CZECH TECHNICAL UNIVERSITY IN PRAGUE

FACULTY OF ELECTRICAL ENGINEERING
DEPARTMENT OF CYBERNETICS

Surrogate Machine Learning Model Workflow for Prediction of Aerodynamic Forces on Vehicles

Bachelor's Thesis

Robert Bakhtigariev

Prague, May 2025

Study programme: Open Informatics
Branch of study: Artificial Intelligence and Computer Science

Supervisor: doc. Ing. Karel Zimmermann, Ph.D

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my advisor at Porsche Engineering, Ing. Jan Sedláček, as well as to the entire team at Porsche Engineering for their expert guidance, provision of computational resources, workspace, and overall support throughout the course of this project.

I would also like to thank my academic supervisor, doc. Ing. Karel Zimmermann, Ph.D., for his valuable advice in the field of Machine Learning (ML) and for his support during the writing of this thesis.

BACHELOR‘S THESIS ASSIGNMENT

I. Personal and study details

Student's name: **Bakhtigariev Robert** Personal ID number: **516262**
Faculty / Institute: **Faculty of Electrical Engineering**
Department / Institute: **Department of Cybernetics**
Study program: **Open Informatics**
Specialisation: **Artificial Intelligence and Computer Science**

II. Bachelor’s thesis details

Bachelor’s thesis title in English:

Surrogate Machine Learning Model Workflow for Prediction of Aerodynamic Forces on Vehicles

Bachelor’s thesis title in Czech:

Proces pro samostatný model strojového učení pro predikci aerodynamických sil na vozidla

Guidelines:

The goal of this thesis is to investigate state-of-the-art surrogate machine learning (ML) models for aerodynamics, focusing on wind pressure prediction. The study will extend the existing MeshGraphNet-based model developed by the student during the semestral work at Porsche Engineering Services s.r.o. by incorporating uncertainty quantification techniques. The enhanced model will be benchmarked against experimental or Computational Fluid Dynamics (CFD) results.
1. Literature Review
a) Study and describe the current state-of-the-art surrogate ML models in aerodynamics [1,2].
b) Study techniques for uncertainty prediction in ML models [3,4,5].
2. Model Extension
a) Integrate selected ML models into existing CFD simulation Porsche Engineering workflow with pre and post-processing.
b) Build upon the existing MeshGraphNet-based wind pressure prediction model developed in your semestral project.
c) Incorporate uncertainty quantification using approaches such as ensemble sampling or uncertainty loss [3,4,5].
3. Model Evaluation
a) Estimate aerodynamic forces using the enhanced model on a test dataset.
b) Benchmark the model predictions against experimental or CFD results.
c) Analyze generalization of the proposed model ability to unseen data.

Bibliography / sources:

[1] Baque, Pierre, et al. "Geodesic Convolutional Shape Optimization." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018. Link: https://arxiv.org/pdf/1802.04016
[2] Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., & Battaglia, P. (2021). "MeshGraphNets: Machine Learning on Meshes with Graph Networks." NeurIPS.
[3] Gal, Y., & Ghahramani, Z. (2016). "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning." Link: https://arxiv.org/abs/1506.02142
[4] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). "Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles." Link: https://arxiv.org/abs/1612.01474
[5] Kendall, A., & Gal, Y. (2017). "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" Link: https://arxiv.org/abs/1703.04977

Name and workplace of bachelor’s thesis supervisor:

doc. Ing. Karel Zimmermann, Ph.D. Vision for Robotics and Autonomous Systems FEE

Name and workplace of second bachelor’s thesis supervisor or consultant:

Date of bachelor’s thesis assignment: **12.02.2025** Deadline for bachelor thesis submission: __________

Assignment valid until: **20.09.2026**

prof. Dr. Ing. Jan Kybic
Head of department’s signature

prof. Mgr. Petr Páta, Ph.D.
Vice-dean´s signature on behalf of the Dean

III. Assignment receipt

The student acknowledges that the bachelor’s thesis is an individual work. The student must produce his thesis without the assistance of others, with the exception of provided consultations. Within the bachelor’s thesis, the author must state the names of consultants and include a list of references.

Date of assignment receipt

Student’s signature

FAKULTA ELEKTROTECHNICKÁ
FACULTY OF ELECTRICAL ENGINEERING
Technická 2
166 27 Praha 6

DECLARATION

I, the undersigned

Student's surname, given name(s): Bakhtigariev Robert
Personal number: 516262
Programme name: Open Informatics

declare that I have elaborated the bachelor’s thesis entitled

Surrogate Machine Learning Model Workflow for Prediction of Aerodynamic Forces on Vehicles

independently, and have cited all information sources used in accordance with the Methodological Instruction on the Observance of Ethical Principles in the Preparation of University Theses and with the Framework Rules for the Use of Artificial Intelligence at CTU for Academic and Pedagogical Purposes in Bachelor’s and Continuing Master’s Programmes.

I declare that I used artificial intelligence tools during the preparation and writing of this thesis. I verified the generated content. I hereby confirm that I am aware of the fact that I am fully responsible for the contents of the thesis.

In Prague on 23.05.2025

Robert Bakhtigariev
....
student's signature

Abstract

This thesis investigates the use of Machine Learning (ML) models for aerodynamic prediction in the early stages of automotive design. The objective is to replace or complement traditional methods with faster and more computationally efficient alternatives.

The Graph Neural Network (GNN) MeshGraphNet is evaluated as a baseline for predicting surface aerodynamic quantities. Its limitations, including the lack of global geometric information about the vehicle shape and insufficient encoding of local geometric features, are analyzed. A new architecture called LaplaceAeroNet is proposed to address these issues and demonstrates improved accuracy and robustness. The model reliably predicts aerodynamic behavior even for vehicle geometries containing local features such as wheels, spoilers, or variations in ride height that were not included in the training data. To estimate the model's confidence, an ensemble-based uncertainty quantification method is employed.

The second part of the thesis is an optional extension beyond the primary focus and addresses the prediction of the full three-dimensional aerodynamic flow around the vehicle. A proposed model called FourierAeroNet combines a Fourier Neural Operator (FNO) with U-Net refinement through the U-Fourier Neural Operator (U-FNO) architecture and outperforms a convolutional baseline by more effectively capturing long-range flow structures.

The results show that ML models can provide accurate and efficient surrogate predictions, supporting aerodynamic analysis and design decisions in the early phases of vehicle development.

Keywords Vehicle Aerodynamics, Surrogate Modeling, Graph Neural Networks, Fourier Neural Operator, Computational Fluid Dynamics, Uncertainty Quantification, Physics-Informed Learning

Abstrakt

Tato práce se zabývá využitím modelů strojového učení pro predikci aerodynamických vlastností v raných fázích vývoje vozidel. Cílem je nahradit nebo doplnit tradiční metody rychlejšími a výpočetně efektivnějšími alternativami.

Grafová neuronová síť MeshGraphNet je hodnocena jako výchozí model pro predikci aerodynamických veličin na povrchu vozidla. Její omezení jsou analyzována, včetně absence informace o globální geometrické struktuře vozidla a nedostatečného popisu lokálních geometrických vlastností. Pro řešení těchto nedostatků je navržena nová architektura s názvem LaplaceAeroNet, která vykazuje vyšší přesnost a robustnost. Model spolehlivě predikuje aerodynamické chování i pro geometrie vozidel obsahující prvky jako kola, spoilery nebo změny světlé výšky, které nebyly zahrnuty ve trénovacích datech. Pro odhad důvěry v predikce modelu je použita metoda kvantifikace nejistoty založená na ansámblu modelů.

Druhá část práce je volitelným rozšířením mimo hlavní zaměření a zabývá se predikcí třírozměrného aerodynamického proudění kolem vozidla. Navržený model FourierAeroNet kombinuje Fourierův neuronový operátor s doladěním pomocí U-Netu prostřednictvím architektury U-Fourier Neural Operator a dosahuje lepších výsledků než výchozí konvoluční model díky efektivnějšímu zachycení dlouhodobých struktur proudění.

Výsledky ukazují, že modely strojového učení mohou poskytovat přesné a efektivní náhradní predikce, které podporují aerodynamickou analýzu a rozhodování v raných fázích vývoje vozidla.

Klíčová slova Aerodynamika Vozidel, Náhradní Modelování, Grafové Neuronové Sítě, Fourierův Neuronový Operátor, Výpočetní Dynamika Tekutin, Kvantifikace Nejistoty, Fyzikálně-Informované Učení

Abbreviations

CFD Computational Fluid Dynamics

ML Machine Learning

SDF Signed Distance Function

PINN Physics-Informed Neural Network

CNN Convolutional Neural Network

FNO Fourier Neural Operator

U-FNO U-Fourier Neural Operator

GNN Graph Neural Network

UDA Unsupervised Data Augmentation

MAE Mean Absolute Error

MSE Mean Squared Error

MaxAE Maximum Absolute Error

Contents

1 Introduction

This thesis investigates using the Machine Learning (ML) models to predict vehicle aerodynamic behavior. The introduction outlines the motivation, reviews existing approaches, and highlights their limitations, followed by a summary of the thesis contributions aimed at overcoming these challenges. Chapter 2 provides an aerodynamic background for the models presented in Chapters 3 and 4. While Chapter 3 is devoted to models selected and refined following the objectives defined in the thesis assignment, Chapter 4 presents additional models explored as part of an independent research initiative. Chapter 5 concludes the thesis by summarizing the main findings and assessing the effectiveness of the proposed solutions.

1.1 Motivation

Aerodynamics is a fundamental aspect of vehicle performance, directly influencing key characteristics such as top speed, fuel efficiency, road stability, and structural integrity. These factors are relevant in both motorsport and consumer vehicle design and are monitored by regulatory authorities due to their impact on emissions and environmental compliance [1, pp. 1–2].

Traditionally, aerodynamic performance is assessed using wind tunnel testing and Computational Fluid Dynamics (CFD) simulations. Wind tunnels provide highly accurate measurements of global aerodynamic forces acting on the entire vehicle, which are essential for evaluating overall vehicle performance. However, they offer limited capability to capture detailed aerodynamic flow behavior, which is important for optimizing local design features. In contrast, CFD enables high-resolution simulation of flow fields by solving the governing equations on digital models. Despite this, the accuracy of global force predictions from CFD is generally lower than that of wind tunnel measurements [2].

Figure 1.1: Wind tunnel testing using a hardware vehicle model enables high-fidelity measurement of global aerodynamic forces acting on the entire car, with low associated error. However, local flow structures can only be resolved with limited spatial resolution, and the method requires substantial time and financial investment. Composed by the author using DALL·E 3-generated elements.

Despite the high precision of traditional methods, both approaches are resource-intensive. Wind tunnel testing requires the construction of costly physical prototypes, access to specialized facilities, and considerable time investment for setup. CFD, while offering greater flexibility, demands extensive computational resources, expert meshing, and long simulation times, often ranging from several hours to multiple days per configuration [3]. These constraints limit the applicability of both methods in early-stage design, where rapid evaluation of various variants is essential.

Figure 1.2: CFD simulations based on a digital vehicle model allow for detailed prediction of local aerodynamic flow phenomena with high spatial resolution and low associated error. In contrast, estimates of global aerodynamic forces are generally less accurate. Moreover, the approach entails moderate to high computational demands and execution time. Composed by the author using DALL·E 3-generated elements.

ML methods have recently emerged as a promising alternative. Once properly trained, such models can predict aerodynamic parameters with acceptable accuracy at a significantly lower computational cost, enabling rapid concept evaluation and complementing traditional approaches.

Figure 1.3: Motivation for ML aerodynamic predictions: rapid estimation of global aerodynamic forces acting on the entire vehicle, as well as aerodynamic flow behavior, based on concept-level digital models. At this early stage, approximate predictions are sufficient to guide design decisions, reducing reliance on traditional evaluations that are costly and time-consuming. Composed by the author using DALL·E 3-generated elements.

Considering the critical importance of aerodynamics in vehicle performance, the ability to rapidly evaluate multiple design configurations during early development stages is essential. However, the resource-intensive nature of traditional methods limits their feasibility for iterative exploration. This work proposes ML techniques to enable fast aerodynamic prediction with sufficient accuracy for engineering applications.

1.2 Related works

For simulation purposes, digital or conceptual vehicle models typically represent the car geometry as a surface mesh [1, pp. 841–843], a collection of discrete points (nodes) connected by edges and faces to approximate the body's surface. An example of such a mesh structure for a car model is shown in Figure 1.4.

A detailed review of existing ML approaches operating on mesh-based vehicle representations is provided in the following subsections. These approaches can be broadly categorized into three groups: (i) models predicting the aerodynamic flow, (ii) models predicting surface aerodynamic quantities derived from the flow, and (iii) models predicting aerodynamic forces, with machine learning approaches typically targeting normalized aerodynamic coefficients, which represent the corresponding forces in a scalable and consistent form.

Figure 1.4: Example of a surface mesh representation of a vehicle, illustrating nodes, edges, and faces used to approximate the body geometry. Adapted from [4] with modifications.

Models predicting the aerodynamic flow

Figure 1.5: Overview of ML methods used for predicting aerodynamic flow around a car geometry. A structured voxel grid is generated around the vehicle, where aerodynamic quantities are predicted for each voxel. Composed by the author using DALL·E 3-generated elements.

A common strategy for predicting the aerodynamic flow field around a vehicle is to first construct a regular voxel grid surrounding the surface mesh, where each voxel is encoded with the Signed Distance Function (SDF) representing the distance to the nearest face of the vehicle mesh [5]. The voxel grid can then either be sliced into 2D planes and processed individually using a 2D U-Net, as proposed in [6], or fed directly into a 3D U-Net, as in XaeroNet [7]. In both cases, the U-Net processes the SDF representation to extract geometric features and predict the corresponding aerodynamic flow fields, with XaeroNet further employing a physics-informed loss, following the principles of Physics-Informed Neural Networks (PINNs), to enforce consistency with the governing equations of fluid dynamics, thereby promoting better generalization to unseen flow conditions.

In addition to using the SDF to represent vehicle geometry, Neural Operator architectures have been developed to learn adaptive projections of the surface mesh onto structured voxel grids. For instance, a Decomposable Multi-scale Iterative Neural Operator (DoMINO) [8] samples a point cloud from the vehicle's surface mesh, treating each mesh node as a point. This point cloud is then projected onto a computational voxel grid using point convolution operations based on nearest-neighbor searches. Specifically, for each voxel in the grid, the nearest points on the vehicle surface are identified, and their features are aggregated through a learnable function. Standard Convolutional Neural Networks (CNNs) operations, such as convolution and pooling, are subsequently applied to the computational grid to process the aggregated features. The predicted flow field is obtained by applying a similar nearest-neighbor aggregation procedure, where for each voxel in the predicted domain the nearest points in the computational grid are identified and aggregated, allowing the model to estimate the final aerodynamic quantities.

In general aerodynamic modeling, Fourier Neural Operators (FNOs) [9] have recently emerged as a strong alternative to traditional U-Nets for flow field prediction tasks. The primary motivation for this shift is that FNOs captures global dependencies more efficiently through spectral methods, leading to improved generalization. For example, FNOs have been successfully applied to the simulation of three-dimensional turbulence [10]. In addition, several extensions of the original FNO framework have been developed for flow modeling tasks. One such variant is the U-Fourier Neural Operator (U-FNO), which integrates the U-Net architecture into the spectral processing pipeline. This combination allows the model to simultaneously leverage the global receptive field of Fourier layers and the strong local feature extraction capabilities of U-Nets. The U-FNO has been successfully applied to turbulent channel flow prediction [11], achieving superior performance compared to both 3D U-Nets and standard FNOs. However, to the best of our knowledge, Fourier-based operators have not yet been applied to aerodynamic flow prediction around vehicles.

In summary, a common baseline for aerodynamic flow prediction around vehicles involves representing the geometry using the SDF and applying U-Net architectures to predict flow fields. More recently, due to the inherent geometric complexity of vehicle shapes, Neural Operator–based methods have gained popularity, offering a scalable and efficient framework for complementing the SDF representations by capturing both large-scale and fine-scale geometric features. In parallel, to better capture global dependencies during feature processing in flow modeling tasks, FNOs have been proposed as an alternative to conventional convolutional architectures.

- Models predicting the surface aerodynamic quantities

Figure 1.6: Overview of ML methods used for predicting aerodynamic coefficients of a car. A model predicts aerodynamic quantities at each mesh node (in this illustration, the values are further interpolated across the surface to produce a smooth distribution). Composed by the author using DALL ·E 3-generated elements.

Graph Neural Networks (GNNs), where the vehicle surface mesh is naturally interpreted as a graph by treating mesh nodes as vertices and mesh connectivity as edges, represent a widely adopted class of models for predicting surface aerodynamic quantities. In this formulation, the prediction of surface quantities is performed at the level of individual graph nodes, each corresponding to a vertex of the original mesh. A representative example is MeshGraphNet [12], which follows the classical encoder{processor{decoder architecture, where the processor stage facilitates local information exchange between neighboring mesh nodes. In its original formulation, MeshGraphNet performs predictions at each node based on its 3D spatial coordinates. This limited input feature space restricts the model's ability to learn high-frequency patterns, especially in regions with rapidly varying geometry. Additionally, the model struggles to capture global geometric dependencies and to maintain scalability for high-resolution meshes due to substantial memory requirements during message passing.

X-MeshGraphNet [13] addresses these limitations by introducing a graph partitioning algorithm that decomposes large surface meshes into smaller, overlapping subgraphs with halo regions to enable efficient inter-subgraph communication. In addition, it incorporates multi-scale message-passing mechanisms to improve the modeling of global geometric dependencies. Crucially, the input features are extended beyond node coordinates to include geometric descriptors such as surface normals, which enrich the representation of local shape and allow the model to more effectively capture high-frequency aerodynamic patterns.

Another approach to addressing the challenge of capturing global geometric dependencies is a Graph-Transformer Operator (AeroGTO) [14], which extends conventional graph-based architectures by incorporating a graph-transformer design. In this model, global self-attention mechanisms are applied in addition to localized information exchange between neighboring nodes, enabling accurate prediction of scalar surface quantities by capturing both short- and long-range interactions across the vehicle geometry.

Alternatively, global geometric information can be captured by projecting surface meshes onto structured voxel grids using Neural Operator architectures, which are capable of automatically extracting relevant features from geometric input representations. The previously mentioned Decomposable Multi-scale Iterative Neural Operator (DoMINO) also provides a mode for predicting surface quantities by applying a nearest-neighbor aggregation procedure, where neighbors are searched not for each voxel in the predicted flow field, but for each node on the predicted surface mesh. Similarly, the Geometry-Informed Neural Operator (GINO) [15] projects the surface geometry onto a voxel grid using a point convolution operator, but instead of applying convolutional layers as in DoMINO, it employs a FNOs to learn latent feature representations across the voxel grid. The final surface quantities are then recovered by mapping features back from the voxel grid to the original surface mesh using a nearest-neighbor projection.

Overall, most existing models rely on three main approaches: GNNs, transformer-based architectures, and Neural Operators. The most fundamental baseline is to employ GNNs, which naturally model local interactions by operating on mesh connectivity, but face significant challenges in capturing global geometric dependencies, maintaining scalability to high-resolution meshes, and effectively representing local geometric shape. For relatively small graphs, transformer-based extensions, such as graph-transformers, have been introduced to enhance global feature aggregation through self-attention mechanisms. However, for large-scale graphs arising from high-resolution vehicle meshes, Neural Operator{based methods are increasingly preferred, as they enable efficient projection of complex surface geometries onto structured voxel grids, thereby improving scalability and capturing both local and global aerodynamic features more effectively. PINNs are generally not employed in this context, as aerodynamic governing equations describe volumetric flow behavior rather than quantities restricted to the vehicle surface. Consequently, the lack of direct physical constraints on surface quantities, combined with the limited size of available datasets resulting from the high computational cost of CFD simulations, often constrains the generalization ability of models in this domain. Moreover, standard data augmentation methods are also limited, as modifications may induce inconsistencies in the surface aerodynamic behavior.

Aerodynamic coefficients prediction

Figure 1.7: Overview of ML methods used for predicting aerodynamic forces of a car, which are typically expressed in terms of normalized coefficients, based on the input geometry. Composed by the author using DALL·E 3-generated elements.

Typically, models developed for surface aerodynamic quantity prediction or flow field prediction are also extended to predict aerodynamic coefficients. This can be achieved by adding a dedicated learnable branch for coefficient regression, as exemplified by the Factorized Implicit Global Convolution (FIGC) model [16]. Alternatively, aerodynamic coefficients can be computed mathematically based on the predicted flow fields or surface distributions.

In addition to these multitask or post-processing strategies, there are also models designed specifically for direct aerodynamic coefficient prediction. A recent example is the DrivAer Transformer [17], which integrates a transformer mechanism within a graph-based neural network framework to model both local and global interactions among mesh nodes for accurate coefficient estimation.

Thus, aerodynamic coefficient prediction is either based on models originally developed for surface quantity and flow field prediction or performed using networks specifically designed for direct coefficient estimation, where the architectural trends closely mirror those observed in surface and flow prediction tasks.

1.3 Contribution

This work proposes several improvements to enhance the accuracy and generalization of machine learning models for aerodynamic flow and surface quantity prediction tasks around vehicles, and demonstrates the application of an uncertainty estimation method in this context. The contributions are as follows:

- Laplacian of SDF: Augmenting the SDF representation with its Laplacian to provide a richer description of local surface curvature.
- Application of U-FNO: Extending the U-FNO architecture to the task of aerodynamic flow prediction around vehicles, leveraging its spectral properties to better capture global dependencies in complex flow fields.

- Laplacian differential coordinates: Enhancing the local encoding of surface geometry by supplementing raw node positions with Laplacian differential coordinates, thereby providing information about local shape variations on the vehicle mesh.
- Unsupervised Data Augmentation (UDA): Applying a consistency-based regularization technique to improve generalization in surface aerodynamic quantity prediction tasks.
- Application of ensemble-based uncertainty prediction: Incorporating an ensemble-based method to estimate predictive uncertainty in aerodynamic surface quantity prediction, improving reliability in model outputs.