AI-Controlled Deployable Vortex Generators for Adaptive Airflow Optimization

Ishan Kasam

Mass Academy of Math and Science

Advanced STEM with Scientific and Technical Writing

Instructor: Kevin Crowthers, Ph.D.

Worcester, MA. 01605

Abstract

Parasitic drag from vortex generators (VGs) imposes persistent efficiency penalties during aircraft cruise, particularly for UAVs and small aircraft where fixed flow control devices operate outside their useful envelope for most missions. While VGs effectively delay boundary layer separation, their fixed height and permanent deployment cause unnecessary drag outside critical maneuvers and high-angle-of-attack situations. Current research gaps include non-optimized operational VG heights and deployable systems lacking truly adaptive control. This work designed and evaluated a deployable VG system capable of actively minimizing drag while preserving separation control when needed. A mechanically actuated VG array was developed using servo actuators and a lightweight machine learning algorithm for microcontroller implementation. The controller combines Gaussian Process regression with a neural network to capture uncertainty in separation onset and nonlinear aerodynamic response. Training data were generated through a multi-fidelity approach, pairing medium-fidelity RANS simulations with high-fidelity URANS simulations via Co-Kriging. Without fully integrated wind tunnel testing, system performance was assessed using a highly realistic flight simulation. Results show adaptive height control improved average lift-to-drag ratio (Cl/Cd) by 6.5% and saved 2.55% fuel compared to standard fixed VG airfoils. These findings demonstrate that VG height optimization, beyond simple deployment or retraction, is critical for reducing parasitic drag. The system is directly applicable to UAVs and light aircraft operating at moderate Reynolds numbers, where adaptive flow control can yield measurable efficiency gains without compromising stability.

Keywords: active flow control, vortex generators, Gaussian regression, machine learning, CFD

Acknowledgements

I would like to thank my parents for supporting me financially in this project. I would like to thank Dr. Crowthers for giving me advice on research methods and how to effectively support a narrative. I would like to thank Professor Aswin Gnanaskandan for unofficially mentoring me and giving me advice on CFD methods.

AI-Controlled Deployable Vortex Generators for Adaptive Airflow Optimization

Aircraft fuel consumption accounts for approximately 2.5% of global CO2 emissions (Ritchie, 2024), with aerodynamic drag being a primary contributor. This project develops deployable vortex generators with machine learning-based control to reduce drag during cruise while maintaining stall-prevention benefits during takeoff and landing.

Boundary Layer Separation

The boundary layer is the thin region adjacent to a surface where velocity transitions from zero at the wall to free-stream values (Leishman, 2023). Under adverse pressure gradients, flow can separate from the wall, reducing lift and increasing drag, with extensive separation causing stall that severely degrades aerodynamic performance (Leishman, 2023). Differences in levels of separation can develop with parameters such as Reynolds number, angle of attack, and duration of time in certain conditions.

Vortex Generator Geometry Fundamentals

Vortex generators are small fin-shaped devices that create vortices to energize the boundary layer and delay separation. Lin (2002) demonstrated that optimal VG performance occurs at height-to-boundary-layer-thickness ratios between 0.1 and 0.5. Jayanarasimhan & Balasubramanian (2025) found that triangular vortex generators achieved 48.7% lift coefficient improvements at 18° angle of attack with height-to-thickness ratios between 0.8 and 1.5, demonstrating sensitivity to geometric parameters that motivates adaptive height optimization.

Static Vortex Generator Limitations

Static vortex generators introduce continuous parasitic drag throughout flight. Lin (2002) noted that VGs sized for takeoff and landing effectiveness reduce cruise efficiency by protruding into attached boundary layers, but their effect has not yet been quantified. Titchener & Babinsky (2015) emphasized

that device drag penalties are rarely quantified despite their significance for aircraft efficiency, highlighting a critical knowledge gap in cost-benefit trade-offs. Ultimately, vortex generators cannot be omitted from aircraft because of their necessity as a safety feature, but they still cause parasitic drag.

Previous Deployable Systems and Their Research Gaps

Le Pape et al. (2012) designed mechanically actuated deployable vortex generators with continuous height control from 0.1 to 3 mm using hydraulic linear motors, successfully preventing dynamic stall but introducing drag penalties at low angles of attack when deployed. Mamman et al. (2023) developed an alternative using twisted spiral artificial muscles (TSAMs) consuming 0.5-2 W per unit versus 5-10 W for mechanical systems, but relied on lookup tables rather than actuating across a continuum of values, and also actuated for the purpose of correcting the temperature difference. Shimomura et al. (2020) demonstrated that deep reinforcement learning could achieve 3-4 times longer flow attachment with plasma actuators compared to fixed-frequency control, but focused only on delaying stall rather than optimizing cruise efficiency. These systems demonstrate feasibility of precise actuation and adaptive control independently, but lack integration of real-time optimization with deployable vortex generators that balance stall prevention and cruise drag reduction across all flight regimes.

Design Requirements and Competitive Advantages

A key criterion for deployable vortex generator systems is that they must achieve sub-millimeter deployment accuracy comparable to the 0.5 mm precision demonstrated by Le Pape et al. (2012) while also running light enough to have low weight. The system must provide actuation response compatible with boundary layer time scales to prevent separation onset, requiring control latency under 100 milliseconds for typical approach and landing conditions.

Multi-Fidelity CFD for Training Data

Three-dimensional CFD is required to capture vortex generator effects. He et al. (2020) demonstrated that deep neural networks could learn correlations between low-fidelity Euler simulations and high-fidelity Navier-Stokes data, requiring only a small number of expensive simulations. Fuchi et al. (2022) achieved 22-58% convergence speedup using machine-learning-predicted flow fields as initial conditions, requiring approximately 1000 training points from high-fidelity simulations. This multi-fidelity approach enables generation of large training datasets capturing the complex relationship between vortex generator height and separation behavior while maintaining computational feasibility for controller development.

Machine Learning for Real-Time Control

Portal-Porras et al. (2022) showed that convolutional neural networks could predict airfoil flow fields with 0.827% lift coefficient error while running approximately 10,000 times faster than conventional CFD, training on 158 Reynolds-Averaged Navier-Stokes simulations. Zhang et al. (2025) reconstructed full surface pressure fields from only 10 optimal sensor locations with 2.43% mean relative error for lift coefficient prediction. These approaches demonstrate that machine learning can provide rapid flow state estimation necessary for closed-loop vortex generator control. While making high fidelity predictions from deep learning, fast predictions may not be possible, but the values can be used to create a cheaper algorithm. The combination of sparse pressure sensing and neural network prediction enables practical implementation of model-based control with latency compatible with boundary layer dynamics.

Problem Statements:

1. Static vortex generators create parasitic drag during cruise conditions and are not optimized for all given sets of conditions.
2. Current active flow control systems for airfoils lack robust actuation techniques, feasible implementation, and optimized, adaptive control.

Objective:

Obj. 1: Develop a machine learning model to predict the optimal height of vortex generators given discrete inputs of pressure at chordwise points.

Obj. 2: Develop a lightweight control algorithm to deploy on a microcontroller paired with sensor fusion and an actuator to physically deploy vortex generators.

Section II: Methodology

This project had no mentor, and the entirety of the work was completed by Ishan Kasam.

Equipment and Materials

The simulation phase uses SolidWorks for modeling, ANSYS Fluent for generating meshes and CFD simulations, and Python libraries including Pandas, SkLearn, and PyTorch for data processing and machine learning, and airfoil geometries were generated from Airfoil Tools and the UIUC database.

Medium Fidelity Reynolds-Averaged Navier–Stokes Simulations

Computational Domain and Geometry

Training data for the control model was generated using Reynolds-Averaged Navier–Stokes simulations. This approach qualifies as medium fidelity because steady-state RANS turbulence simulations with meshes of ~1.8 million cells balance accuracy and short run times for the machine learning model. The geometry of a NACA 4415 airfoil was downloaded from the UIUC airfoil coordinates database and then scaled to a 250 mm chord to achieve Reynolds numbers of 0.8×10^6 to 2.0×10^6 at realistic freestream velocities of 50–120 m/s in air. Rectangular counter-rotating vortex generators were modeled in SolidWorks by replicating device

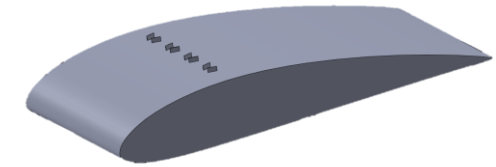

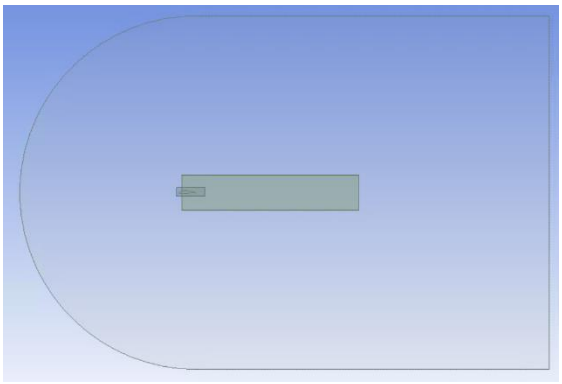

dimensions from Reuss et al. (1995), and scaling them proportionally for the 250 mm chord. The external flow domain was constructed using a C-shaped topology to provide sufficient downstream space for wake development during separated-flow conditions. The inlet boundary was positioned 1.25 m upstream of the leading edge, the outlet 5 m downstream of the trailing edge, and far-field boundaries 2.5 m above and below the chord line to minimize blockage (the effective blockage less than 1%) and reduce boundary-induced errors in lift and drag prediction. Boundary conditions were defined as velocity inlet upstream, pressure outlet downstream, no-slip walls on the airfoil and vortex generators, and symmetry conditions at far-field boundaries to maintain uniform flow properties at domain extents.

Mesh Generation and Refinement

Targeted mesh refinement was applied to capture vortex formation and wake recovery while keeping the cell count low enough to have fast solver times. A bounding box for refinement was placed 0.1c upstream, 0.5c downstream, and ±0.1c vertical to the airfoil to capture vortex generator-induced shear layers and near-wall gradients. A second refinement region that was larger extended from upstream of the vortex generators to 9c downstream and ±1c normal chord to capture wake development and improve the predictions of lift and drag.

The mesh was generated in ANSYS Fluent using a poly-hexcore topology, meaning a predominantly hexahedral mesh with polyhedral cells in regions that transition to the inflation layers. 28 inflation layers and a first layer height of 1.5e-6 meters were chosen to achieve $y^+ < 1$ to resolve the viscous sublayer and near-wall boundary layer structure (Leishman, 2023), and accurate resolution near the wall is required because vortex generators directly impact wall shear stress and separation onset (Lin, 2002). Complete settings and parameters for the mesh are provided in Appendix X. The resulting meshes contained approximately 1.84 million cells for vortex-generator configurations and 1.60 million cells for the clean airfoil, and both had

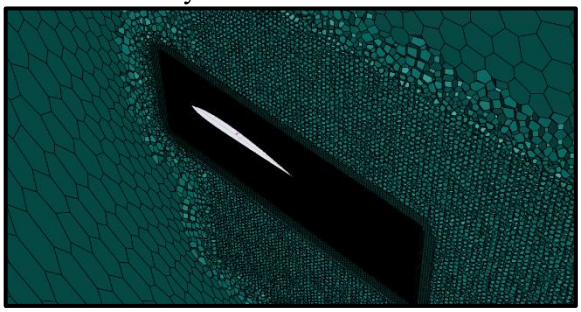

maximum skewness values of 0.151 (irrelevant as very few cells had this), but average skewness of ~0.97.

Simulation Setup and Execution

The CFD simulations were solved in ANSYS Fluent using the k–ω SST turbulence model because it captures interactions near the wall better than k-ε but remains efficient in the free stream like k-ε. Simulations were run to residual convergence with a criterion of 10^{-5} because this would allow the simulations to stop only once non-fluctuating values of lift and drag were reached. Turbulence intensity at inlet and outlet boundaries was specified as 0.1% to represent low-turbulence freestream conditions that would be seen in real life.

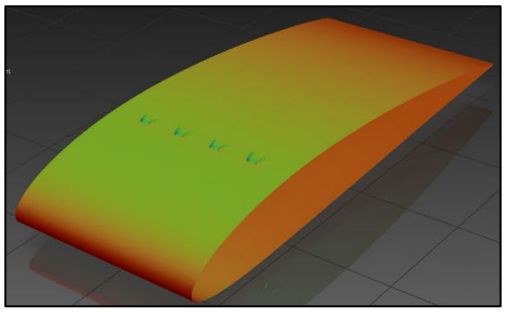

An automated script executed 154 steady RANS simulations spanning 22combinations of Reynolds-numbers and angle-of-attack and seven vortex-generator heights (0–3 mm). Reynolds numbers from 0.8×10^6 to 2.0×10^6 and angles of attack from 0° to 20° were selected using Latin Hypercube sampling to ensure space-filling stratified sampling across each parameter dimension and avoid clustering (Fuchi et al., 2022). In the simulations, Reynolds number and angle of attack were manipulated by varying the velocity magnitude and direction in the solver. The full simulation matrix is provided in Appendix X. For each unique case, lift coefficient, drag coefficient, and pressure coefficient at 20 chordwise locations were recorded.

Validation

The CFD simulations were validated against experimental data from Reuss et al. (1995) for the 2.71 mm vortex generator configuration at angles of attack of 12°, 18°, and 20°. The predicted lift and drag coefficients showed errors of approximately 5–10% compared to wind tunnel measurements, which was considered acceptable for generating training data, and the validated mesh was then changed to

simulate other vortex generator heights (0–3 mm) for which experimental validation data was not available.

High Fidelity URANS Simulations

Geometry, Domain, and Mesh Generation

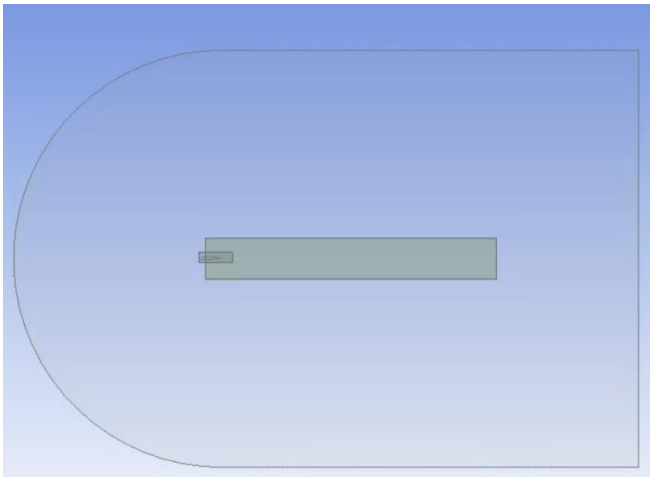

The geometries for the airfoils used for the high-fidelity simulations were the same as those used for the medium fidelity simulations. The bounding C-shaped domain also remained the same, but this time, the mesh refinement boxes changed, with the box for the direct wake of the vortex generators starting directly upstream of the VGs and extending 0.1c downstream of the trailing edge. The larger wake bounding box started at the middle of the chord, extending downstream 15c, and ±1c vertical to the chord line of the airfoil.

The mesh was made with ANSYS Fluent and included much finer face sizes for the airfoil and VG walls, as well as finer element sizes for the refinement boxes for VG wake and the far wake. The meshes had cell counts of around ~10.3 million cells.

Solver

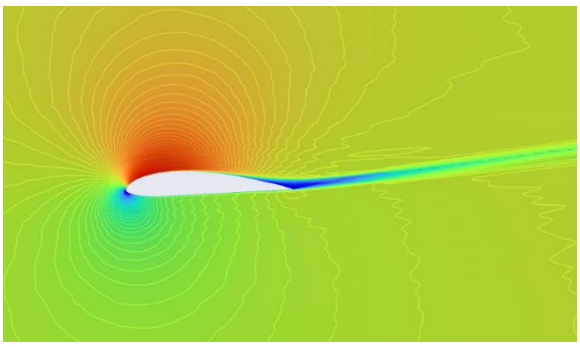

10 cases from the lower fidelity simulations were chosen to be ran for the high-fidelity simulations. The solver settings remained nearly identical to the medium-fidelity simulations, but the model was now URANS k-omega SST. First a steady RANS case was converged as an initialization,

and then the URANS was used and converged. The time step was 0.0004s and 2000 time steps were used, but convergence usually occurred at around 1700 time steps.

Gaussian Process Regression Training and Usage

Surrogate Modeling Strategy and Rationale

A surrogate aerodynamic model was developed to replace expensive CFD evaluations during control and mission simulation. Real-time control requires millisecond-level predictions, while CFD solutions require minutes to hours per case. Machine-learning surrogates trained on CFD data have been shown to reproduce aerodynamic coefficients with high accuracy while providing orders-of-magnitude speedup, making them suitable for real-time aerodynamic prediction and control applications (Negoita & Hothazie, 2024; Portal-Porras et al., 2022).

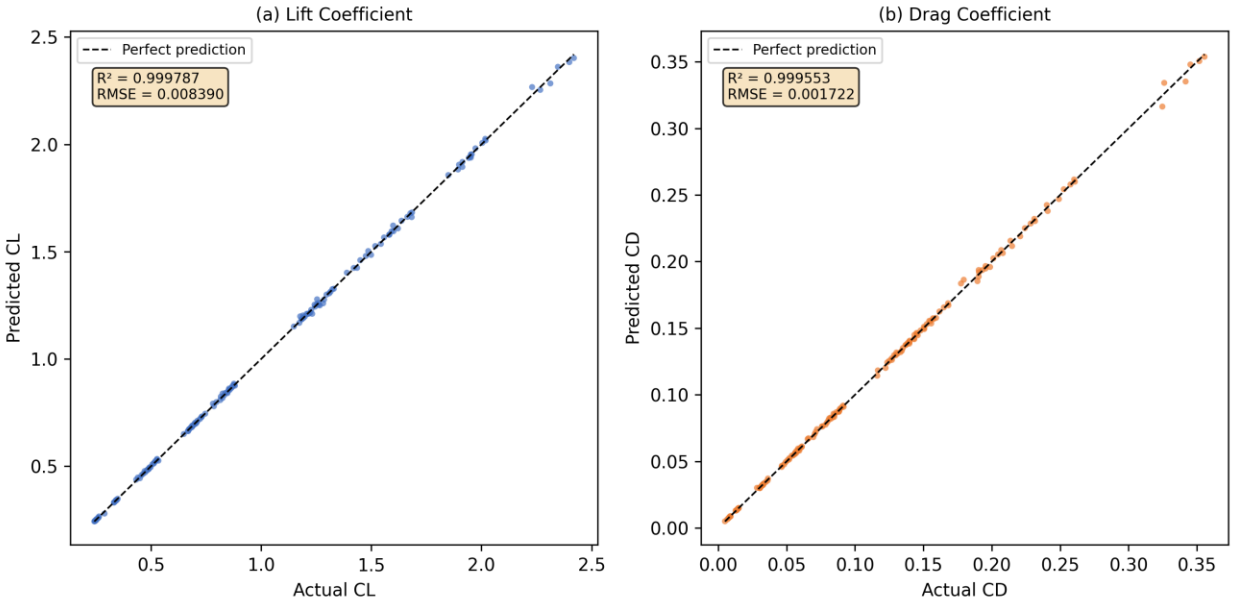

Gaussian Process Regression (GPR) was selected as the primary surrogate modeling technique because it performs well with limited datasets and produces smooth interpolation across nonlinear aerodynamic parameter spaces. The available dataset consisted of CFD simulations parameterized by angle of attack (α), Reynolds number (Re), and vortex generator height (h), which are the primary variables governing boundary-layer behavior and separation onset (Leishman, 2023). Vortex-generator

height was treated as a continuous variable because device effectiveness is highly sensitive to geometric scale and boundary-layer thickness ratios (Lin, 2002).

Training Data Preparation and Feature Processing

The training dataset consisted of medium-fidelity CFD simulations spanning Reynolds numbers representative of small aircraft and UAV flight. Reynolds number was rescaled into units of millions prior to normalization to prevent it from dominating kernel distance calculations. All inputs were standardized using z-score normalization so that Euclidean distance in feature space represented comparable physical variation across α, Re, and h.

This preprocessing step was required because distance-based kernels are sensitive to feature magnitude. Without scaling, Reynolds number would dominate similarity calculations and reduce sensitivity to VG height and angle of attack. Dataset structure and preprocessing statistics are provided in Appendix A (Dataset and preprocessing).

Gaussian Process Model Architecture

Separate Lift and Drag Models

Two independent GPR models were trained: one predicting lift coefficient (CL) and one predicting drag coefficient (CD). Separate models were used because lift and drag respond differently to vortex generator deployment and have different variance structures. This prevents drag predictions from being biased by lift-dominated behavior.

Kernel Selection and Hyperparameters

A Matérn kernel with smoothness parameter $\nu = 1.5$ was used. Aerodynamic responses vary smoothly across operating conditions but exhibit sharper gradients near stall and separation onset. The Matérn kernel captures this behavior more realistically than infinitely smooth kernels. Hyperparameters were optimized using multiple restarts to reduce sensitivity to local minima. Full hyperparameter values are listed in Appendix B (GPR configuration).

Regularization and Numerical Stability

A small ridge-type noise term was introduced to stabilize matrix inversion and prevent overfitting to CFD noise. Without regularization, clustered training points can lead to ill-conditioned covariance matrices. Sensitivity analysis is provided in Appendix C (Hyperparameter sensitivity).

Multi-Fidelity Residual Learning Framework

To improve prediction accuracy near critical flow regimes, a multi-fidelity framework was implemented. A secondary GPR model was trained on the residual between high-fidelity and medium-fidelity simulations. This residual-learning approach allows a small number of high-fidelity simulations to improve predictions while limiting computational cost, consistent with multi-fidelity strategies demonstrated for aerodynamic flows (Fuchi et al., 2022).

A spatial gating rule was implemented so the residual correction is applied only near high-fidelity training points. This prevents sparse high-fidelity data from introducing global bias. Details of gating threshold selection are provided in Appendix D (Multi-fidelity gating).

Control Algorithm and Height Optimization Logic

The surrogate model was embedded into a control algorithm that selects vortex-generator height in real time. Candidate heights spanning the actuator range were evaluated at each state, and the height that maximized aerodynamic efficiency while maintaining required lift was selected. This approach was motivated by research showing strong sensitivity of vortex-generator performance to geometric parameters and boundary-layer thickness ratios (Lin, 2002; Jayanarasimhan & Balasubramanian, 2025).

The optimizer enforced a stall-margin constraint to prevent operation near separation. Objective weights, regime logic, and pseudocode are provided in Appendix F (Control algorithm).

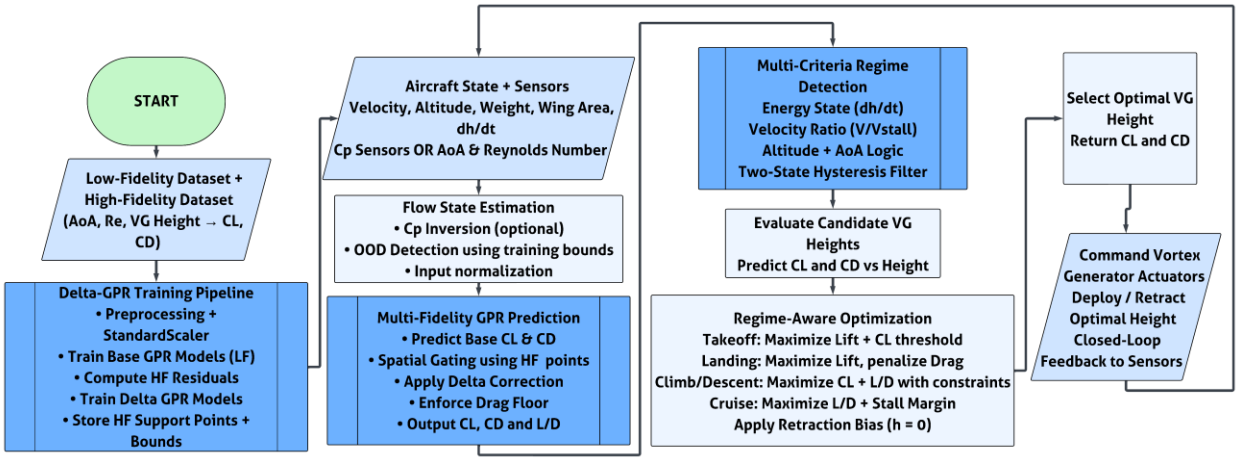

Mission Flight Simulation and Performance Evaluation

Flight Path Generation and Time Discretization

The simulation generates a piecewise flight path composed of sequential flight phases. Each phase defines altitude, velocity, and duration, and the trajectory is discretized into time steps. Smaller time steps were used during takeoff and landing to capture rapid changes, while larger steps were used during cruise to reduce computational cost. Time-step validation is documented in Appendix H (Time-step convergence).

Atmospheric Modeling and Reynolds Number Calculation

Atmospheric properties were updated at each time step to compute Reynolds number and dynamic pressure. This ensures aerodynamic predictions remain consistent with the surrogate training variables. The full atmospheric model is provided in Appendix I (Atmospheric model).

Conversion of Aerodynamic Coefficients to Fuel Consumption

At each time step, the surrogate model predicts CL and CD. Drag is then computed from dynamic pressure and wing area, linking aerodynamic efficiency to required propulsive power. Required power is calculated from drag and velocity and converted to fuel flow using propulsion efficiency and brake-specific fuel consumption. This approach links aerodynamic efficiency directly to fuel burn. Model parameters are listed in Appendix J (Fuel model). Fuel mass is updated at each time step using forward Euler integration. Aircraft weight therefore decreases during the mission, which alters lift requirements

and closes the feedback loop between aerodynamics and fuel consumption. Numerical convergence testing is provided in Appendix K (Integration validation).

Configuration Comparison and Drag Penalty Modeling

Three configurations were simulated under identical mission conditions: clean airfoil, fixed-height vortex generators, and deployable adaptive vortex generators. A constant drag penalty was applied to represent actuator and mechanism drag not captured in CFD. Deployable vortex-generator systems have been shown to introduce drag penalties when deployed, which must be included to obtain realistic mission-level results (Le Pape et al., 2012). Comparison methodology and sensitivity analysis are provided in Appendix L (Configuration comparison).

Statistical Tests

Repeated Measures ANOVA (Flight Simulation L/D)

One-way ANOVA comparing lift-to-drag ratio across all eight configurations in the 629-point flight simulation found significant differences ($F = 44.231$, $p = 4.50 \times 10^{-61}$), with mean values ranging from 14.21 (Fixed 1.2%) to 16.56 (Fixed 0.0%).

Repeated Measures ANOVA (Flight Simulation Fuel Rate)

One-way ANOVA comparing fuel consumption rate across all eight configurations found no significant differences ($F = 1.559$, $p = 0.143$), with mean fuel rates ranging from 0.002050 kg/s (Fixed 0.0%) to 0.002285 kg/s (Fixed 1.2%).

Paired t-Tests (Deployable vs Fixed Configurations)

Paired t-tests compared the deployable system (mean L/D = 15.802) against each fixed configuration at 629 matched time points: vs Fixed 0.0% ($t = -50.059$, $p = 2.20 \times 10^{-221}$, $d = -1.996$), vs Fixed 0.3% ($t = 25.953$, $p = 1.83 \times 10^{-101}$, $d = 1.035$), vs Fixed 0.5% ($t = 36.164$, $p = 1.18 \times 10^{-155}$, $d = 1.442$), vs Fixed 0.7% ($t = 40.960$, $p = 1.62 \times 10^{-179}$, $d = 1.633$), vs Fixed 0.9% ($t = 43.419$, $p = 2.83 \times$

10^{-191}, d = 1.731), vs Fixed 1.2% (t = 43.874, p = 2.05×10^{-193}, d = 1.749), and vs Fixed 1.5% (t = 39.870, p = 3.34×10^{-174}, d = 1.590), with all comparisons except Fixed 0.0% showing significant improvements ($p < 0.001$).

Wilcoxon Signed-Rank Test

The Wilcoxon signed-rank test comparing deployable and fixed 3.0 mm configurations found the deployable system produced significantly higher lift-to-drag ratios (W = 66.0, p = 0.017*).

Levene's Test for Homogeneity of Variance

Levene's test confirmed equal variance across vortex generator height groups (W = 0.231, p = 0.966). This result satisfied the homoscedasticity assumption required for valid ANOVA inference.

Section III: Results

CFD Aerodynamic Dataset Results

The CFD dataset spanning angle of attack (0–20°), Reynolds number (0.85×10^6–1.98×10^6), and vortex generator height (0–3 mm) comprised 151 simulations. Maximum lift coefficient ranged from 2.02 to 2.42, occurring at angles of attack between 16.1° and 19.0°. Minimum drag coefficient ranged from 0.00496 to 0.00911, and peak lift-to-drag (coefficients) ratio ranged from 28.66 to 58.35 across all configurations.

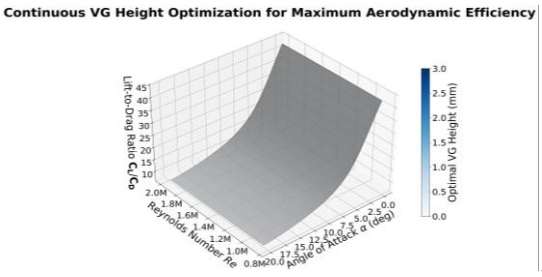

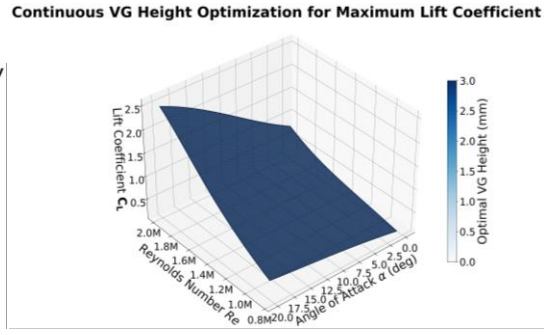

Surrogate Model Validation Results

Cross-validation accuracy metrics for the Gaussian Process Regression models are summarized in Table 2. For lift coefficient prediction, 10-fold cross-validation yielded RMSE = 0.00839, MAE = 0.00550, and R^2 = 0.99979. Leave-one-out cross-validation produced RMSE = 0.00791, MAE = 0.00509, and R^2 = 0.99981. For drag coefficient prediction, 10-fold cross-validation yielded RMSE = 0.00172, MAE = 0.00105, and R^2 = 0.99955. Leave-one-out cross-validation produced RMSE = 0.00151, MAE = 0.00088, and R^2 = 0.99966.

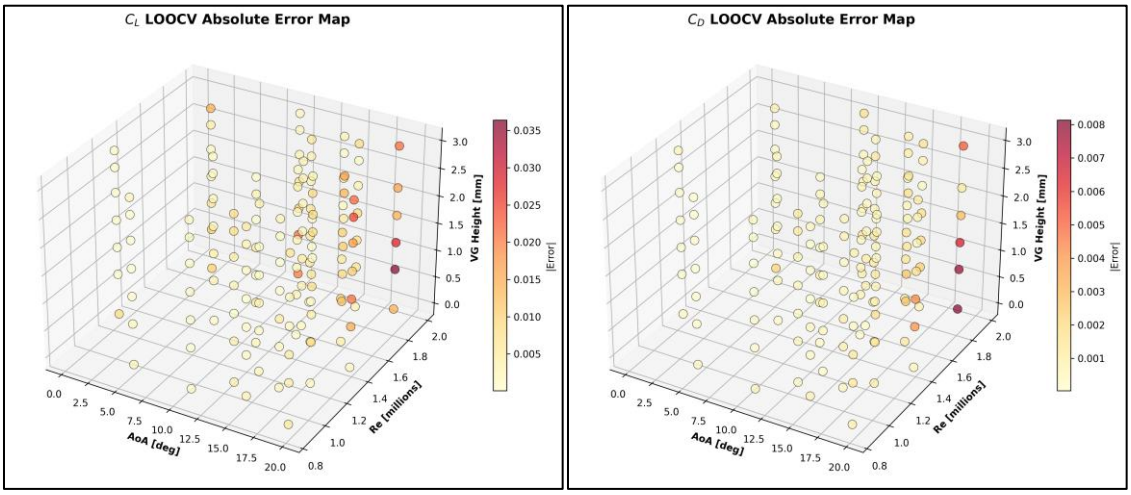

VG Deployment Behavior in Flight Simulation

The control algorithm selected vortex generator heights from 0.00 mm to 3.00 mm across the 629-time-step mission. The deployable system was active for 2.1% of mission duration with a deployed height of 3.00 mm when active. The vortex generators remained fully stowed for 97.9% of the mission, primarily during cruise where clean configuration maximized efficiency.

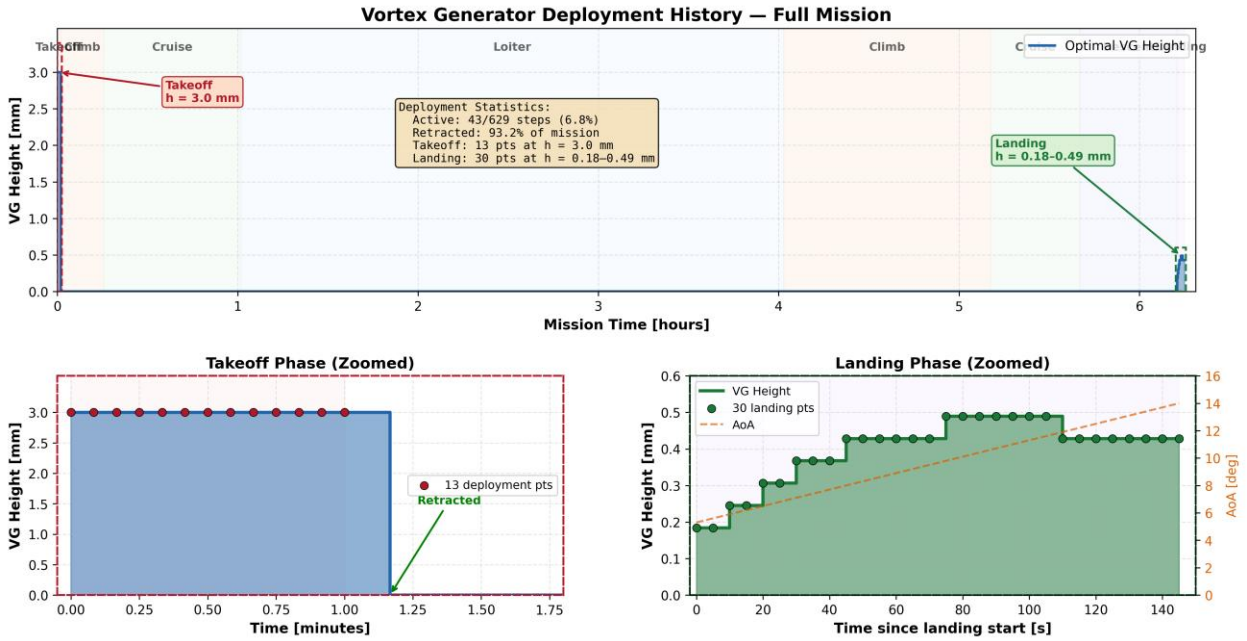

Flight Simulation Results

Mission-averaged lift-to-drag ratio ranged from 14.21 (Fixed 1.2%) to 16.56 (Fixed 0.0%) across configurations. The deployable configuration achieved 15.80, outperforming all fixed vortex generator configurations. Total fuel consumption ranged from 43.04 kg (Fixed 0.3%) to 47.69 kg (Fixed 1.2%). The deployable configuration consumed 44.27 kg, achieving 2.66% fuel savings relative to Fixed 0.7% ($p = 3.28 \times 10^{-72}$***) and 7.18% savings relative to Fixed 1.2%.

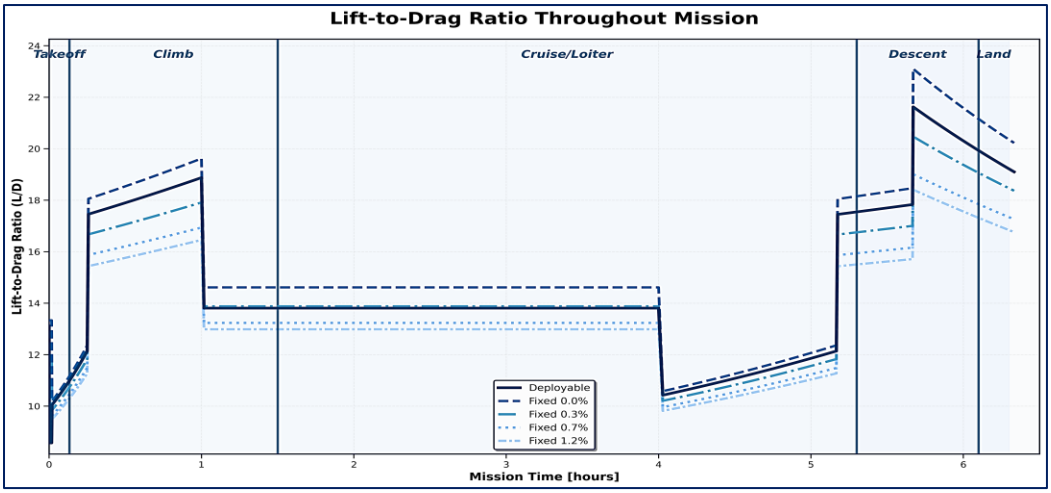

Section IV: Discussion

Both project objectives were successfully accomplished. The first objective, developing a machine learning model to predict optimal vortex generator heights, was achieved through Gaussian Process Regression models demonstrating excellent predictive accuracy with minimal error. The second objective, creating a control algorithm for VG deployment, was implemented in simulation, successfully automating vortex generator height selection across the mission profile.

The deployable system outperformed fixed vortex generator configurations in both aerodynamic efficiency and fuel consumption. Paired t-tests confirmed statistically significant improvements with large effect sizes, validating that adaptive height control provides measurable benefits. The system remained stowed during most of the mission, deploying only during high-angle-of-attack phases where separation control was needed, confirming that vortex generator benefits are highly condition-specific rather than universal.

Comparison to Previous Studies

Similar to Le Pape et al. (2012), this work demonstrated that deployable vortex generators reduce drag penalties while maintaining separation control but extended their binary deployment approach to continuous height optimization using real-time machine learning predictions rather than lookup tables. Unlike Shimomura et al. (2020) who used reinforcement learning requiring online training to predict flow conditions, the Gaussian Process framework provides uncertainty quantification without online learning, making it more practical and simpler for safety-critical aircraft applications. The multi-fidelity CFD approach builds upon Fuchi et al. (2022) by applying delta learning to predicting flow characteristics with accuracy, but this time focusing on airfoils with vortex generators rather than bluff bodies.

Limitations

Several limitations affected this study. The control system was validated using flight simulation rather than wind tunnel testing or flight trials, meaning performance in real life is uncertain. RANS turbulence modeling introduces uncertainty compared to experimental data, with validation showing errors of approximately 5-10%. The system was evaluated only for a single airfoil geometry at moderate Reynolds numbers, so evidence of generalizability is not present.

Implications and Applications

These findings demonstrate that adaptive vortex generator control can improve fuel efficiency on small aircraft and large UAVs operating at moderate Reynolds numbers. Fuel savings translate directly to extended range, reduced operating costs, and lower emissions, all of which are critical for any venture in any field. The multi-fidelity surrogate modeling approach is broadly applicable to other CFD prediction causes where high-fidelity simulations are prohibitively expensive, and this can be applied not only to bluff bodies, but to any flow control device as well. Results suggest that flow control device parameters should be treated as continuous control variables rather than discrete states, with implications for other flow control devices like morphing wings, fluidic actuators, and boundary layer suction systems.

Future Research

This work shows great potential for improvement and increased validation. Wind tunnel validation using a fully integrated physical model could provide insight to the robustness and effectiveness of sensors and servo actuators on the airfoil itself and show how much the gap in the airfoil surface would impact performance. Pressure-sensitive paint or particle image velocimetry could visualize vortex formation and boundary layer development for model refinement, which could be used to inform positions of VGs on the airfoil for different configurations. Flight testing on UAV wings would demonstrate real-world performance, assess reliability and durability, and validate fuel consumption predictions.

Additional airfoil geometries, higher Reynolds numbers, and compressibility effects should be investigated to extend applicability beyond small aircraft to commercial aviation. Coordinating vortex generator control with trajectory optimization would enable autopilots to adjust flight paths while maintaining optimal flow control effectiveness. Exploring alternative actuators like shape-memory alloys or piezoelectric devices could be explored to find the best method of actuation for deployable vortex generators.

Section V: Conclusion

This project created an AI-controlled deployable vortex generator system using multi-fidelity computational fluid dynamics and Gaussian Process regression for real-time aerodynamic simulation. The approach produced medium-fidelity RANS and high-fidelity URANS training data, resulting in surrogate models of outstanding accuracy with negligible prediction error. Flight simulation results showed that adaptive height control enhanced aerodynamic efficiency and decreased fuel burn compared to fixed vortex generators, with statistical significance established through a series of tests. The deployable system was kept retracted for the majority of the mission, only being deployed for separation control, proving the condition-specific relevance of vortex generator benefits. These results validate that continuous height optimization offers tangible efficiency benefits beyond binary deployment approaches, directly applicable to UAVs and light aircraft where adaptive flow control can mitigate parasitic drag without sacrificing safety. In a world where aviation emissions currently contribute 2.5% of global CO2 emissions, adaptive flow control systems can revolutionize passive safety components into active optimization platforms, making a significant impact on emissions reduction when implemented fleet-wide.

Section VI: References

Fuchi, K. W., Wolf, E. M., Makhija, D. S., Schrock, C. R., & Beran, P. S. (2022). Multi-fidelity machine learning applied to steady fluid flows. *International Journal of Computational Fluid Dynamics*, *36*(7), 618–640. https://doi.org/10.1080/10618562.2022.2154758

Jayanarasimhan, K., & Balasubramanian, N. K. (2025). An overview of flow control in aerodynamic surfaces using vortex generators. *Physics of Fluids*, *37*(3). https://doi.org/10.1063/5.0260937

Le Pape, A., Costes, M., Richez, F., Joubert, G., David, F., & Deluc, J.-M. (2012). Dynamic stall control using deployable leading-edge vortex generators. *AIAA Journal*, *50*(10), 2135–2145. https://doi.org/10.2514/1.j051452

Leishman, J. G. (2023, January 1). *Boundary layer flows*. Introduction to Aerospace Flight Vehicles. https://eaglepubs.erau.edu/introductiontoaerospaceflightvehicles/chapter/introduction-to-boundary-layers/

Lin, J. C. (2002). Review of research on low-profile vortex generators to control boundary-layer separation. *Progress in Aerospace Sciences*, *38*(4–5), 389–420. https://doi.org/10.1016/s0376-0421(02)00010-6

Mamman, R., Kotak, P., Weerakkody, T., Johnson, T., Krebill, A., Buchholz, J., & Lamuta, C. (2023). Deployable vortex generators for low Reynolds numbers applications powered by cephalopods inspired artificial muscles. *iScience*, *26*(12), 108369. https://doi.org/10.1016/j.isci.2023.108369

Negoita, M.-F., & Hothazie, M.-V. (2024). A machine learning-based approach for predicting aerodynamic coefficients using deep neural networks and CFD Data. *INCAS BULLETIN*, *16*(4), 91–104. https://doi.org/10.13111/2066-8201.2024.16.4.9

Portal-Porras, K., Fernandez-Gamiz, U., Zulueta, E., Ballesteros-Coll, A., & Zulueta, A. (2022). CNN-based flow control device modelling on aerodynamic airfoils. *Scientific Reports*, *12*(1). https://doi.org/10.1038/s41598-022-12157-w

Titchener, N., & Babinsky, H. (2015). A review of the use of vortex generators for mitigating shock-induced separation. *Shock Waves*, *25*(5), 473–494. https://doi.org/10.1007/s00193-015-0551-x

Zhang, Y., Zhang, S., Jiang, F., Wu, Z., & Li, W. (2025). Predicting surface pressure fields and lift coefficients of subsonic airfoils by machine-learning-enhanced compressive sensing framework. *Physics of Fluids (1994)*, *37*(9). https://doi.org/10.1063/5.0282416

Ritchie, H. (2024, April 8). *What share of global co_2 emissions come from aviation?*. Our World in Data. https://ourworldindata.org/global-aviation-emissions

Reuss, R., Hoffman, M., & Gregorek, G. (1995). *Effects of Surface Roughness and Vortex Generators on the NACA 4415 Airfoil*. https://doi.org/10.2172/206541