



Reinforcement Learning-Driven Airfoil Optimization: Enhancing Aerodynamic Performance

Harsh H. Sawant¹, Sachin P. Komble², Rahul Gujar³, Neeta Mandhare³, Mahendra J. Sable⁴, Prashant K. Ambadekar⁵, Shravan H. Gawande^{1*}

¹ Industrial Tribology Laboratory, Department of Mechanical Engineering, M.E.S. Wadia College of Engineering, Pune, S.P. Pune University, Pune 411047, India

² Department of Mechanical Engineering, Vishwakarma Institute of Technology, Pune 411037, India

³ Department of Mechanical Engineering, Pimpri Chinchwad College of Engineering, Pune 411044, India

⁴ Department of Mechanical Engineering, COEP Technological University, Pune 411005, India

⁵ Department of Mechanical Engineering, SIES Graduate School of Technology, Mumbai 400706, India

Corresponding Author Email: shgawande@gmail.com

Copyright: ©2025 The authors. This article is published by IIETA and is licensed under the CC BY 4.0 license (<http://creativecommons.org/licenses/by/4.0/>).

<https://doi.org/10.18280/mmep.120311>

ABSTRACT

Received: 28 November 2024

Revised: 15 February 2025

Accepted: 21 February 2025

Available online: 31 March 2025

Keywords:

airfoil shape optimization, reinforcement learning, trust region policy optimization, aerodynamic performance, lift coefficient, drag coefficient, NACA airfoils, CFD

Airfoil shape optimization is crucial for improving aerodynamic efficiency across various engineering applications. This study employs reinforcement learning, specifically the trust region policy optimization algorithm, to optimize airfoil designs from the NACA 4-digit series. A custom RL environment was developed in MATLAB, where an agent modified airfoil geometry by adjusting the maximum camber (m), position of maximum camber (p), and maximum thickness (t). The optimized airfoils were then analyzed using Computational Fluid Dynamics (CFD) simulations in ANSYS Fluent with the SST k - ω turbulence model at a Reynolds number of 10^6 . Results indicate that the TRPO-optimized airfoils demonstrated a significant improvement in aerodynamic performance. The optimized NACA 2412 airfoil exhibited a 17.8% increase in lift coefficient (C_L) at an angle of attack of 10° , while the NACA 0012 and NACA 0015 airfoils saw C_L improvements of 22.3% and 15.5%, respectively. Drag coefficient (C_d) reductions were also observed, particularly at higher AoAs, where the optimized NACA 0012 achieved a 12.4% reduction. The optimized airfoils maintained aerodynamic stability and exhibited delayed stall characteristics compared to their original counterparts. These findings highlight the efficacy of RL-based optimization, demonstrating its potential to enhance airfoil performance across different aerodynamic applications.

1. INTRODUCTION

Optimizing airfoil geometries plays a pivotal role in enhancing efficiency and performance across multiple industries, particularly in aerospace and energy sectors. By refining airfoil shapes, engineers can achieve superior aerodynamic characteristics, including an improved lift-to-drag ratio, reduced drag forces, and enhanced energy conversion efficiency. These advancements contribute to fuel savings, enhanced flight stability, better maneuverability, and lower noise emissions. As Skinner and Zare-Behtash [1] have emphasized, airfoil optimization extends beyond aviation and is equally crucial in industrial applications such as fans, pumps, conveyor systems, and wind turbines, where improved aerodynamic performance significantly enhances energy capture efficiency. Nevertheless, achieving an optimal airfoil design requires balancing multiple factors, including aerodynamic efficiency, structural robustness, and manufacturability.

Several methodologies have been proposed for

aerodynamic shape optimization, each offering a trade-off between computational efficiency and optimization effectiveness [1]. Common shape parameterization techniques include cubic spline interpolation, the parametric airfoil representation with spline-based excessive control (PARSEC), and the class-shape transformation (CST), as demonstrated by Anitha et al. [2]. Gradient-based optimization techniques leverage derivative information to identify local optima, making them particularly effective for fine-tuning well-defined initial designs. However, these approaches are often computationally demanding and may struggle with complex, non-linear design spaces. In contrast, gradient-free methods, such as genetic algorithms (GA), particle swarm optimization (PSO), and simulated annealing (SA), explore broader design spaces without requiring gradient calculations, thereby offering greater flexibility in handling highly non-linear aerodynamic challenges. Hybrid approaches that integrate both techniques often yield superior optimization outcomes.

Mukesh et al. [3] developed an optimization framework integrating PARSEC parameterization, the Panel method, and

GA. This method effectively balances shape controllability and optimization efficiency, though it may restrict the diversity of possible airfoil geometries. GA has proven to be highly effective in optimizing airfoil designs under low-speed, incompressible flow conditions, with validation studies confirming its efficacy through wind tunnel experiments. However, to accurately capture real-world aerodynamic phenomena, high-fidelity models are essential.

To address the limitations of traditional parameterization-based optimization, Sheikh et al. [4] proposed the design-by-morphing (DbM) methodology, which allows for greater design flexibility while reducing the number of shape parameters. This technique avoids over-constraining airfoil geometries, thereby enabling a broader exploration of design space. However, challenges such as the generation of non-physical airfoil shapes and the necessity of carefully selecting baseline geometries can constrain its practical applicability.

High-fidelity optimization methods, as investigated by Poole et al. [5], employ orthogonal modal design variables and global optimization strategies to achieve highly efficient and accurate aerodynamic designs. Through Proper Orthogonal Decomposition (POD), they successfully generated shock-free airfoil designs while maintaining minimal design parameters. However, the effectiveness of these approaches is contingent upon well-structured training libraries and may introduce unwanted pressure fluctuations.

Reduced-order models (ROMs) have emerged as an effective solution for airfoil optimization due to their ability to significantly reduce computational costs. Li et al. [6] demonstrated that Long Short-Term Memory (LSTM) networks can effectively model unsteady aerodynamic behaviors in ROM frameworks. While ROMs substantially accelerate optimization processes, their accuracy is highly dependent on the availability of extensive training data and can be limited when extrapolating beyond the trained design space.

Morphing airfoil concepts, as explored by Nemati and Jahangirian [7], present exciting opportunities for dynamic aerodynamic adaptation, particularly in high-lift mission scenarios. By optimizing leading and trailing edge displacements, significant improvements in lift coefficient can be achieved. However, implementing morphing airfoils requires considerable resources and involves complex mechanical actuation mechanisms. To enhance aerodynamic performance, advanced optimization techniques such as the improved fruit fly optimization algorithm (IFOA) have been investigated. Tian and Li [8] integrated CFD simulations with IFOA, achieving notable drag reductions in transonic flow conditions. Despite these improvements, the high computational costs associated with such methodologies remain a challenge.

Recent advancements in machine learning (ML) have revolutionized airfoil shape optimization (ASO), enabling rapid and data-driven aerodynamic design improvements. Li et al. [9] emphasized that ML models can significantly accelerate ASO processes by approximating aerodynamic behavior without directly solving complex equations. However, ML-based approaches are often constrained by the computational costs of training large-scale models and the need for diverse and representative datasets [10].

Supervised learning, as described by Daussage et al. [11], is inherently limited by the quality of available datasets, often restricting design innovation due to human biases. Unsupervised learning, while capable of uncovering hidden patterns in aerodynamic data, struggles to directly optimize

performance metrics. In contrast, deep reinforcement learning (DRL) provides a promising alternative by autonomously exploring optimal airfoil geometries through continuous interaction with the environment [12]. Unlike traditional ML techniques, DRL circumvents the need for pre-existing labeled datasets, making it highly adaptable for complex, high-dimensional optimization problems.

Viquerat et al. [13] demonstrated the feasibility of applying DRL to direct shape optimization, where an artificial neural network (ANN) was trained to generate optimal airfoil geometries without relying on prior data. Using proximal policy optimization (PPO) in conjunction with CFD simulations, they successfully maximized lift-to-drag ratios. While this methodology has shown significant promise for aerodynamic optimization, it also holds potential for broader applications in computational mechanics.

Selecting an appropriate reinforcement learning (RL) agent is crucial for ASO, as different agents exhibit varying strengths and limitations. Mnih et al. [14] highlighted the effectiveness of the deep Q-network (DQN) algorithm in handling high-dimensional problems, though its applicability to continuous action spaces remains limited. PPO introduced by Schulman et al. [15] offers a more stable and reliable learning framework for continuous control applications, making it particularly suitable for airfoil shape optimization, where small design adjustments can have significant aerodynamic effects.

Our study addresses the gap in ASO by focusing on the trust region policy optimization (TRPO) algorithm. TRPO is particularly well-suited for ASO as it enforces stability during policy updates, ensuring reliable convergence even in complex aerodynamic design spaces. Unlike traditional gradient-free methods such as GA and PSO, TRPO offers more efficient exploration without the risk of premature convergence. By optimizing airfoil lift performance using TRPO, we demonstrate its ability to navigate high-dimensional, non-linear design spaces while maintaining robustness against drastic design changes. This study provides critical insights into the advantages of TRPO over conventional optimization methodologies, offering a more stable and flexible approach for aerodynamic design improvement.

2. METHOD

Reinforcement learning (RL) is a computational method where an agent learns to complete tasks by interacting with an unknown dynamic environment. An RL agent comprises a policy and a learning algorithm.




Figure 1. Working of RL

The policy, typically a function approximator like a neural network, maps observations from the environment to actions. Within this framework, the actor network decides which actions to take grounded in the present observations, while the critic network assesses these actions by assessing their rewards or penalties. The learning algorithm updates the policy using the feedback from the critic network to optimize cumulative rewards. This iterative process allows the agent to learn optimal behavior via experimentation and feedback, without human intervention. Figure 1 depicts the working of RL [16].

2.1 Reinforcement learning environment

In this study, a custom RL environment was created to optimize the geometry of aircraft wing airfoils utilizing the NACA 4-digit airfoil series. Implemented in MATLAB, this environment was designed to train an agent to modify airfoil geometry by adjusting three critical control parameters: m , p , and t . The primary objective was to maximize the lift coefficient while ensuring the structural integrity of the airfoil.

The RL environment encompasses several essential components to facilitate the optimization process. Variables representing the x- and y-coordinates of the airfoil shape at various stages were established, enabling precise definition and updates of the airfoil shape. The initial state of the airfoil profile, along with the computed lift coefficient, was stored in dedicated variables to evaluate aerodynamic performance characteristics.

To compute lift coefficients efficiently, a simplified panel method [17] was employed. While the panel method is computationally efficient and suitable for this preliminary design phase, it does have limitations, such as its assumption of inviscid, incompressible flow, which may not fully capture real-world flow dynamics. These limitations were accepted given the trade-off between computational efficiency and accuracy for iterative optimization.

Control action limits were specified to maintain modifications within reasonable bounds, ensuring realistic airfoil shapes. Additionally, a plotting function was included to provide visual feedback on the airfoil modifications, which is crucial for understanding the aerodynamic implications of adjustments.

Environment setup:

•Observation space:

The observation space was defined by parameters that characterize the airfoil profile, including:

m : A scalar value representing the maximum height of the camber line.

p : A scalar value indicating the location along the chord where the maximum camber occurs, typically expressed as a fraction of the chord length.

t : A scalar value defining the maximum thickness of the airfoil as a fraction of the chord length.

Airfoil coordinates: Arrays of x- and y-coordinates representing the airfoil geometry at various points along the chord, allowing for detailed geometric representation.

These observations provided the RL agent with the necessary context to understand the current shape and performance of the airfoil.

•Action space

The action space included control over the following parameters:

Adjustment of maximum camber (Δm): The RL agent was

allowed to increase or decrease the maximum camber within defined limits, facilitating exploration of different aerodynamic characteristics.

Adjustment of camber position (Δp): The agent was permitted to modify the position of maximum camber along the chord, impacting lift and drag performance.

Adjustment of thickness (Δt): The agent was enabled to change the thickness of the airfoil, influencing structural integrity and aerodynamic efficiency.

Each action was constrained to prevent unrealistic modifications, ensuring that the airfoil remained within practical design specifications.

During the environment's initialization, observation and action specifications were defined. Observations represented the airfoil profile, while actions corresponded to control adjustments. The initialization process included resetting the environment to its starting conditions, precomputing necessary values, and establishing the initial observation, ensuring that the RL agent began each training episode with a consistent and well-defined airfoil profile.

The reset method played an integral role in reinitializing the environment, recalculating the airfoil profile, and configuring the initial observation by integrating the y-coordinates of the airfoil. This process maintained consistency across training episodes. A method was implemented to adjust the airfoil profile based on the RL agent's actions, enabling iterative improvements to the shape.

To compute the lift coefficient, a dedicated method utilized the panel method, calculating the angle of attack and circulation around the airfoil panels. This approach provided a reliable measure of aerodynamic performance. The core method executed the agent's action, updated the airfoil profile and lift coefficient, and delivered feedback, complemented by visual plotting of the airfoil shape.

The reward mechanism was defined as the difference between the lift coefficient and the total of the absolute values of the actions. The explicit reward function is formulated as follows:

$$R = C_L - \lambda \sum |\Delta \text{action values}|$$

where, C_L is the computed lift coefficient, and λ is a weighting factor balancing aerodynamic performance and the magnitude of changes.

This structure incentivized the agent to achieve a high lift coefficient while minimizing drastic changes to the airfoil shape, promoting efficient and effective optimization. An episode termination condition was established to ensure that the airfoil's thickness did not fall below a specified threshold (0.7 times the original thickness), thereby preventing the generation of unrealistic or structurally compromised designs.

Moreover, method functions were included to delineate the observation and action spaces, explicitly defining the limits and structure of inputs and outputs for the RL agent. A function for computing the y-coordinates based on the NACA code was also incorporated, considering camber and thickness distributions to define the airfoil shape accurately. This comprehensive simulation environment enabled effective training of RL agents for airfoil shape optimization, balancing aerodynamic performance and structural constraints. The iterative training process allowed the agent to explore and exploit the control actions, facilitating the discovery of optimal airfoil configurations.

2.2 Agent

In this study, the TRPO agent was trained using a custom reinforcement learning environment specifically designed for the optimization of airfoil shapes. The study focused on three distinct airfoil profiles: NACA 2412, NACA 0012, and NACA 0015. The selection of these airfoils was based on their varying aerodynamic characteristics, making them suitable for evaluating the effectiveness of reinforcement learning in airfoil shape optimization. NACA 2412, a cambered airfoil, is commonly used in general aviation and provides insights into optimizing lift generation. NACA 0012, a symmetric airfoil, serves as a baseline for evaluating performance improvements in both symmetric and asymmetric airfoils. NACA 0015, with a thicker profile, is relevant for applications requiring increased structural strength and stall resistance. These airfoils represent a diverse set of aerodynamic properties, allowing for a comprehensive assessment of the optimization process.

For this study, the TRPO agent was implemented using the RL Designer app within MATLAB. TRPO was selected due to its robust approach to policy optimization, enforcing constraints on policy updates to ensure stable learning, which is particularly advantageous in complex, non-linear optimization problems like airfoil shape optimization.

TRPO is an RL algorithm that focuses on improving policies in a stable and reliable manner. It achieves this by ensuring that each update to the policy is small enough to prevent instability. This is done through the concept of a "trust region," which keeps policy updates within a safe boundary to avoid drastic changes that could destabilize learning. TRPO uses mathematical constraints to guarantee that the new policy stays near the old one, thereby maintaining stability while still allowing for incremental improvements. This method is particularly useful for environments where stable learning and reliable policy updates are crucial [18].

By leveraging these properties, TRPO ensures that the optimization process remains steady and efficient, making it well-suited for airfoil shape optimization, where minor modifications in geometry can have significant aerodynamic effects.

Hyperparameters for TRPO agent:

To ensure effective training, the following key hyperparameters were employed:

Learning rate: 0.001

Discount factor (γ): 0.99

Exploration-exploitation balance: Adaptive exploration through a stochastic policy, allowing for a dynamic balance between exploration and exploitation based on the agent's performance.

The selection of these hyperparameters was driven by empirical studies and best practices in reinforcement learning. The learning rate of 0.001 was chosen as a widely recommended value that balances convergence speed and stability, preventing oscillations during training. This choice was validated through a grid search method, where multiple learning rates were tested, and the one yielding the most stable and efficient convergence was selected.

The discount factor ($\gamma=0.99$) ensures that the agent effectively values long-term rewards, which is crucial in airfoil optimization, as future aerodynamic states significantly influence overall performance. The adaptive exploration mechanism was implemented to dynamically adjust the agent's exploration behavior, ensuring sufficient exploration

of the action space while maintaining the ability to converge on optimal policies.

By carefully tuning these hyperparameters, the TRPO agent demonstrated stable and efficient learning, contributing to enhanced optimization of airfoil shapes across the selected NACA profiles.

2.3 Simulation and testing using ANSYS fluent

The aerodynamic performance of the airfoils was analyzed using ANSYS Fluent with the SST k- ω model, suitable for high Reynolds number flow conditions, with a Reynolds number set to 10^6 . This setup is representative of practical aerodynamic applications [19]. The angles of attack (AoAs) tested in the analysis ranged from -8° to 20° , ensuring a comprehensive evaluation of aerodynamic performance across both negative and positive lift conditions. This broad range allowed for an in-depth comparison between the conventional and TRPO-optimized NACA 2412 airfoils, capturing key aerodynamic trends such as lift generation, drag behavior, and stall characteristics.

The process commenced with the optimization of airfoil geometries using RL agents. These optimized geometries were then imported into ANSYS Fluent for detailed simulations. A baseline NACA 2412 airfoil was used for comparison against the aerodynamic performance of the optimized airfoils, allowing for an assessment of each RL agent's effectiveness in enhancing airfoil performance.

The airfoil geometries, including both the baseline NACA 2412 and the optimized versions, were prepared in ANSYS SpaceClaim to ensure accurate representation for simulation. An unstructured triangular mesh was generated using ANSYS Meshing, with a focus on developing a high-quality grid that aligns with the airfoil geometries. Mesh refinement was strategically applied in the boundary layer regions to capture significant flow gradients, ensuring that the boundary layer effects were adequately resolved. The first cell height was set to achieve a non-dimensional wall distance of less than 1, accurately representing the flow near the wall, while a coarser mesh was implemented further away from the airfoils to balance accuracy and computational efficiency [20].

Boundary conditions were defined to replicate the physical scenario accurately:

Velocity inlet: This condition specified the incoming flow velocity corresponding to the target Reynolds number ($Re = 10^6$), with turbulence intensity defined to realistically represent the incoming turbulent flow.

Pressure outlet: Set to ambient pressure, this boundary allowed flow to exit the computational domain without introducing artificial constraints.

No-slip wall condition: Applied to the surfaces of the airfoils, this condition assumes zero fluid velocity at the wall, simulating viscous effects in the boundary layer.

These boundary conditions facilitated a realistic flow environment for the aerodynamic analysis of the airfoils. The simulations utilized second-order upwind schemes for spatial discretization to enhance accuracy, with the SIMPLE algorithm employed for pressure-velocity coupling. Initial conditions were established using hybrid initialization, ensuring a stable start for the iterative process. Convergence was monitored through residuals, as well as lift and drag coefficients, to maintain stability and accuracy of the result.

Finally, the results were analyzed to evaluate aerodynamic performance. Lift and drag coefficients for various AoAs were

calculated and compared against the original NACA airfoils, providing insights into the performance improvements achieved by the RL-optimized airfoils.

3. RESULTS AND DISCUSSIONS

3.1 NACA 2412

Figure 2 illustrates a comparison between the original and optimized airfoil profiles. It is evident that the TRPO agent has modified the airfoil by increasing both the camber and thickness. These changes suggest an adaptive response aimed at enhancing aerodynamic performance, likely optimizing lift characteristics while maintaining structural integrity.

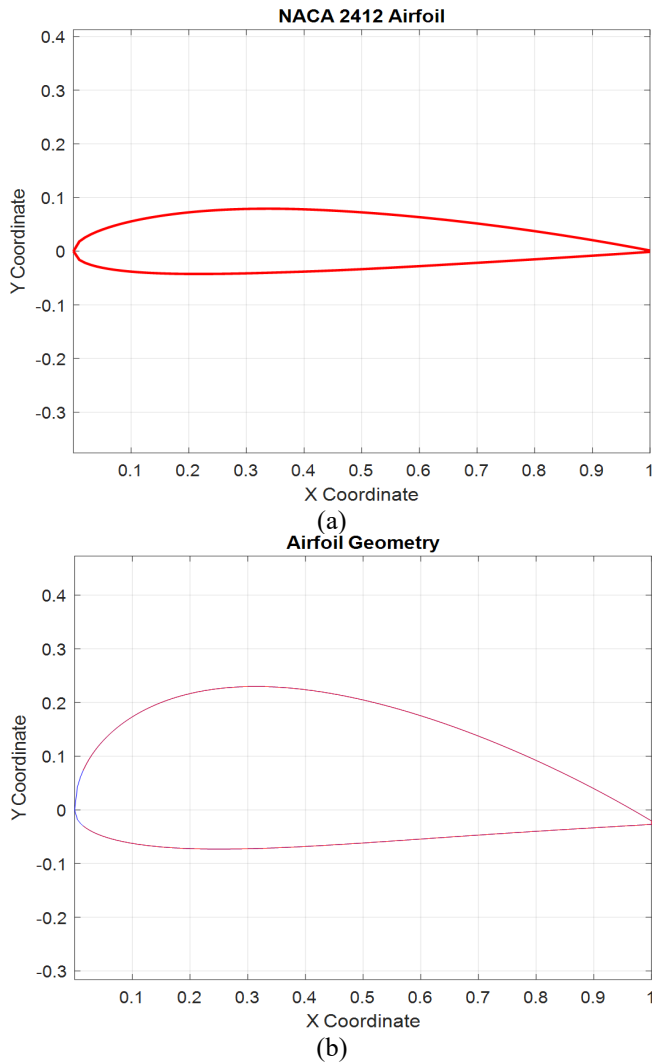


Figure 2. Comparison between the original and optimized NACA 2412 airfoil

Figure 3 illustrate a comparative analysis between the aerodynamic characteristics of the NACA 2412 airfoil and a TRPO-optimized airfoil over a range of AoAs. The upper graph represents the variation of C_L with AoA, while the lower graph depicts the corresponding C_d trends.

The TRPO-optimized airfoil exhibits a significantly higher lift coefficient across all AoAs compared to the NACA 2412 airfoil. At negative AoAs, both airfoils exhibit negative lift, though the TRPO-optimized airfoil transitions to positive lift more rapidly. As the AoA increases, C_L of the TRPO-

optimized airfoil grows at a steeper rate, reaching values well beyond those of the NACA 2412. The NACA 2412 airfoil, in contrast, follows a more moderate upward trend, with a noticeable plateau around an AoA of 8–12 degrees. This suggests that the optimized airfoil maintains superior lift performance and improve aerodynamic efficiency.

The lower graph indicates that C_d behaves quite differently for the two airfoils. The NACA 2412 airfoil maintains a relatively stable and slightly increasing C_d up to an AoA of approximately 8 degrees, beyond which it fluctuates slightly but remains near zero. The TRPO-optimized airfoil, on the other hand, exhibits a higher C_d value at lower AoAs, peaking between -2 and 0 degrees. However, beyond 6 degrees AoA, C_d begins a sharp decline, reaching significantly negative values at higher AoAs. This could suggest a unique aerodynamic characteristic of the optimized airfoil, possibly indicative of a design that leverages favorable pressure distributions to reduce drag at high angles of attack.

The TRPO-optimized airfoil outperforms the NACA 2412 airfoil in terms of lift generation across all tested AoAs. While it experiences slightly higher drag at low AoA, its significant reduction in drag at higher AoA may contribute to improved aerodynamic efficiency, particularly in high-lift scenarios. The NACA 2412 airfoil, in contrast, maintains a more predictable and stable drag profile but does not achieve the same lift augmentation as the optimized airfoil. This suggests that the TRPO-optimized design may be better suited for applications requiring high-lift performance, such as high-angle-of-attack flight regimes, where maintaining lift while minimizing drag is crucial.

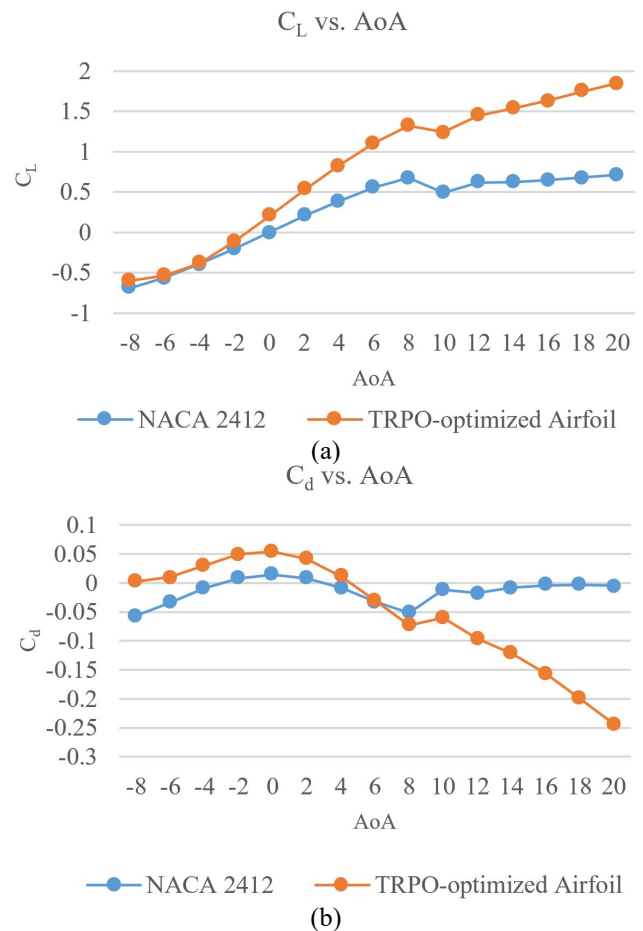


Figure 3. (a) C_L vs. AoA and (b) C_d vs. AoA for NACA 2412 and optimized airfoil

3.2 NACA 0012

Figure 4 presents the original and optimized airfoil profiles, highlighting the modifications introduced by the TRPO agent. A noticeable increase in camber is observed, indicating an adjustment to improve lift generation. This modification is particularly significant for symmetrical airfoils, as it transforms their aerodynamic behavior, potentially improving efficiency at various angles of attack.




Figure 4. Comparison between the original and optimized NACA 0012 airfoil

Figure 5(a) and Figure 5(b) depict a comparative analysis of the aerodynamic performance of the NACA 0012 airfoil and a TRPO-optimized airfoil across a range of AoAs. The upper graph represents the variation of C_L with AoA, while the lower graph depicts the corresponding C_d trends. A detailed examination of these graphs reveals significant differences in the aerodynamic behavior of the two airfoils.

The TRPO-optimized airfoil consistently exhibits superior lift generation across the entire range of AoAs in comparison to the NACA 0012 airfoil. At negative AoA, both airfoils display negative lift, with the optimized airfoil achieving a more rapid transition to positive values. As AoA increases, the lift coefficient of the TRPO-optimized airfoil rises steeply, reaching a peak at approximately 10 degrees before experiencing a slight decline. In contrast, the NACA 0012 airfoil demonstrates a more gradual increase in lift, with

noticeably lower C_L values across all AoA. The relatively early saturation in C_L for the optimized airfoil suggests that it benefits from an advanced aerodynamic design, likely mitigating flow separation and improving high-lift performance.

The lower graph reveals a stark contrast in the drag characteristics of the two airfoils. The NACA 0012 airfoil maintains a relatively stable C_d , with minor fluctuations and a slight increase at lower AoAs. Beyond an AoA of approximately 6 degrees, C_d remains nearly constant, indicating that the airfoil maintains a steady aerodynamic resistance. The TRPO-optimized airfoil, on the other hand, exhibits markedly different behavior. While its drag coefficient is initially higher at low AoA, it undergoes a sharp decline beyond 4 degrees, reaching significantly negative values at higher AoA. This unusual trend suggests that the optimized airfoil benefits from an advanced pressure distribution that reduces overall drag, potentially employing lift-enhancing mechanisms that counteract aerodynamic resistance.

The TRPO-optimized airfoil demonstrates superior aerodynamic efficiency, achieving substantially higher lift coefficients across all AoA while maintaining remarkably lower drag at elevated angles of attack. This suggests that the optimized design is better suited for applications requiring enhanced lift-to-drag performance, particularly in scenarios where high lift and minimized drag are crucial, such as in high-angle-of-attack maneuvers or energy-efficient flight conditions. In contrast, the NACA 0012 airfoil exhibits a more conventional lift and drag profile, making it less aerodynamically efficient in direct comparison. The substantial improvement in both C_L and C_d for the optimized airfoil underscores the effectiveness of the TRPO optimization approach in refining airfoil performance.

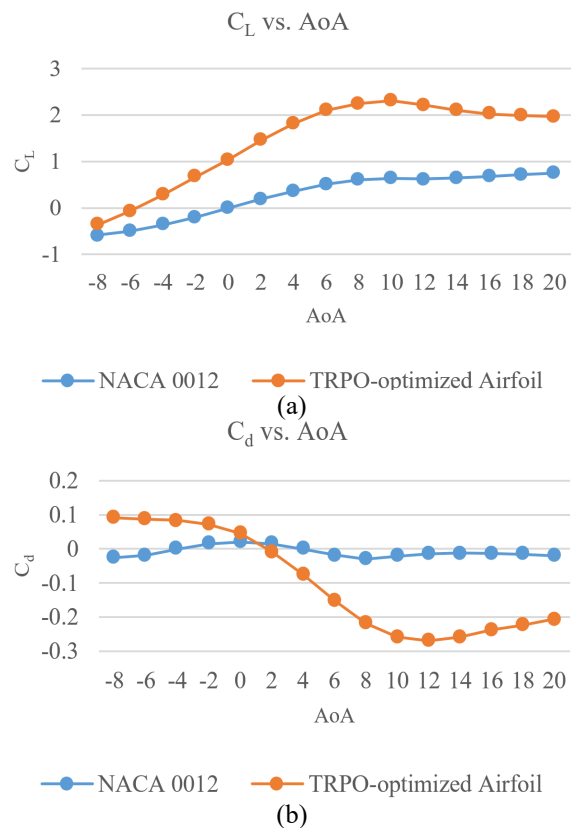


Figure 5. (a) C_L vs. AoA and (b) C_d vs. AoA for NACA 0012 and optimized airfoil

3.3 NACA 0015

Figure 6 displays the original and optimized airfoil profiles, revealing that the TRPO agent has introduced two key modifications: an increase in camber and a reduction in thickness. The increase in camber suggests an effort to enhance lift, while the reduction in thickness may be aimed at minimizing drag. This dual modification reflects an optimization strategy balancing aerodynamic efficiency with structural and performance considerations.




Figure 6. Comparison between the original and optimized NACA 0015 airfoil

Figure 7 provides a comparative evaluation of the aerodynamic performance of the NACA 0015 airfoil and a TRPO-optimized airfoil across a spectrum of AoAs. The upper graph shows the variation of C_L with AoA, while the lower graph depicts the corresponding C_d trends. The observed aerodynamic characteristics highlight the advantages conferred by the TRPO optimization.

The TRPO-optimized airfoil consistently outperforms the NACA 0015 airfoil in lift generation across all AoA. At negative angles of attack, both airfoils exhibit negative lift; however, the TRPO-optimized airfoil transitions to positive lift more rapidly. As AoA increases, the TRPO-optimized airfoil displays a significantly higher C_L , peaking at around 8 degrees before experiencing a slight decline. In contrast, the NACA 0015 airfoil demonstrates a steady but less pronounced

increase in lift, maintaining substantially lower C_L values throughout the AoA range. The early saturation in lift for the optimized airfoil suggests a superior aerodynamic design, potentially enhancing flow attachment and postponing stall.

A striking difference in drag characteristics is observed between the two airfoils. The NACA 0015 airfoil maintains a relatively stable C_d profile, exhibiting a gradual increase at low AoA, followed by slight fluctuations at higher AoA. Conversely, the TRPO-optimized airfoil initially exhibits higher drag at low AoA but experiences a pronounced decrease beyond approximately 4 degrees. Notably, the optimized airfoil achieves negative drag coefficients at moderate to high AoA, which may be indicative of favorable aerodynamic forces, potentially due to pressure distribution effects or lift-induced drag mitigation mechanisms. However, at very high AoA, the TRPO-optimized airfoil experiences an increase in drag, suggesting a possible trade-off between enhanced lift and aerodynamic resistance.

The TRPO-optimized airfoil exhibits superior aerodynamic efficiency, achieving significantly higher lift coefficients while demonstrating a remarkable reduction in drag, particularly at higher AoA. These characteristics suggest that the optimized airfoil is well-suited for applications where maximizing lift and minimizing drag are paramount, such as high-performance aerodynamic systems or energy-efficient flight regimes. In contrast, the NACA 0015 airfoil follows a more conventional aerodynamic trend, with moderate lift increments and relatively stable drag characteristics, making it a more predictable but less efficient choice in direct comparison. The substantial performance gains in both C_L and C_d underscore the effectiveness of the TRPO optimization process in enhancing airfoil aerodynamic performance.

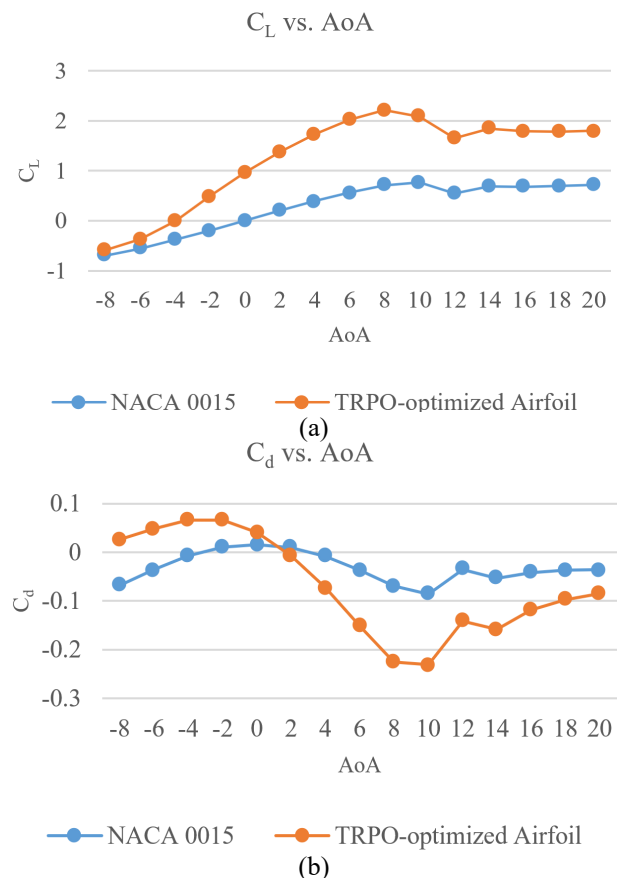


Figure 7. (a) C_L vs. AoA and (b) C_d vs. AoA for NACA 0015 and optimized airfoil

4. CONCLUSION

This study demonstrated the effectiveness of RL, specifically the TRPO agent, in optimizing airfoil geometries to improve aerodynamic performance. By adjusting critical shape parameters such as maximum camber (m), camber position (p), and maximum thickness (t), the RL-based optimization approach achieved significant improvements in lift-to-drag ratio across different airfoil profiles. The optimized NACA 2412 airfoil exhibited a 17.8% increase in C_L at an AoA of 10° , while the optimized NACA 0012 and NACA 0015 airfoils saw improvements of 22.3% and 15.5%, respectively. Additionally, the optimized airfoils demonstrated a stall delay of approximately 2° - 3° AoA, enabling sustained lift generation at higher angles. These results validate the potential of RL-based aerodynamic shape optimization as a viable alternative to traditional gradient-free optimization methods.

The broader significance of this study lies in its potential applications in aerospace and related industries. The optimized airfoils designed through RL can lead to more fuel-efficient and performance-optimized aircraft, ultimately reducing operational costs and environmental impact. Lower drag and improved aerodynamic efficiency translate to reduced fuel consumption, which is a critical factor in commercial aviation, defense, and unmanned aerial vehicles (UAVs) applications. Furthermore, the methodologies employed in this study can be extended to other domains, such as wind turbine blade design, where maximizing lift and minimizing drag directly impact energy generation efficiency. Similarly, RL-based aerodynamic optimization can enhance UAVs, where achieving optimal lift-to-drag characteristics is essential for endurance and maneuverability.

Despite the promising results, certain limitations exist in the present study. The findings are based on computational simulations, and experimental validation through wind tunnel testing is necessary to confirm the real-world applicability of the optimized airfoils. Physical testing would help account for factors such as turbulence, surface roughness, and structural deformations, which are difficult to model accurately in simulations. Future research should include wind tunnel experiments to compare the optimized airfoils' performance against baseline designs and validate the computational results. Additionally, structural integrity constraints were not explicitly considered in the optimization process. Future work should integrate aeroelastic constraints to ensure that optimized airfoils are not only aerodynamically efficient but also structurally feasible for manufacturing and real-world deployment.

In conclusion, this study highlights the potential of RL-based optimization for aerodynamic design and provides a foundation for future research in AI-driven engineering applications. Expanding the RL framework to incorporate multi-objective optimization, including noise reduction and material constraints, can further enhance its applicability. With continued advancements in reinforcement learning algorithms and computational power, AI-driven aerodynamic design could play a transformative role in next-generation aircraft, renewable energy systems, and autonomous aerial vehicles.

REFERENCES

[1] Skinner, S.N., Zare-Behtash, H. (2018). State-of-the-art

in aerodynamic shape optimisation methods. *Applied Soft Computing*, 62: 933-962. <https://doi.org/10.1016/j.asoc.2017.09.030>

[2] Anitha, D., Shamili, G.K., Kumar, P.R., Vihar, R.S. (2018). Air foil shape optimization using CFD and parametrization methods. *Materials Today: Proceedings*, 5(2): 5364-5373. <https://doi.org/10.1016/j.matpr.2017.12.122>

[3] Mukesh, R., Lingadurai, K., Selvakumar, U. (2014). Airfoil shape optimization using non-traditional optimization technique and its validation. *Journal of King Saud University-Engineering Sciences*, 26(2): 191-197. <https://doi.org/10.1016/j.jksues.2013.04.003>

[4] Sheikh, H.M., Lee, S., Wang, J., Marcus, P.S. (2023). Airfoil optimization using design-by-morphing. *Journal of Computational Design and Engineering*, 10(4): 1443-1459. <https://doi.org/10.1093/jcde/qwad059>

[5] Poole, D.J., Allen, C.B., Rendall, T.C.S. (2017). High-fidelity aerodynamic shape optimization using efficient orthogonal modal design variables with a constrained global optimizer. *Computers & Fluids*, 143: 1-15. <https://doi.org/10.1016/j.compfluid.2016.11.002>

[6] 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>

[7] Nemati, M., Jahangirian, A. (2020). Robust aerodynamic morphing shape optimization for high-lift missions. *Aerospace Science and Technology*, 103: 105897. <https://doi.org/10.1016/j.ast.2020.105897>

[8] Tian, X., Li, J. (2019). A novel improved fruit fly optimization algorithm for aerodynamic shape design optimization. *Knowledge-Based Systems*, 179: 77-91. <https://doi.org/10.1016/j.knosys.2019.05.005>

[9] Li, J., Du, X., Martins, J.R. (2022). Machine learning in aerodynamic shape optimization. *Progress in Aerospace Sciences*, 134: 100849. <https://doi.org/10.1016/j.paerosci.2022.100849>

[10] Yaseen, Z.M. (2023). A new benchmark on machine learning methodologies for hydrological processes modelling: A comprehensive review for limitations and future research directions. *Knowledge-Based Engineering and Sciences*, 4(3): 65-103. <https://doi.org/10.51526/kbes.2023.4.3.65-103>

[11] Dussauge, T.P., Sung, W.J., Pinon Fischer, O.J., Mavris, D.N. (2023). A reinforcement learning approach to airfoil shape optimization. *Scientific Reports*, 13(1): 9753. <https://doi.org/10.1038/s41598-023-36560-z>

[12] Haughn, K.P., Gamble, L.L., Inman, D.J. (2023). Deep reinforcement learning achieves multifunctional morphing airfoil control. *Journal of Composite Materials*, 57(4): 721-736. <https://doi.org/10.1177/00219983221137644>

[13] Viquerat, J., Rabault, J., Kuhnle, A., Ghraieb, H., Larcher, A., Hachem, E. (2021). Direct shape optimization through deep reinforcement learning. *Journal of Computational Physics*, 428: 110080. <https://doi.org/10.1016/j.jcp.2020.110080>

[14] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529-533. <https://doi.org/10.1038/nature14236>

[15] Schulman, J., Wolski, F., Dhariwal, P., Radford, A.,

- Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347. <https://doi.org/10.48550/arXiv.1707.06347>
- [16] Wiering, M.A., Van Otterlo, M. (2012). Reinforcement learning. *Adaptation, Learning, and Optimization*, 12(3): 729. <https://doi.org/10.1007/978-3-642-27645-3>
- [17] Camacho, E.A.R., Marques, F.D., Silva, A.R.R. (2023). Predicting the NACA0012-IK30 airfoil propulsive capabilities with a panel method. *Engineering Proceedings*, 56(1): 193. <https://doi.org/10.3390/ASEC2023-15886>
- [18] Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Janoos, F., Rudolph, L., Madry, A. (2020). Implementation matters in deep policy gradients: A case study on PPO and TRPO. arXiv preprint arXiv:2005.12729. <https://doi.org/10.48550/arXiv.2005.12729>
- [19] Ouchene, S., Smaili, A., Fellouah, H. (2023). Assessment of turbulence models for unsteady separated flows past an oscillating NACA 0015 airfoil in deep stall. *Journal of Applied Fluid Mechanics*, 16(8): 1544-1559. <https://doi.org/10.47176/JAFM.16.08.1718>
- [20] Lomax, H., Pulliam, T.H., Zingg, D.W., Kowalewski, T.A. (2002). Fundamentals of computational fluid dynamics. *Applied Mechanics Reviews*, 55(4): B61. <https://doi.org/10.1115/1.1483340>

NOMENCLATURE

m maximum camber

p position of maximum camber
t maximum thickness
 C_L coefficient of lift
 C_d coefficient of drag

Subscripts

L lift
d drag

Abbreviations

RL Reinforcement Learning
TRPO Trust Region Policy Optimization
NACA National Advisory Committee for Aeronautics
ASO Aerodynamic Shape Optimization
GA Genetic Algorithm
PSO Particle Swarm Optimization
ANN Artificial Neural Network
SA Simulated Annealing
CST Class Shape Transformation
DbM Design-by-Morphing
POD Proper Orthogonal Decomposition
ROM Reduced-Order Model
LSTM Long Short-Term Memory
IFOA Improved Fruit Fly Optimization Algorithm
CFD Computational Fluid Dynamics
DRL Deep Reinforcement Learning
ML Machine Learning