

Advancing Road Vehicle Design: Machine Learning-based Aerodynamic Shape Optimization

PhD Candidate:

Alberto Vilariño Tarrío

Supervisors:

Carlos Sanmiguel Vila (INTA-UC3M)

Rodrigo Castellanos (UC3M)



PhD Aerospace Engineering UC3M

Doctoral Meetings 2026


2nd and 3rd June 2026

INTRODUCTION

Current aerodynamic development is still an iterative loop: geometry change, mesh, CFD/WT assessment, expert interpretation and redesign.



Constrains



THE OPPORTUNITY

In recent times the development cycles had shortened, with a need for faster solutions the integration of AI is being promoted.



GOAL: Can we use ML to explore a realistic 3D vehicle design space while keeping the aerodynamic feedback sufficiently accurate to support design decisions?

Recent literature map: datasets and surrogate models are moving very fast

The field has shifted from small ad-hoc CFD studies to large public/industrial ML datasets.

Datasets

DrivAerNet++ 8k+ designs; multimodal DrivAer variants; drag, surface and flow-field data

DrivAerML 500 morphed DrivAer notchbacks; validated scale-resolving workflow

AhmedML 500 Ahmed-body variants; time-accurate hybrid RANS/LES, rich wake physics

DrivAerStar 12k industrial STAR-CCM+ simulations; CAD/FFD parameters and wind-tunnel validation

Surrogate methods

Geometry networks RegDGCNN, PointNet/PointTransformer, mesh and point-cloud models

Neural operators FNO, Transolver/Transolver++, AB-UPT, graph neural operators

Implicit fields TripNet/triplane models for Cd, surface fields and volume queries

Industrial surrogates Interactive race-car/vehicle design with expert-validated CFD and GIST-type models

Interpretation

The opportunity is no longer “can ML predict drag?” The open question is how to couple ML, high-fidelity physics, constraints and global optimization in a reliable workflow.

Sources: DrivAerNet++ (2024), DrivAerML (2024), AhmedML (2024), TripNet (2025), DrivAerStar (2025), CarBench (2025), Dallara/IBM GIST (2026).

Current status

~65 candidate geometry parameters defined, including dependencies, constraints, expected aerodynamic relevance and implementation complexity.

Why morphing?


It preserves a coherent vehicle shape while enabling design-space exploration under geometric constraints.

Why DrivAer?

It is a realistic open reference geometry, more relevant than oversimplified bodies for transient automotive aerodynamics.

Next decision

Select the first reduced set of parameters for LES feasibility tests and optimization-loop validation.



Checkpoint message: the design space is being narrowed scientifically before launching expensive 3D campaigns.

Scientific bottleneck: global optimization needs trustworthy flow information

Why this is not a simple “optimize Cd” problem

Geometry


Realistic vehicles have many coupled parameters. A local change can affect the underbody, wheel flow, separation and rear wake.

Physics

The wake is three-dimensional, separated and unsteady. Similar drag can hide different pressure and lift mechanisms.

Learning


ML models learn the bias of their data. A biased CFD layer can make the optimizer converge to non-physical trends.




RESEARCH PROGRESS



- Geometry preparation
- Meshing
- Morphing
- Automation



- Genetic algorithm
- Well known
- Demonstrated results
- “In house” developed
- Great exploration-exploitation



SOD2D



- CFD solver
- High fidelity turbulence simulations (LES)
- GPU native
- Strong HPC and AI integration

STRATEGY AND TOOLBOX

ACHIEVEMENTS AND FUTURE WORK

Achievements

- Establish and document the workflow
- Define vehicle parameters
- Evaluate and select appropriate software tools:
 - Ansa
 - Develop meshing criteria and quality standards
 - Design a robust morphing strategy
 - Explore automation opportunities for preprocessing workflow
 - SOD2D
 - Execute initial simulation studies
 - Investigate automation approaches for simulation setup and execution
 - HyGO
- Dissemination of Research Findings
 - Presented at EUROGEN 2025
 - Accepted for presentation at the 17th World Congress on Computational Mechanics (WCCM 2026), July 2026
- Co-authored the paper:
 - *“Optimization-embedded active multi-fidelity surrogate learning for multi-condition airfoil shape optimization”*
- Presenter at “TCS-Siemens executive Forum on AI in industries, R&D and manufacturing”.

Ongoing activities

- Finalizing CFD simulation setup and validation.
- Implementing the complete workflow using the selected toolchain.

Future work

- Integrate physics-based data to improve agent performance
- Explore reinforcement learning and other advanced AI agents



Questions?

Advancing Road Vehicle Design: Machine Learning-based Aerodynamic Shape Optimization

PhD Student: Alberto Vilariño Tarrío

Supervisors: Carlos Sanmiguel Vila (INTA-UC3M), Rodrigo Castellanos (UC3M)

Where the PhD can contribute

What the literature already does well

Large datasets, fast surrogate inference, mesh/point-cloud/transformer models, and increasingly realistic vehicle geometries.

Remaining gap

Most work is still centered on geometry-to-Cd, pressure or mean-field prediction. Less attention is paid to unsteady wake modes, ROMs and robust global optimization loops.

Proposed thesis position

An end-to-end workflow coupling constrained morphing, multi-fidelity CFD, LES-informed ROMs/surrogates and HyGO-based global search.

1 Static → dynamic

include force fluctuations, wake asymmetry and modal amplitudes

2 Single fidelity → multi-fidelity

use cheap CFD for breadth and LES anchors for physics

3 Prediction → decision

connect surrogates to actual design-space search

4 Cd only → wake-aware


optimize aerodynamic performance plus stability/robustness proxies

One-sentence contribution

Develop and validate a multi-fidelity, LES-informed global optimization workflow for realistic vehicle geometry, where the learned model is constrained by wake physics rather than only trained to interpolate Cd.

This positioning is consistent with recent benchmark studies showing progress in surrogate accuracy but persistent questions about physical consistency, generalization and workflow integration.

What LES adds to the optimization problem



ROMs enabled by LES

POD / DMD / SPOD for interpretable wake modes; autoencoder or neural-ODE latent models for nonlinear reduced dynamics; modal coefficients as compact design descriptors.

Examples of objective quantities

Mean Cd remains the primary target, but LES enables base-pressure recovery, lift/yaw RMS, pressure RMS, wake asymmetry, dominant POD-mode amplitude and switching frequency.

Why this is relevant for cars

Vehicle wakes show three-dimensional separation, coherent structures and bi-stability. Geometry changes can alter the wake state without being visible in a steady scalar metric.

Recommended message for the committee

LES is not only a more expensive CFD option. It is the data source that makes the thesis physics-aware: ROMs, wake-state diagnostics and robust surrogate validation.