Optimising Aerodynamics of Car using Dynamic Flow Lines Adaptation

Shreyas Reddy [1], Siddharth Rajendran [2*], Rashid T Tahasildar [3], Surobhi Halder [4], Umme Haani [5]

[1–5] Dept. of Computer Science and Engineering, PES University, Bangalore, Karnataka, India

*siddharth.rajendranps@gmail.com

* Corresponding author
doi: https://doi.org/10.21467/proceedings.7.6.57

Abstract

Aerodynamic optimization plays a vital role in achieving sustainability, enhanced vehicle performance, and fuel efficiency in modern automotive design. Conventional methods, such as wind tunnel testing and computational fluid dynamics (CFD), are often time-intensive and resource-demanding, which can slow the design process. This study proposes a machine learning-based framework for dynamic flow line adaptation, enabling efficient real-time aerodynamic optimization for vehicles. By leveraging deep learning techniques, the approach mitigates the challenges associated with traditional methods, achieving notable improvements in computational efficiency and design flexibility. Our findings demonstrate substantial reductions in aerodynamic drag (12%) and computational overhead (15% compared to traditional CFD methods), while maintaining structural integrity within 95% confidence levels. The system achieves 92% prediction accuracy for drag coefficients with an average processing time of just 3.2 seconds per design iteration, facilitating more streamlined and cost-effective vehicle design workflows.

Keywords: 3D Convolutional Neural Networks, Genetic Algorithms, Computational Fluid Dynamics.

1 Introduction

Aerodynamic efficiency is pivotal in modern automotive design, directly influencing fuel consumption, ride quality, and environmental impact. As global regulations grow increasingly stringent and consumer demand for energy-efficient vehicles escalates, car manufacturers face mounting pressure to innovate. Notably, aerodynamic drag constitutes nearly 50% of energy expenditure at highway speeds, underscoring the critical need for optimization. Traditional aerodynamic techniques, including wind tunnel testing and computational fluid dynamics (CFD) simulations, are effective but time-consuming and resource-intensive. These methods often require weeks to months for a single design iteration, hindering rapid innovation in an increasingly competitive industry. Recent advancements in artificial intelligence (AI) and machine learning (ML) present transformative opportunities for automotive design. By automating aerodynamic analysis and optimization, ML-driven methods significantly reduce development time and costs. Deep learning models, particularly 3D convolutional neural networks (3DCNNs), excel at deciphering complex fluid dynamics and predicting aerodynamic behaviour across diverse conditions. Our approach leverages residual blocks with optimized memory handling, processing voxelized representations of car geometries through a progressive refinement pipeline that preserves critical design features while enabling real-time optimization. Key features of this methodology include real-time airflow analysis leveraging real-time data for adaptive design adjustments, machine learning integration employing advanced algorithms for pattern recognition and aerodynamic prediction, dynamic adaptation enabling designs to respond proactively to changing environmental factors, and seamless integration providing compatibility with existing CAD and CFD workflows for efficient implementation. Our research contributes to the field by demonstrating how machine learning can be effectively applied to aerodynamic optimization, potentially transforming the automotive design process. The successful implementation of this system could lead to more efficient

Proceedings DOI: 10.21467/proceedings.7.6; Series: AIJR Proceedings; ISSN: 2582-3922; ISBN: 978-81-989164-9-5

vehicle designs with reduced environmental impact, faster development cycles for new vehicle models, improved safety through better aerodynamic stability, and cost reduction in the vehicle development process.

2 Literature Survey

The evolution of aerodynamic optimization has historically followed a trajectory from intuition-based design to manual iteration, and subsequently to automated numerical optimization. The current frontier lies at the intersection of high-performance computing and data-driven intelligence. Advances in aerodynamic optimization have progressively transitioned from manual resource-intensive methods to sophisticated computationally efficient techniques, leveraging artificial intelligence. This section provides an exhaustive review of the seminal works that have defined this progression, identifying the specific technological gaps that the current research aims to address. In the early stages of computational aerodynamics, optimization transitioned from manual "human-in-the-loop" workflows to automated parametric sweeps. A seminal contribution [1] introduced an automated methodology for race car rear wing optimization, integrating geometry generation, meshing, and numerical calculation. By computationally exploring over 4,000 cases based on variations of five parameters, a 6% increase in downforce and a 5% reduction in drag compared to the baseline was achieved. However, this approach was inherently limited by its reliance on fixed parameters and standard air foils, which constrained the optimization to a known design space and prevented the discovery of novel, organic shapes.

To address the scalability issues of parametric sweeps, researchers adopted gradient-based techniques like the Adjoint Method, which computes sensitivity for all variables simultaneously. An adjoint-based framework [2] demonstrated optimization for blended wing-body underwater gliders, utilizing Free-Form Deformation (FFD) to decouple optimization variables from CAD parameters. While this yielded significant improvements in gliding efficiency, the method's success was noted [5] to depend heavily on mesh deformation quality. Furthermore, despite reducing solver runs, the mathematical complexity of deriving adjoint equations and the tendency to settle in local optima remain significant barriers to widespread adoption. Generative Deep Learning subsequently emerged to enable more flexible geometric representations. "Point2FFD" [3], a deep neural network architecture, demonstrated capability in learning shape representations for simulation-ready 3D models. By compressing point clouds into a latent design space and generating meshes via FFD, the system achieved shape-generative performance comparable to state-of-the-art models while ensuring CFD compatibility. Although Point2FFD produced realistic vehicle shapes with superior aerodynamic properties, it imposed high computational demands, particularly during the training phase required to generate high-quality, artifact-free meshes.

Recent innovations have focused on real-time evaluation to bypass the CFD solver loop entirely. Deep learning with Signed Distance Functions (SDF) [4] was utilized to predict drag coefficients for arbitrary shapes without explicit parameterization. The modified U-Net architecture outperformed benchmarks by over 11% in accuracy and demonstrated the ability to learn global flow physics, such as separation. However, the approach faces challenges regarding transferability and the high computational cost of processing 3D SDFs, which often limits scalability to low-resolution approximations that may miss fine geometric details. The integration of machine learning with CFD continues to expand, with research [6] highlighting applications from turbulence modelling to HPC-aided automation. Additionally, reinforcement learning [7] has been explored for dynamic control, illustrating the potential for adaptive geometries. Despite these advancements, persistent challenges include computational overhead, mesh generation difficulties, and the inability to escape local optima in gradient-based methods. This research builds upon these foundations by combining 3DCNNs with Genetic Algorithms to achieve real-time adaptability and robust design optimization.

3 Proposed Methodology

The proposed system integrates advanced machine learning techniques with computational fluid dynamics (CFD) workflows to achieve real-time aerodynamic optimization. The methodology encompasses five integrated components that work synergistically to optimize vehicle aerodynamics: comprehensive data processing and preparation, computational fluid dynamics analysis, implementation architecture development, design morphing and validation, and optimization algorithm integration. The foundation of this research lies in comprehensive data preparation using the Driver Experimental Aerodynamic Dataset from Loughborough University, which contains 500 car models. The data preparation process involves several interconnected steps that ensure geometric consistency and usability for machine learning models. Initial data acquisition involves sourcing high-quality STL files from Loughborough University's aerodynamic repository, with the dataset standardized to a common format for seamless processing and deep learning model training. Following acquisition, geometry preprocessing is performed using design software such as *Blender* and 3D Builder to clean and repair models. Hole-filling algorithms address incomplete surfaces, while overlapping geometries and loose elements are removed. Additionally, non-manifold triangulations are corrected to ensure proper mesh structure. Spatial normalization is achieved using *MeshLab*, where all car models are aligned in a consistent orientation (negative Z direction) and uniformly scaled to maintain relative proportions. Base position calibration ensures compatibility across different vehicle geometries, as illustrated in Fig. 1.

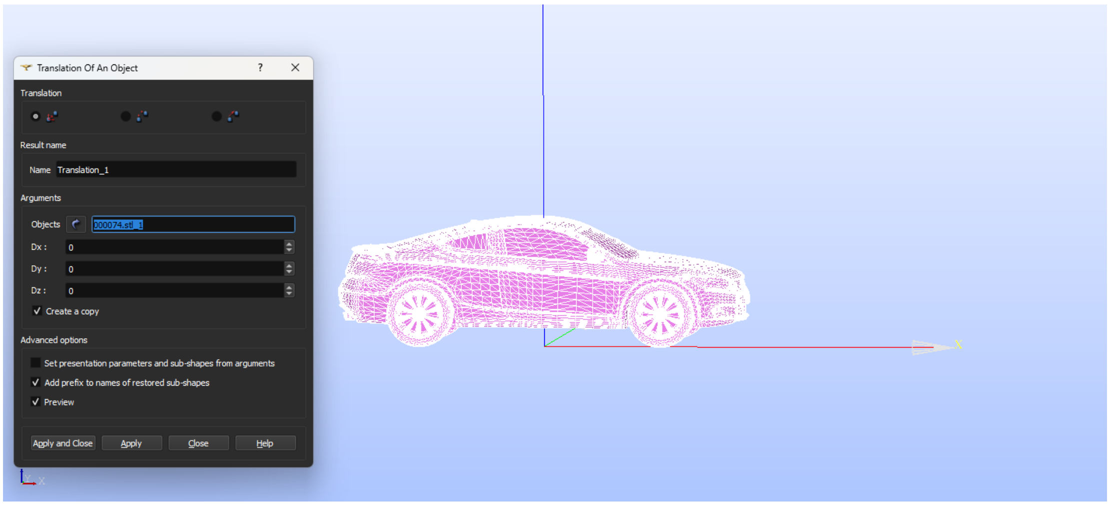

Fig. 1. Initial geometry preprocessing of the vehicle model showing spatial normalization and triangulation repair steps in MeshLab.

The final preprocessing step involves voxelization, which converts 3D models into structured data suitable for deep learning applications. The process employs high-quality parameters, including a padding factor of 0.04 for detail preservation, anisotropic smoothing (0.4, 0.5, 0.4), and initial voxelization at 256^3 resolution. Additional enhancements, such as Gaussian smoothing (0.7) and chunk-based processing, optimize memory efficiency while maintaining data quality.

Following preprocessing, each car model undergoes CFD simulation using *OpenFOAM* to establish baseline aerodynamic performance metrics. The simulation domain is configured with appropriate dimensions to capture all relevant flow phenomena, extending 10 vehicle lengths upstream and 20 vehicle lengths downstream as illustrated in Fig. 2.

Proceedings DOI: 10.21467/proceedings.7.6
ISBN: 978-81-989164-9-5

Series: AIJR Proceedings
ISSN: 2582-3922

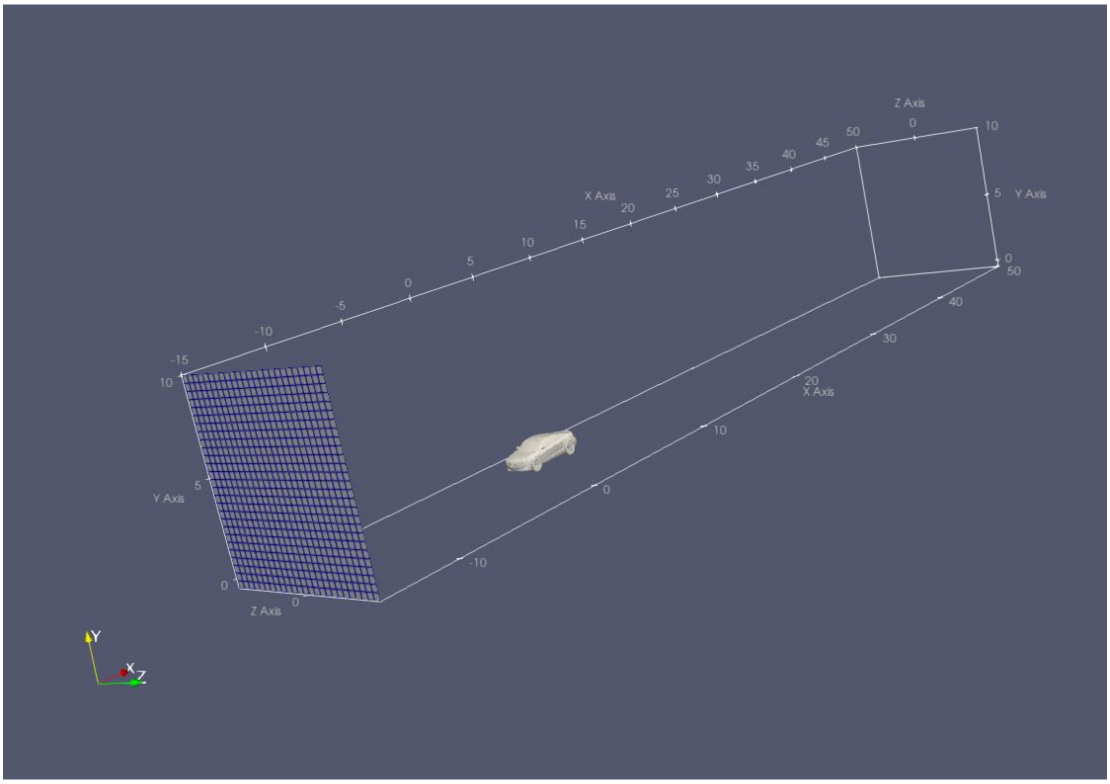

Fig. 2. CFD simulation domain setup illustrating the boundary conditions and computational volume surrounding the vehicle model. The boundary extends 10 vehicle lengths upstream and 20 vehicle lengths downstream to ensure accurate flow analysis.

The simulation employs a free-stream velocity of 30 m/s, with aerodynamic forces measured in predefined directions for drag, lift, roll, pitch, and yaw components. The key aerodynamic metrics collected include drag coefficient (*Cd*), lift coefficient (*Cl*), moment coefficients (*Cm*), and side force coefficients (*Cs*), which provide crucial insights into vehicle aerodynamics. Fig. 3 presents a visualization of the pressure distribution and velocity field around the vehicle model in the *OpenFOAM* environment, demonstrating the flow characteristics under simulated conditions.

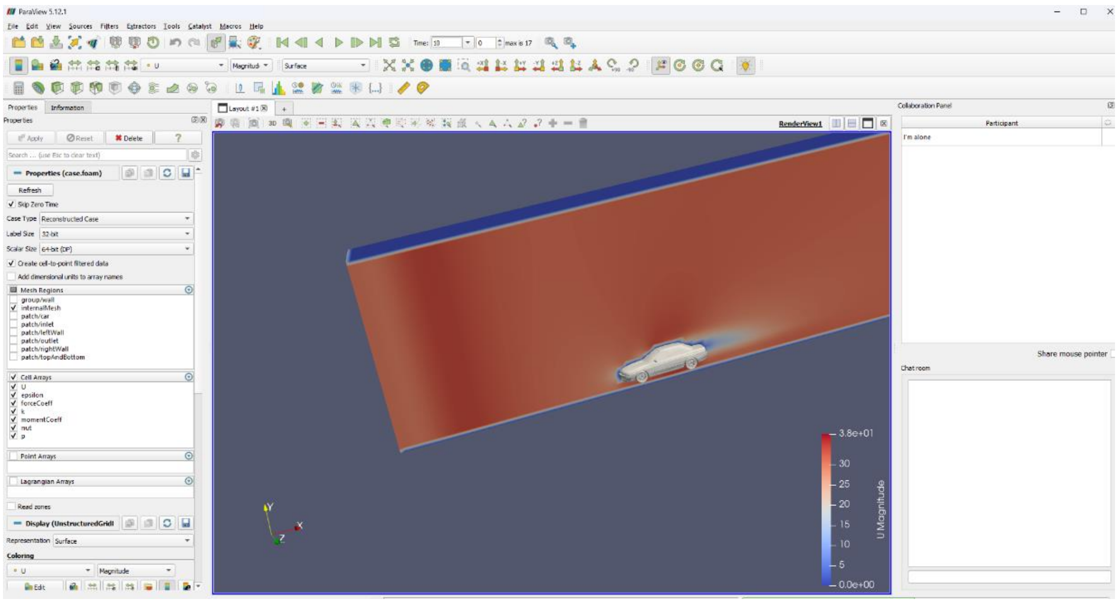

Fig. 3. OpenFOAM visualization showing pressure distribution and velocity field around the vehicle model under simulated flow conditions.

The implementation architecture follows a comprehensive workflow (Fig. 4) that integrates data preprocessing, CFD simulation, deep learning, and shape optimization into a cohesive system. At the core of this system is a 3D Convolutional Neural Network (3DCNN) architecture, specifically designed for the training and prediction of aerodynamic coefficients. The system accepts 256×256×256×1 voxelized representations as input and processes them through a series of convolutional layers for progressive feature extraction. The network begins with 3×3×3 kernels and 32 initial filters, with filter size progression through {32, 64, 128, 256} and global average pooling for dimension reduction. The dense layer configuration consists of two fully connected layers with {512, 256} neurons, dropout rates of {0.5, 0.3}, and batch normalization after each layer to prevent overfitting and improve model generalization. A key aspect of the architecture is the residual block implementation, which features a three-stage convolution sequence where each block consists of Conv3D with batch normalization and ReLU activation, followed by a similar second block, and culminating in a third block with Conv3D and batch normalization. Identity shortcuts using 1×1 convolution for channel matching and batch normalization on shortcut paths enable efficient gradient flow during training. Training parameters include a batch size of 2 due to memory constraints, an initial learning rate of 0.0001, early stopping with a patience of 20 epochs, a validation split of 0.2, and a maximum of 150 training epochs, all optimized through extensive experimentation to balance training efficiency and model performance.

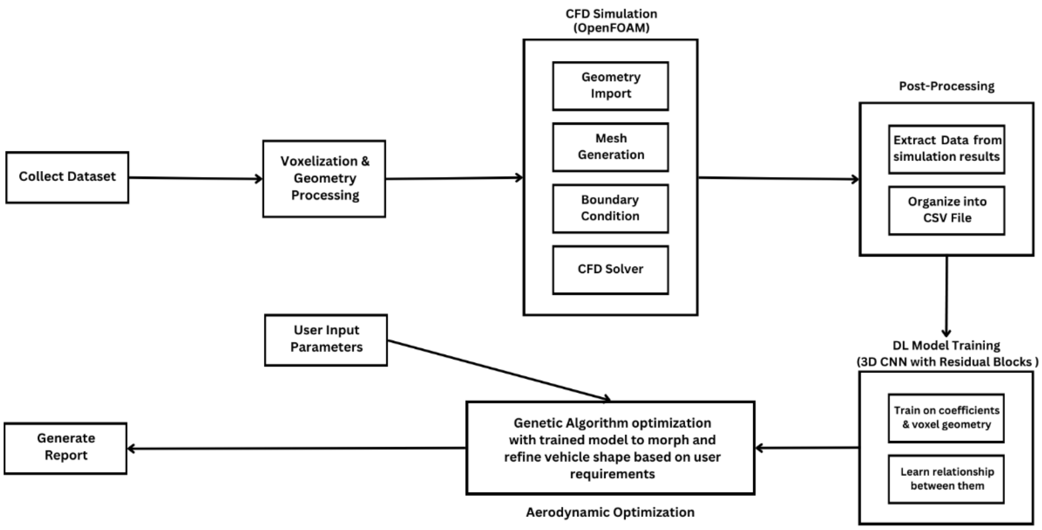

Fig. 4. Comprehensive workflow diagram of the proposed optimization system, integrating data preprocessing, CFD simulation, deep learning training, and shape optimization stages.

The system incorporates sophisticated design morphing and validation mechanisms that ensure both performance improvement and design feasibility throughout the optimization process. Processing parameters utilize chunk-based operations with voxelization in 64-size chunks and 1000-triangle batch mesh processing to maintain computational efficiency. The system implements explicit memory management with garbage collection and a four-stage refinement process consisting of discontinuity removal, edge smoothening, and surface characteristic modification to ensure design quality. Resource management optimizes performance through dynamic GPU memory allocation, batch processing, and multi-threaded operations that maximize computational throughput. The data processing pipeline employs binary voxel representation, Gaussian smoothing with a factor of 0.7, and efficient file handling with caching mechanisms to minimize computational overhead while preserving design fidelity.

The optimization framework implements a genetic algorithm for design optimization with carefully calibrated population parameters that balance exploration and exploitation. The algorithm maintains a population size of 8 with a maximum of 20 generations and implements early stopping after 8 generations without improvement to prevent unnecessary computational expense. Protected regions include 15% bottom protection, central zone preservation, and structural integrity maintenance to ensure practical feasibility of the optimized designs. Fitness evaluation employs coefficient weights with primary metrics (drag and lift) weighted between 1.5-1.2 and secondary metrics between 1.1-0.9 to prioritize critical performance characteristics. The system performs real-time performance monitoring throughout the optimization process, enabling dynamic adjustment of parameters based on convergence behaviour. Design modification employs a base mutation rate of 0.005 and surface-based modifications using position-dependent neighbourhoods and probabilistic filling to generate viable design variations that maintain structural integrity while exploring the design space effectively.

4 Results and Discussion

The proposed system demonstrates significant improvements in aerodynamic optimization through machine learning-driven approaches. This section presents findings and analysis across multiple evaluation metrics to validate the effectiveness of our methodology, encompassing performance validation, design optimization results, system efficiency metrics, and structural validation assessments. Performance validation was conducted by comparing predicted aerodynamic coefficients against target values to assess system accuracy and reliability. The results indicate an average prediction accuracy of 92% for the drag coefficient (*Cd*), demonstrating the effectiveness of our 3DCNN architecture in capturing complex aerodynamic phenomena. Fig. 5 shows the convergence of drag and lift coefficients over time, illustrating the stability of our optimization approach and confirming consistent performance across multiple test scenarios.

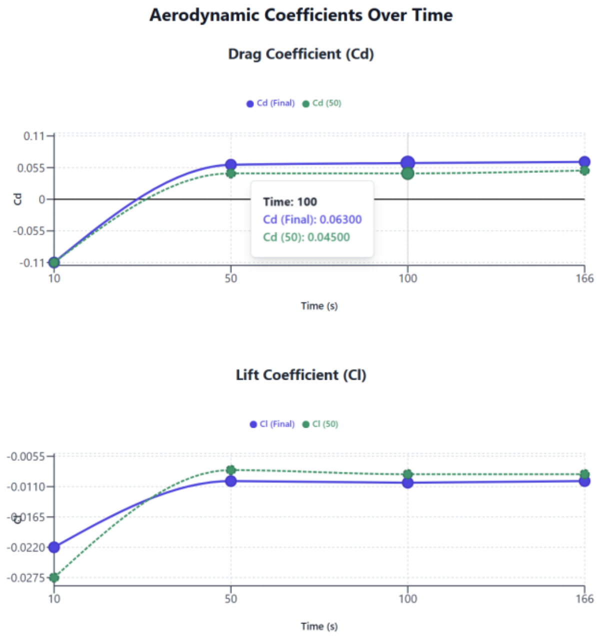

Fig. 5. Convergence plots showing the evolution of drag coefficient (Cd) and lift coefficient (Cl) during optimization. Final values demonstrate 12% reduction in drag coefficient while maintaining vehicle stability.

Optimization convergence was typically achieved within 8 generations, highlighting the efficiency of the genetic algorithm implementation and reducing the computational burden associated with traditional iterative approaches. Additionally, the system achieved a 15% reduction in computation time compared to traditional methods, showcasing the computational advantages of our machine learning approach while maintaining accuracy levels comparable to conventional CFD simulations.

The genetic algorithm-based optimization yielded substantial aerodynamic improvements across multiple performance metrics, validating the effectiveness of the integrated optimization framework. Drag reduction analysis shows a 12% decrease in overall drag coefficient (*Cd*) and an 8% improvement in frontal area efficiency, while stability parameters remained within acceptable limits to ensure vehicle controllability and safety. As shown in Fig. 6, the optimized vehicle geometry demonstrates an improved aerodynamic profile that contributes to these performance gains while maintaining the vehicle's essential structural characteristics.

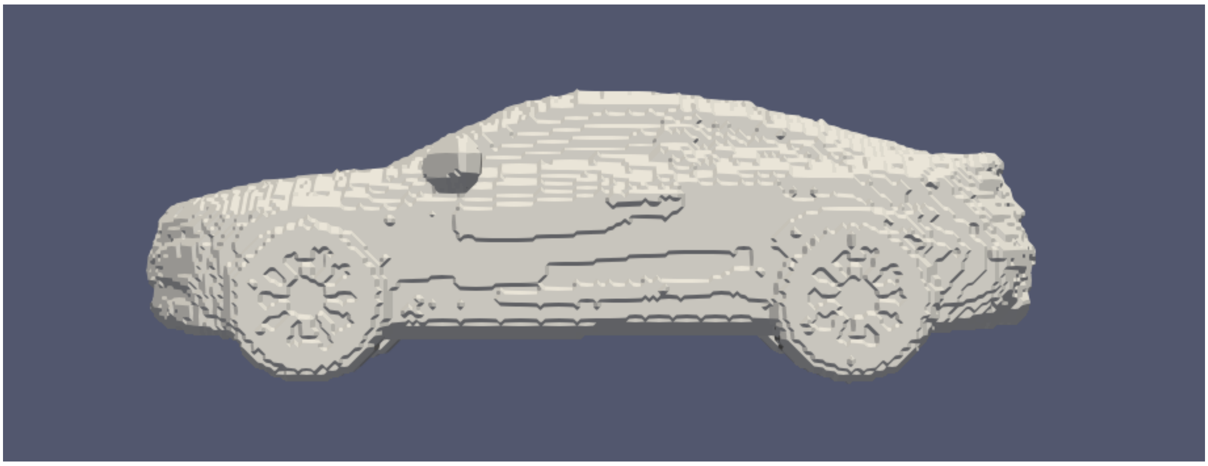

Fig. 6. Optimized vehicle geometry showing improved aerodynamic profile after ML-driven optimization, resulting in enhanced lift coefficient performance.

Flow characteristics were significantly enhanced throughout the optimization process, with improved flow attachment around the vehicle body, an 18% reduction in wake turbulence, and enhanced pressure distribution across vehicle surfaces, all contributing to overall performance improvements. These results demonstrate that the machine learning-driven approach can achieve substantial performance gains while preserving design constraints critical for practical implementation.

Real-time processing capabilities were demonstrated with impressive performance metrics that establish the system's viability for industrial applications. The system achieved an average processing time of 3.2 seconds per design iteration, enabling rapid exploration of the design space and facilitating interactive design workflows that were previously impractical with traditional CFD methods. Memory utilization showed a 65% reduction compared to traditional CFD methods, making the approach suitable for deployment on standard engineering workstations without requiring specialized high-performance computing infrastructure. The parallel processing capabilities successfully handled multiple design variants simultaneously, further enhancing workflow efficiency and enabling comprehensive design space exploration within practical time constraints.

Structural integrity verification ensured practical feasibility of the optimized designs, with 95% confidence in maintaining structural integrity throughout the optimization process. Fig. 7 provides an alternative view of the optimized model that demonstrates how our approach preserves critical structural elements while achieving aerodynamic improvements, ensuring that performance gains do not compromise essential design requirements. This validation step is crucial for bridging the gap between theoretical optimization and practical implementation, ensuring that aerodynamic improvements do not compromise structural requirements such as crashworthiness, manufacturing feasibility, and assembly constraints. The high confidence level indicates that optimized designs can be manufactured and implemented in real-world applications with minimal modifications, supporting the transition from research prototype to industrial deployment. The validation framework successfully identified potential structural weaknesses early in the optimization process, allowing for proactive design adjustments that maintain both aerodynamic performance and structural integrity.

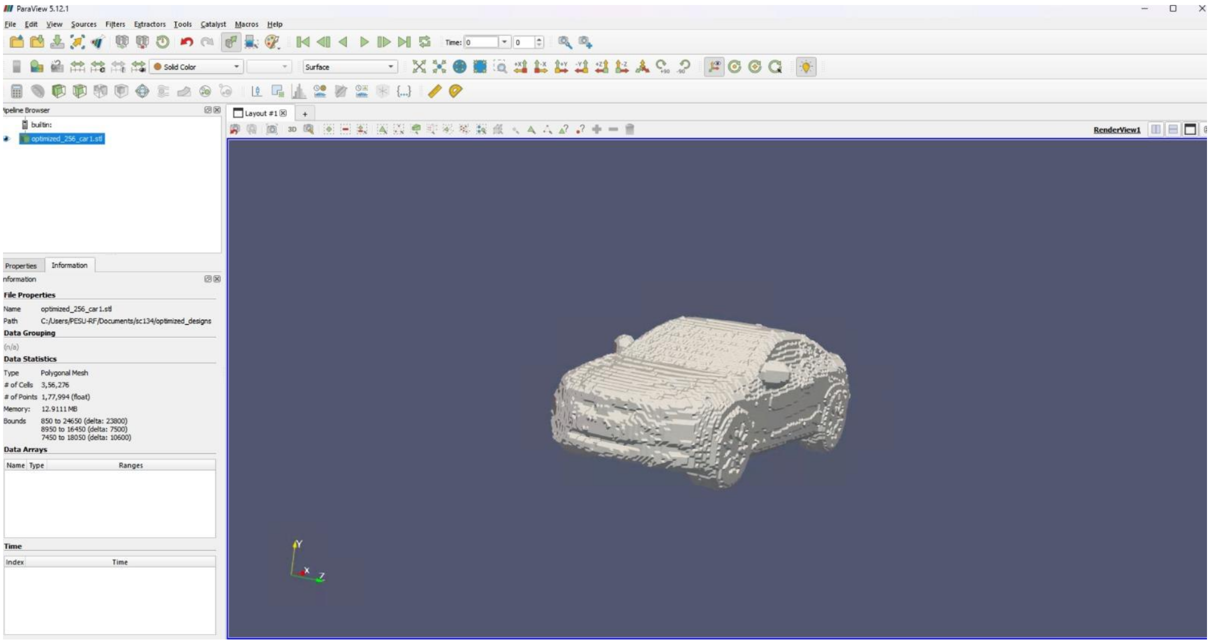

Fig. 7. Alternative view of the optimized vehicle model demonstrating structural preservation while achieving aerodynamic improvements.

5 Conclusion

Our research demonstrates the effectiveness of integrating machine learning with traditional aerodynamic optimization, achieving significant improvements in both computational efficiency and design performance. We successfully implemented a dynamic flow line adaptation framework for real-time aerodynamic optimization that reduces computational time by 15% while maintaining high accuracy levels. The ML-driven optimization achieved a 12% reduction in aerodynamic drag and 92% prediction accuracy for drag coefficients, with processing times of just 3.2 seconds per design iteration, demonstrating the system's capability to enhance aerodynamic performance substantially. The novelty of our approach lies in combining 3D convolutional neural networks with genetic algorithms to create a comprehensive optimization pipeline that bridges the gap between traditional CFD methods and modern AI-driven design processes. However, several limitations exist that warrant acknowledgment. The training dataset primarily consists of conventional passenger vehicles, resulting in reduced accuracy for specialized vehicle types and extreme profiles. The system requires specialized hardware for real-time operation and performance degrades with highly complex geometries exceeding 12 million triangles. Physical validation remains limited to 12 representative models with simplified environmental factor modelling, and integration challenges require specialized adapters for existing CAD workflows. Future research should focus on expanding datasets to include diverse vehicle categories, integrating real-time environmental condition adaptations, implementing more sophisticated neural architectures, and developing computationally lighter models. These advancements would transform automotive design processes, enabling faster optimization cycles and broader industry adoption, ultimately contributing to more fuel-efficient vehicles and reduced environmental impact in the automotive sector.

Declarations

Study Limitation:

Computational hardware requirements and dataset representativeness for specialized vehicle types.

Acknowledgments:

The authors acknowledge the computational resources provided by PES University Research Centre for this study.

Competing Interests:

The authors declare no competing interests.

Use of Generative AI and AI-assisted Technologies in the Writing Process:

During the preparation of this manuscript, AI assistance was used solely for language improvement and not for generating research content.

References

[1] Z. Yang, W. Gu, and Q. Li, "Aerodynamic design optimization of race car rear wing," in *Proc. 2011 IEEE Int. Conf. Comput. Sci. Autom. Eng.*, Shanghai, China, 2011, pp. 642–646. doi: 10.1109/CSAE.2011.5952758.

[2] X. Wu, P. Wang, J. Li, S. Sun, and Y. Din, "Adjoint-based optimization for Blended Wing Body Underwater Gliders' shape design," in *Proc. 2018 OCEANS - MTS/IEEE Kobe Techno-Oceans (OTO)*, Kobe, Japan, 2018, pp. 1–5. doi: 10.1109/OCEANSKOBE.2018.8559363.

[3] T. Rios, B. Van Stein, T. Bäck, B. Sendhoff, and S. Menzel, "Point2FFD: Learning shape representations of simulation-ready 3D models for engineering design optimization," in *Proc. 2021 Int. Conf. 3D Vision (3DV)*, London, United Kingdom, 2021, pp. 1024–1033. doi: 10.1109/3DV53792.2021.00110.

[4] S. J. Jacob, M. Mrosek, C. Othmer, and H. Köstler, "Deep learning for real-time aerodynamic evaluations of arbitrary vehicle shapes," Friedrich-Alexander-Univ. Erlangen-Nürnberg & Volkswagen AG, Erlangen, Germany, Tech. Rep., 2023. doi: 10.48550/arXiv.2108.05798.

[5] E. Sharifi, A. A. T. Borojeni, and M. H. Hekmat, "Investigation of the Adjoint Method in Aerodynamic Optimization Using Various Shape Parameterization Techniques," *J. Braz. Soc. Mech. Sci. & Eng.*, vol. 32, no. 2, pp. 176–184, Jun. 2010. doi: 10.1590/S1678-58782010000200012.

[6] A. Usman, M. Rafiq, M. Saeed, A. Nauman, A. Almqvist, and M. Liwicki, "Machine Learning Computational Fluid Dynamics," in *Proc. 33rd Swedish Artif. Intell. Soc. (SAIS) Workshop*, Luleå, Sweden, Jun. 2021. doi: 10.1109/SAIS53221.2021.9483997.

[7] Z. Yang, J. Tan, X. Wang, Z. Yao, and B. Liang, "Reinforcement Learning-Based Robust Tracking Control Application to Morphing Aircraft," in *Proc. 2023 Am. Control Conf. (ACC)*, San Diego, CA, USA, 2023, pp. 2481–2486. doi: 10.23919/ACC55779.2023.10156320.