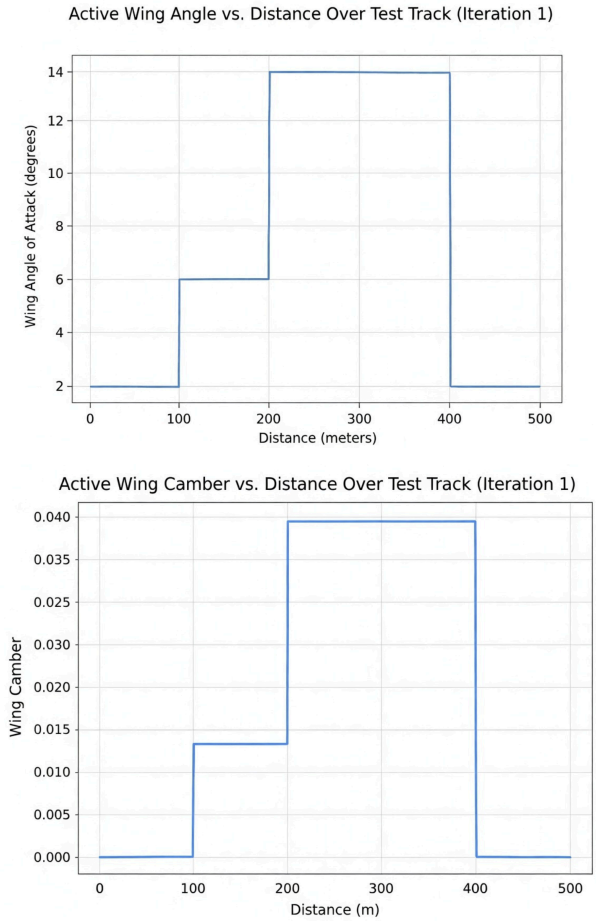

Figure 12 is the results of the first iteration vehicle dynamics simulation and shows how each feature the controller was changing as well as speed evolved over each meter on the track.

In this design, the vehicle dynamics involved simple one-dimensional longitudinal and lateral dynamics. For example, since the rear wing was the only focus, for lateral dynamics of heading angle/yaw, a bicycle model (one front wheel and one rear wheel) was utilized instead of an actual vehicle. This allowed simplicity in the dynamics and for the rear wing to essentially be moving along the test track itself during performance simulation. The track was a function where curvature (κ) was defined as values based on the position on the track (s) in meters (m). The track for testing starts without a curve until 100 meters, gets a slight curve for the next 100 meters, then goes to a very high curve for 120 meters before becoming more and more straight

while still being slightly curved till the finish line. During the track simulation (Figure 12) there is clear room for improvement in a redesign with a lap time of 10.84 seconds (shown in Table 4) being unfeasible in realistic conditions especially with how curved the track was made to be. There is lots of energy (161.9 kJ) lost to drag as well (found in Table 4). The speed to distance curve in Figure 12 shows an ever-increasing speed (meters per second [m/s]) which in reality should be dampened by downforce and have a limit. To be more aligned with reality and not have overestimated performance, the evolution of speed over time on the track must be changed. What is most in need of improvement, though, is the controller which applies the same logic for the camber percentage and angle of attack as seen in their change throughout the track being identical. This is to be improved to be more continuous, realistic, and aerodynamic optimization driven allowing for less energy consumption in the process as seen in iteration 3.

Figure 13:

Actual Drag Polar Curve (De Marco, 2025)

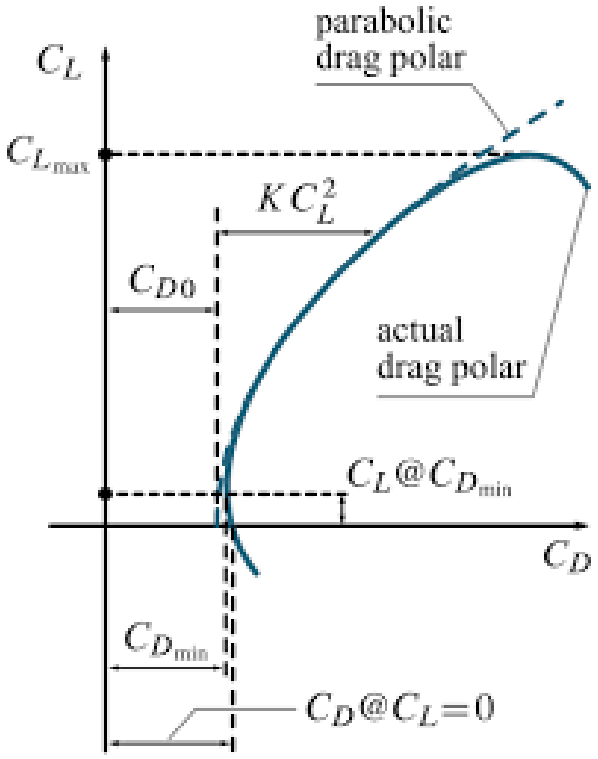

Figure 13 is the ideal drag-polar relationship plotted based on physics and follows a parabolic path.

Figure 14:

PI-ROM Iteration 1 Predictions vs. XFOIL Predictions of CD for Different CLs

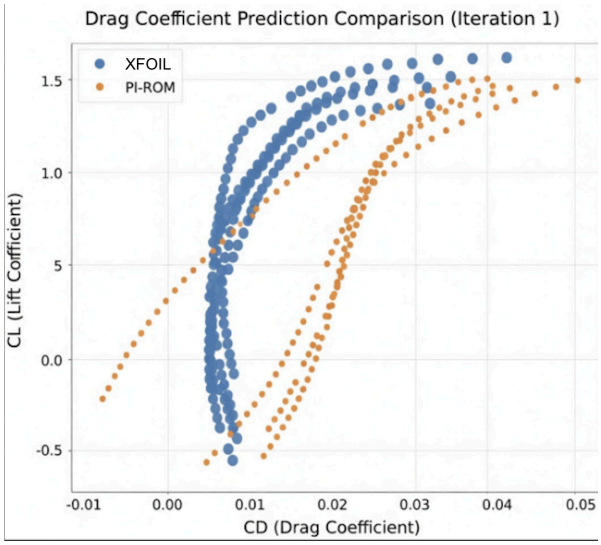

Figure 14 compares the predicted values of drag coefficient based on angle of attack for XFOIL and the design in its first iteration.

Figure 15:

PI-ROM Iteration 1 Drag Prediction True vs. Predicted Regression Scatter Parity Plot

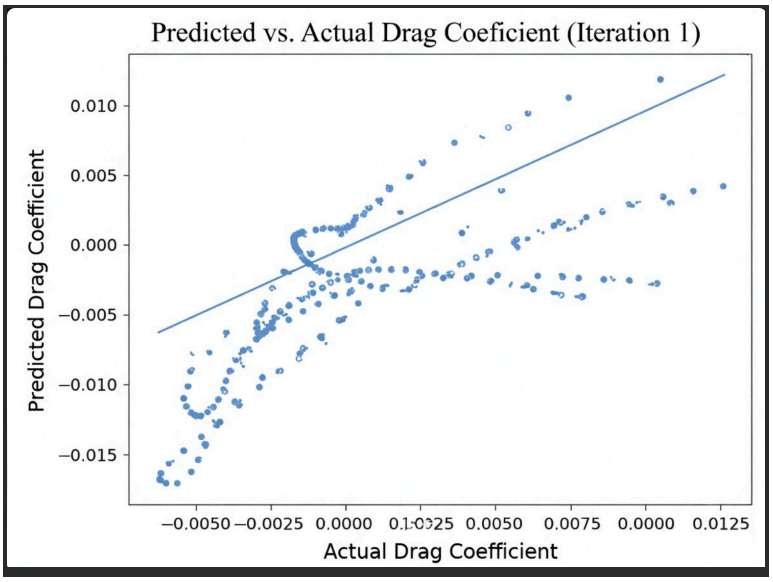

Figure 15 compares true vs predicted values for drag coefficient and the ideal line for this difference in the first iteration.

With the first iteration of the PI-ROM, drag predictions were much poorer especially in the ability to generalize compared to lift's predictions as seen in Table 1. In reality, the relationship between CD and CL should follow the curve depicted in Figure 13, however in iteration 1's predictions they instead follow a different structure shown by the orange scatter points. XFOIL's calculations marked in blue in Figure 14 do follow this reality but the PI-ROM does not as much as shown in orange within Figure 14. Furthermore, the scatter plot of true to predicted (Figure 15) shows the scatter points going in random directions without having a clear pattern as the CL's scatter plot did (Figure 11). The data does not fit as closely to the blue regression line as it should for performing according to the criteria. Finally and most importantly, it is illogical for there to be negative CD values which the model is predicting as seen in Figure 15 unlike XFOIL. This can err the vehicle dynamics and calculation for energy lost making them false. This is clear evidence that proper CD prediction should be prioritized and its improvement is the goal addressed by iteration 2. The results of iteration 2 are shown in Table 3

Redesign and Retest 1 (Iteration 2)

Table 3:
Updated PI-ROM Results Metrics for Iteration 2 and each Aerodynamic Coefficient

Model Evaluation Metric	**Lift, Iteration 1**	**Drag, Iteration 1**	**Drag, Iteration 2**
MAE	0.0298	0.0046	0.0034
R2	0.9959	0.5583	0.7467
Range-Normalized MAPE	1.370%	11.810%	8.730%

Figure 16:
PI-ROM Iteration 2 Drag Prediction True vs. Predicted Regression Scatter Parity Plot

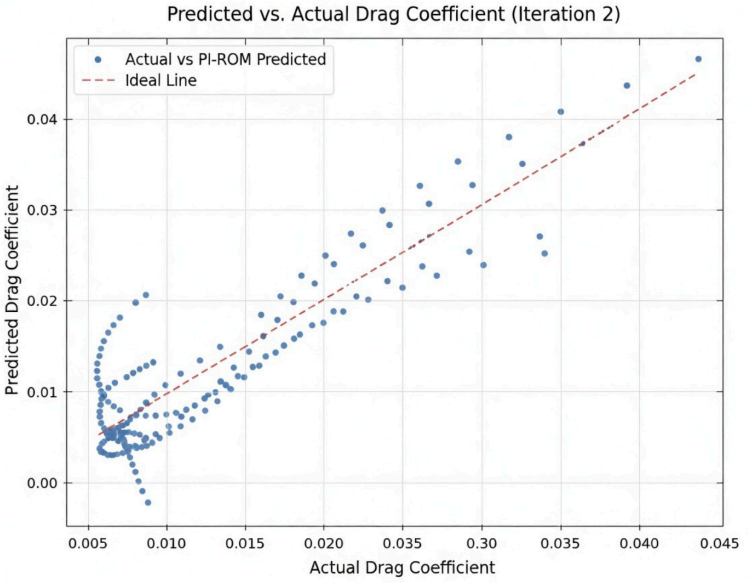

Figure 16 compares true vs predicted values for drag coefficient and the ideal line for this difference in the second iteration.

Figure 17:
PI-ROM Iteration 2 Predictions vs. XFOIL Predictions of CD for Different CLs

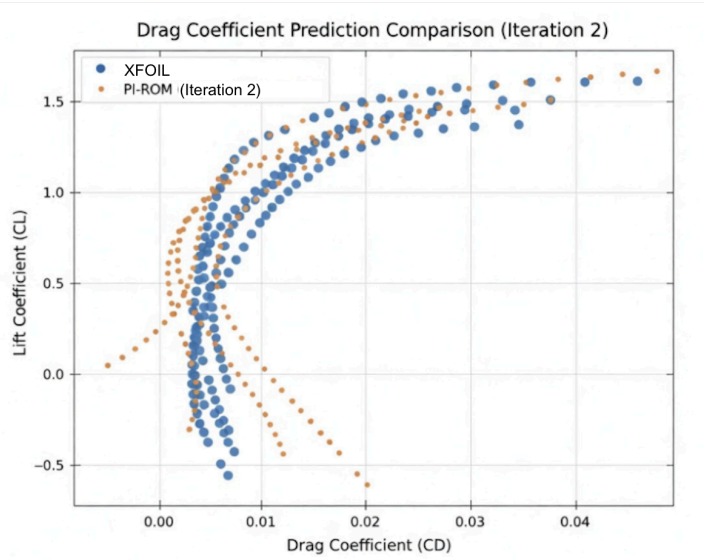

Figure 17 compares the predicted values of drag coefficient based on angle of attack for XFOIL and the design in its second iteration.

In this first redesign, the neural network residual section for predicting drag now takes in the predicted coefficient of lift from the physics equation as an input in addition to angle of attack and camber. Predicting CL did not really need a major residual change since its predictions were already so robust. The results of it in iteration 2 were no different to that of iteration 1 so these metrics nor the associated plots were not shown in Table 3 or as figures; the main improvement was in drag's metrics so prediction of lift was the same in this iteration as for the first. In iteration 2 the neural network residual is primarily focused on improving drag prediction. This did improve drag prediction since it had an MAE of 0.0046 and now it has 0.0034 and most importantly, the R2 score went from 0.5583 to 0.7467. This signifies the ability to fit XFOIL properly and improve accuracy. The scatter plot of true vs. predicted (Figure 16) now portrays a more patterned graph with the points following the trend of the ideal line which characterizes the higher R2 score. The percentage error dropped from around 11% to around 8% which now fits the performance criteria of a percentage error below 10% for accuracy. The PI-ROM still predicts in similar time, thus not requiring the need of an XFoil time comparison

for this iteration. The PI-ROM is also more aligned with XFOIL for depicting the drag-polar relationship and is not as severely predicting negative drag due to an included positive CD constraint (Figure 17). Now not only is XFOIL predicting Figure 13 but so is the PI-ROM making it more physics-aligned. Since only the PI-ROM model was altered, the vehicle dynamics simulation was run (results shown in Table 4) but its speed or wing changes were not plotted since neither the controller nor vehicle differential equations were changed.They were still the same and needed to be improved. These issues of the need for greater realism and speed caps can be addressed with improved drag prediction in iteration 3.

Redesign and Retest 2 (Iteration 3)

Iteration 3 improved on the vehicle dynamics simulation by including ordinary differential equations for tire friction which places a maximum allowed speed. Furthermore the control loop of the controller places a priority on maximizing downforce in corners and minimizing drag in straights through an objective-driven loop instead of purely changing camber/angle of attack based on the curvature of the track. In the updated simulation, downforce has a greater role by applying a limit to the allowed speed for a vehicle. This means that speed cannot just increase infinitely as it was earlier and the vehicle should not just maintain or increase speed at curves especially when they get steep. In real life, tire friction requires the need of braking and so with it as an implementation, braking plays more of a role. The controller now has different optimizations such as maximizing downforce for long lateral sections (high curvature) or when curvature is really small, minimize drag by lowering angle of attack. Since this iteration focused on the simulation and controller aspects, the metrics related to performance time/accuracy of the PI-ROM were not repeated as those weights/parameters were saved and used. This includes the error heatmaps, drag polar curves, and thin airfoil curves in addition to

test MAPE or R2. These aspects were not improved upon from iteration 3 and are kept the same; iteration 3 only improved the vehicle dynamics and controller design which was evaluated with the simulation over the test track. Furthermore, since the simulation is now realistic and producing viable results, the fixed wing baseline was tested in this environment to compare for criteria. Its results are shown in Figure 19 while the active aero controller's results are shown in Figure 18.

Figure 18:
Simulation Wing Changes and Downforce/Drag Changes for Iteration 3 Dynamics/Controller

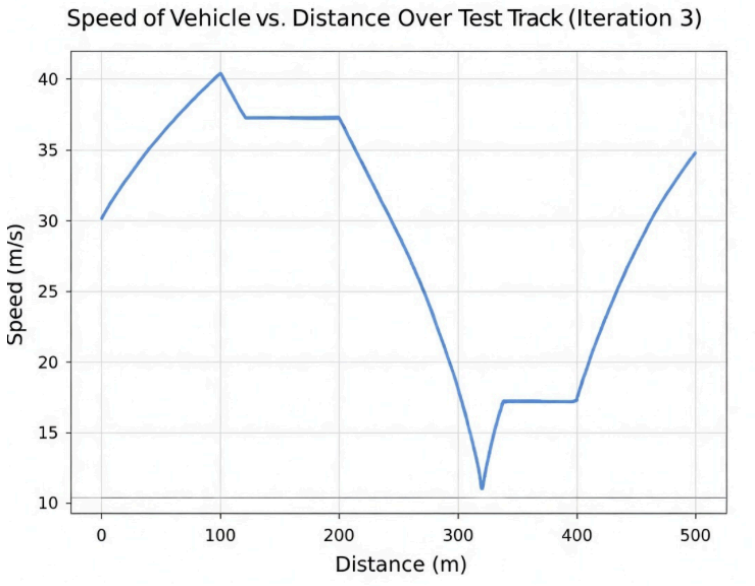

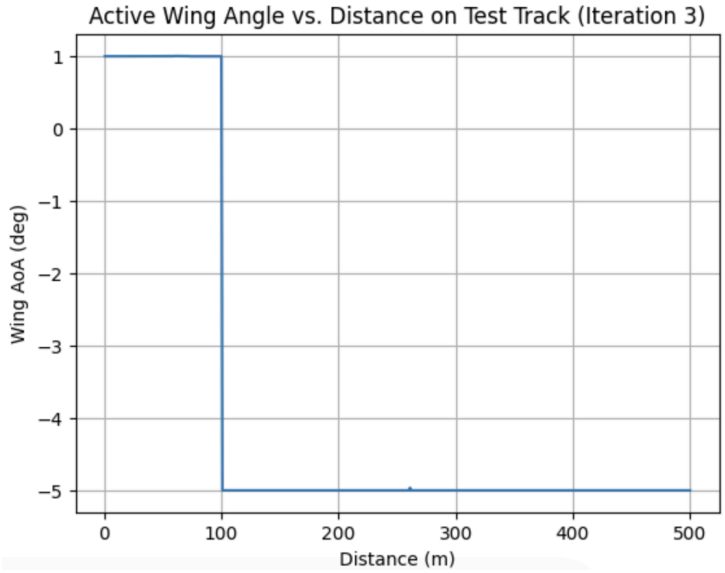
Active Wing Angle vs. Distance on Test Track (Iteration 3)
Wing AoA (deg)
1
0
−1
−2
−3
−4
−5
0
100
200
300
400
500
Distance (m)

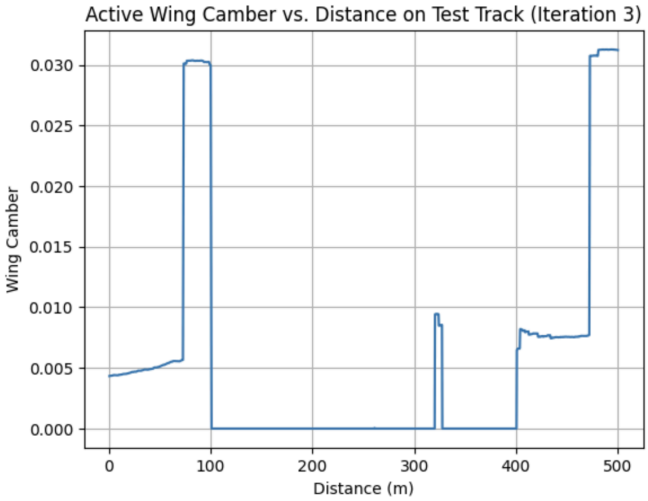
Active Wing Camber vs. Distance on Test Track (Iteration 3)
Wing Camber
0.030
0.025
0.020
0.015
0.010
0.005
0.000
0
100
200
300
400
500
Distance (m)

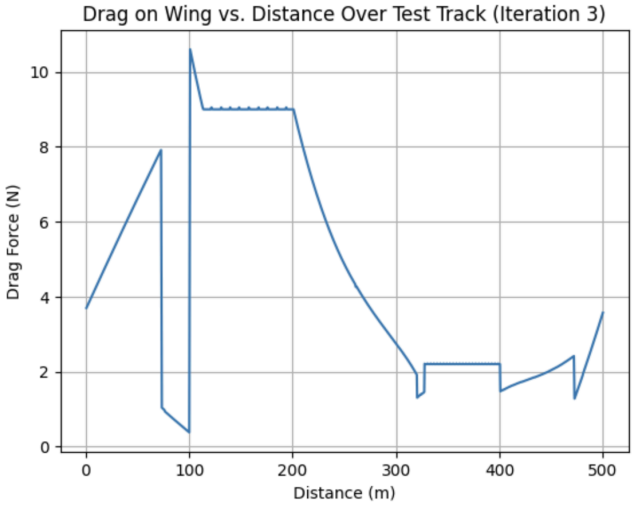
Drag on Wing vs. Distance Over Test Track (Iteration 3)
Drag Force (N)
10
8
6
4
2
0
0
100
200
300
400
500
Distance (m)

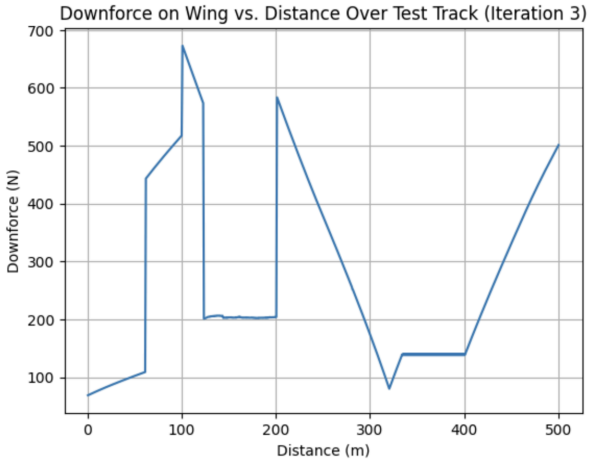

Figure 18 is the results of the third iteration vehicle dynamics simulation and shows how each feature the controller was changing as well as speed evolved over each meter on the track. It also includes downforce and drag which had to be recorded for the third iteration since they were crucial to the objective of the controller now.

The results in Figure 18 clearly demonstrate improvements in the overall realism of the simulation. The speed is shown to decrease often and a speed that changes like this is more realistic with the highly curved track and the addition of tire dynamics. Forces like drag and downforce (measured in Newtons [N]) are shown to be minimized or maximized with clear peaks and minimums to show that the wing's changes actually had effects. These peaks also correspond to when the track is straight or especially curved. For example, 100-200 meters is the most curved and thus downforce is the highest then at the peaks of 100 and 200 meters. Similarly, 400-500 meters is completely straight so drag has to be minimized by the wing there, which is shown in the plot. In previous iterations there was no logic to drive the wing to change these forces so they were essentially constant and their variance was not included as a metric of success for previous iterations. These values were not as dominating in other iterations and thus the wing's changes had no effect on them. It is also interesting to see that the controller placed

emphasis on camber changes while keeping the wing angle relatively small and constant, this could show the potential benefits of an active camber system.

Figure 19 shows the fixed wing's downforce, drag, and speed over time on the same track. To prove the necessity of an active aero wing that is based on the predictions of a reduced-order model, the simulation had to be run on the same track setup for a fixed wing. Now that the design is in its highest performance iteration, it could be compared with a fixed wing for true evaluation (Figure 19). This wing is of camber 0.02 and angle of attack of 15 degrees which is common in motorsports and well-equipped especially for a track this curved. Since the angle of attack is so high, it can be expected to have a much higher downforce production but also higher drag as well which is not beneficial.

For this wing, the drag and downforce are not clearly being maximized or minimized in certain sections of the track and are instead following a simple declining or increasing plot so even though the downforce may be higher due to the large angle of attack, this will not correlate well for lap time or energy saved. Additionally for the downforce plots, not only is the active wing trying to maximize it in the most cornered part of the track, the fixed wing's downforce is actually at its lowest during the corners (200-320 meters). It is not able to maximize and change the downforce production unlike the active wing which can maximize downforce to allow for greater safety. The drag plot is nearly identical to downforce plot and during the straights it is very high which makes it lose more energy and increases lap time.

The active wing was able to reduce the energy consumption by drag, showing it was better in drag management for performance, but also downforce management for stability and safety on curves. It struck a balance between downforce and drag where necessary since increasing one also increases the other which can be detrimental. This highlights a significant

difference between the intention of the active wing versus the normal, relative aerodynamic properties of a fixed wing that does not take advantage of the rear wing's geometry. The speed vs. distance graph shows similar results to the active aero wing and this goes to demonstrate the realism of the dynamic simulation since there are speed limits and it varies according to the curvature of the track like a real vehicle. The active wing should have no effect on the speed, only the aerodynamic forces, so it is valid that both speed plots are the same.

Figure 19:
Simulation Drag and Downforce Changes over Track for Fixed Wing

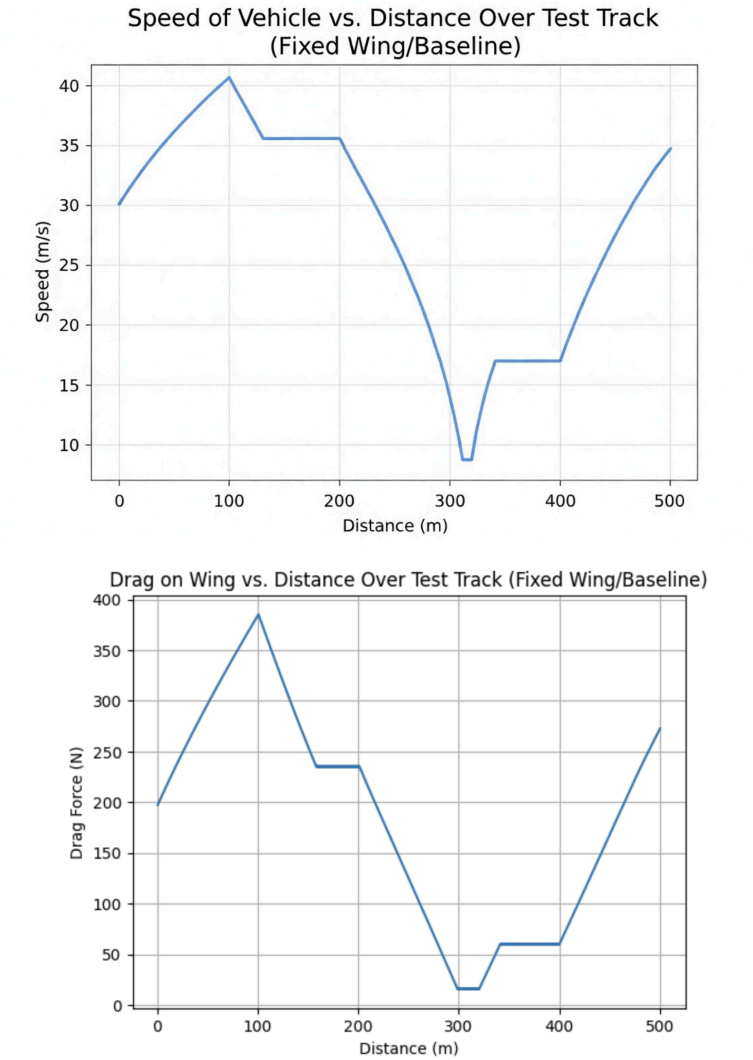

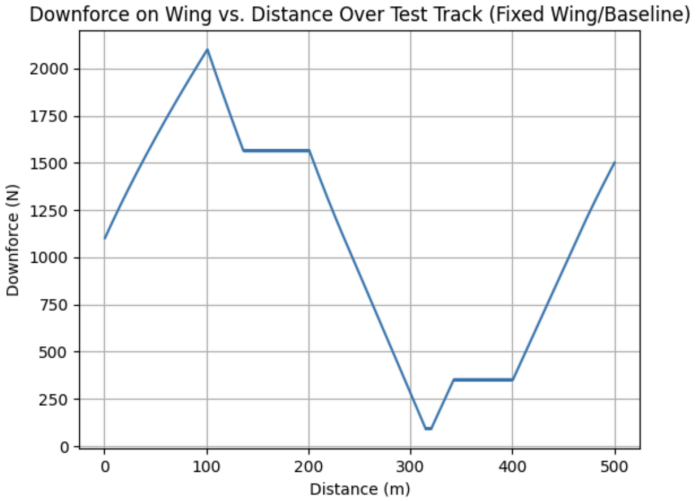

Figure 19 is the results of the fixed wing vehicle dynamics simulation and shows how each aerodynamic force was changing as well as speed over each meter on the track. It did not include camber or angle of attack since these are constant for a fixed wing.

Table 4:
Simulation Results for Lap Time and Energy Lost Across Tested Wing Types

Simulation Metrics	**Fixed**	**Iteration 1**	**Iteration 2**	**Iteration 3**
Lap Time (s)	19.44	10.84	10.68	18.24
Energy Lost (kJ)	7.77	161.90	20.65	2.16

Table 4 is a comparison for all iterations and the fixed wing baseline on the same track and for the metrics of energy lost and lap time.

Iterations 1 and 2 were run using the less realistic dynamics simulation which accounts for their unrealistic lap times. That low of a time is unfeasible with these track conditions and in order to accurately evaluate the controller, these times had to be feasible by revamping the vehicle dynamics in iteration 3. Even though its time is higher, it makes more sense and is thus an improvement in the design. Iterations 1 and 2 had similar results but iteration 2 had a drastically lower energy loss since drag prediction was improved which is the main reason

energy is lost throughout the track as it is doing negative work. Energy is primarily lost to drag on straight sections where speed can be maximized and since drag is proportional to speed, drag maximizes too. The lap time was higher by 1.2 seconds for iteration 3 compared to a fixed wing which is substantial in a motorsports environment. Since it has the capability to manage drag effectively, this allows speed to be the same but the drag force created to be less and minimized for improved performance and an immense reduction in energy consumption.The differences in lap time could be further improved upon in future simulations.

There was also less energy lost to drag with there being 5.61 kJ more saved compared to the fixed wing. Any energy saved is valuable and especially with this lap-time reduction, an active aero wing of this design has great potential. Overall in comparing each iteration there are improvements in terms of realism which is an important criteria. In comparison of the final iteration, iteration 3, with the fixed wing it performed better in terms of lap time reduction and energy consumption which aligns with the performance criteria initially set. The PI-ROM also succeeded in terms of the criteria of accuracy and time but also in maintaining the physics. Together, the iteration 2 PI-ROM which would have the saved weights and iteration 3 controller/simulation could have potential to alter the future of active aerodynamics not only to save drivers but also in future aerodynamic applications.

Conclusion

The design met the need of providing a quick and accurate replacement for predicting aerodynamic coefficients instead of constant aerodynamic solver calculations done for an active aero rear wing. An active aero controller framework was also paired with the physics-informed reduced-order model (PI-ROM) to demonstrate the full capabilities for real-time usage and

optimization. Following the necessary redesigns, the PI-ROM/active aero controller rear wing design was able to not only improve lap times, but was also efficient and accurate in its separate regimes.

The results back this understanding of the design since the PI-ROM, which was primarily based on the thin airfoil and drag polar equations but had a neural network residual, achieved a 26.58 times speedup factor compared to XFOIL predictions on the same NACA 2412 airfoil. This speedup is what allowed the controller and vehicle dynamic simulation to perform numerous, complex calculations in consecutive order to keep track of the various states. Further backing the validity of the PI-ROM is the predictive accuracy for the coefficients of lift and drag. The R2 scores were near 1.00 and positive and the MAE values were consistently low, serving as a benchmark for the design's potential. But most importantly, in graphing the drag polar relationship (CD vs. CL) and the thin airfoil relationship (AoA vs. CL), the PI-ROM maintained overlap with the XFOIl's predictions and the graphs followed a curve agreeing with the literature. This strongly indicates that the model is interpretable with respect to physics and is not purely phenomenological.

However, it is important to note the model's highest MAE near boundary points or when values like camber rise. This is not only when the wing starts to become less ideal so the physics relationships are not the most accurate, but also could be due to the model not having enough exposure to the boundaries because values there were not present in the sample set. Additionally, since the data was collected from XFOIL which only marks down converged cases, the model could have an incomplete view of all possible aerodynamic conditions since it was only exposed to those that converge properly introducing an aspect of bias to its predictions.

This paired with other reasons could have led to why drag was initially much more difficult to predict accurately compared to lift. While lift was being captured well by the physics baseline, drag often needed a strong negative residual from the neural network which pushed the model into predicting negative numbers. This was redesigned to have the coefficient of lift predicted by the physics baseline be an input for the neural network residual in its prediction of drag since lift and drag are very interconnected especially in the drag polar relationship. As far as the simulation, the controller initially had trouble making proper changes to optimize or minimize the aerodynamic forces of downforce and drag as it should and was mostly based off track curvature. This was modified to ensure the controller prioritized changes to the aerodynamic forces by changing the camber or angle of attack accordingly rather than just changing each wing configuration in unison based on how curved the track was.

It is more common and available for wings to change their angle of attack but one that changes camber in real-time could positively impact the future of active aerodynamics. The design's controller largely changed camber to alter the aerodynamic forces for optimization while leaving angle relatively untouched. An active aero rear wing that manipulates camber could have great potential for future applications and would require more research into its development.

This design can be expanded upon in the future through greater inclusion of boundaries possibly with softwares like OpenFOAM in order to prevent converging issues. While physical constraints to ensure drag did not go below zero were added, this could be changed to make sure the physics baseline predictions are more accurate and try to prevent negative drag prediction in a different manner. The results of this study are deterministic and immediate future work can include stochastic disturbances. The vehicle dynamic simulation also assumed simple

one-dimensional differential equations so this could be updated to include more complex equations that are more closely aligned with the real world. This would need more complex computer systems, however, to keep track of more states. Finally, data for a 3D rear wing could be collected and used to train the model and controller. This would be more beneficial for the real world but would need more expensive equipment and experimentation.

In terms of real world application, this project is the baseline for active aerodynamic work that requires quick aerodynamic property prediction especially for the rear wing. The ROM could replace expensive aerodynamic evaluations during real motorsports and so there could be real-time optimization on the track. The PI-ROM method is crucial in many other areas as well such as digital simulations or real-time control of any physics systems such as military equipment. The design sets up a plausible framework for active aero in motorsports with a controller that relies on an accurate and quick physics-informed reduced-order model with a neural network residual.

Reference List

Andrés-Pérez, E., & Paulete-Periáñez, C. (2021). On the application of surrogate regression models for aerodynamic coefficient prediction. *Complex & Intelligent Systems*. https://doi.org/10.1007/s40747-021-00307-y

Cadence CFD Solutions. (2023). *An Overview of Thin Airfoil Theory*. Resources.system-Analysis.cadence.com. https://resources.system-analysis.cadence.com/blog/2023-an-overview-of-thin-airfoil-theory

Cakir, M. (2012). *Scholar Commons CFD study on aerodynamic effects of a rear wing/ spoiler on a passenger vehicle*.

https://scholarcommons.scu.edu/cgi/viewcontent.cgi?params=/context/mech_mstr/article/1000/&path_info=CFD_study_on_aerodynamic_effects.pdf

Daily, E. (2025, June 23). *How CFD is killing Wind Tunnels*. Medium. https://medium.com/@Engineering_Daily/how-cfd-is-killing-wind-tunnels-2335b88ae551

de Buck, P., & Martins, J. R. R. A. (2022). Minimum lap time trajectory optimisation of performance vehicles with four-wheel drive and active aerodynamic control. *Vehicle System Dynamics*, 1–17. https://doi.org/10.1080/00423114.2022.2101930

De Marco, A. (2020, May 24). *Aircraft Drag Polar*. Flight Mechanics for Pilots. https://agodemar.github.io/FlightMechanics4Pilots/mypages/drag-polar/

Dean, H., Decker, K., Robertson, B., & Mavris, D. (2025). Reduced-Order Modeling Approach to Blunt-Body Aerodynamic Modeling | Journal of Spacecraft and Rockets. *Journal of Spacecraft and Rockets*. https://doi.org/10.2514/jsr.2025.62.issue-4;article:article:10.2514/1.A36095;page:string:Article

Georgiev, P., De Filippis, G., Gruber, P., & Sorniotti, A. (2023). On the Benefits of Active Aerodynamics on Energy Recuperation in Hybrid and Fully Electric Vehicles. *Energies*, *16*(15), 5843. https://doi.org/10.3390/en16155843

Ghoreyshi, M., Jirásek, A., & Cummings, R. M. (2014). Reduced order unsteady aerodynamic modeling for stability and control analysis using computational fluid dynamics. *Progress in Aerospace Sciences*, *71*, 167–217. https://doi.org/10.1016/j.paerosci.2014.09.001

Jiménez Elbal, A., Zarzuelo Conde, A., & Siampis, E. (2024). Simultaneous Optimisation of Vehicle Design and Control for Improving Vehicle Performance and Energy Efficiency

Using an Open Source Minimum Lap Time Simulation Framework. *World Electric Vehicle Journal*, *15*(8), 366. https://doi.org/10.3390/wevj15080366

Li, K., Kou, J., & Zhang, W. (2021). Unsteady aerodynamic reduced-order modeling based on machine learning across multiple airfoils. *Aerospace Science and Technology*, *119*, 107173. https://doi.org/10.1016/j.ast.2021.107173

Massaro, M., & Limebeer, D. J. N. (2021). Minimum-lap-time optimisation and simulation. *Vehicle System Dynamics*, *59*(7), 1069–1113. https://doi.org/10.1080/00423114.2021.1910718

Motorsport. (2023, March 1). *What is DRS in F1, how does it work and is it automatic?* Www.motorsport.com. https://www.motorsport.com/f1/news/what-is-drs-in-f1-how-does-it-work-is-it-automatic/10437677/

NASA Glenn Research Center. (2024, July 10). *Lift Equation | Glenn Research Center | NASA*. Glenn Research Center | NASA. https://www1.grc.nasa.gov/beginners-guide-to-aeronautics/lift-equation/

Noble, J. (2024, June 6). *How F1's new active aero will work in 2026 as DRS is dropped*. Motorsport.com. https://www.motorsport.com/f1/news/how-f1s-new-active-aero-will-work-in-2026/10620106/

Occamsracers, A. (2023, February 11). *Active Aero Wings and Spoilers*. Occam's Racer. https://occamsracers.com/2023/02/11/active-aero-on-a-miata/

Rijns, S., Teschner, T. R., Blackburn, K., Siampis, E., & Brighton, J. (2025). Optimising vehicle performance with advanced active aerodynamic systems. *Vehicle System Dynamics*, 1–29. https://doi.org/10.1080/00423114.2025.2505619

Renn, P. I., & Gharib, M. (2022). Machine learning for flow-informed aerodynamic control in turbulent wind conditions. *Communications Engineering*, *1*(1), 1–9. https://doi.org/10.1038/s44172-022-00046-z

Zhang, T., Sun, Y., Wang, Y., Li, B., Tian, Y., & Wang, F.-Y. (2024). A Survey of Vehicle Dynamics Modeling Methods for Autonomous Racing: Theoretical Models, Physical/Virtual Platforms, and Perspectives. *IEEE Transactions on Intelligent Vehicles*, 1–24. https://doi.org/10.1109/tiv.2024.3351131