

**AUTOMOTIVE AERODYNAMIC DESIGN
USING SENSITIVITY-BASED METHODS**

**CONCEPTION AÉRODYNAMIQUE
AUTOMOBILE À L'AIDE DE MÉTHODES
BASÉES SUR LE GRADIENT**

A Thesis Submitted to the Division of Graduate Studies
of the Royal Military College of Canada
by

Maurice Noah Nayman, B.A.Sc.

In Partial Fulfillment of the Requirements for the Degree of
Master of Applied Science in Aeronautical Engineering

April, 2024

© This thesis may be used within the Department of National Defence but copyright
for open publication remains the property of the author.

You'll have plenty of CFD to do. Let's see if you still like it in a year...

To my sister, I miss you every day.

Acknowledgements

I would like to thank Dr. Ruben Perez, for his support throughout my six years in his research group. He is a supervisor, a mentor, and a friend. He saw my potential as a first year student, and helped shape my academic career. He always encouraged my research pursuits, and challenged me to achieve more. His support of my personal interests and how they could be incorporated into my research, have helped form me into the engineer I am today.

I would also like to thank Elio Xhemalaj, Lucas Heathcote and the entire Multimatic CFD team. Your constant support through all my questions helped me greatly throughout my research. Your cumulative industry experience provided me with so many tools to shape the course of my thesis and my understanding of automotive aerodynamics.

Additionally, I would like to thank Dr. Mohsen Ferchichi and Dr. Jean-Marc Noël. Your support helped ground me as a person, and allowed me to realize my academic potential.

Last but not least I must thank my family, whose unconditional support of my education helped shape my scholastic career. Any question I could ask was met with a yes. If there was any way you could build me up, you were the first to be there. I am eternally grateful for you.

Abstract

Nayman, Maurice Noah. M.A.Sc. Royal Military College of Canada, March 2024. *Automotive Aerodynamic Design Using Sensitivity-Based Methods*. Supervised by Ruben E. Perez, B.Eng., M.A.Sc., Ph.D., P.Eng., Professor.

With the shift towards electrification in the automotive industry, aerodynamic drag has a larger influence on a car's range as compared to internal combustion vehicles. Automotive aerodynamic design is largely aesthetic driven, but manufacturers are beginning to shift towards more drag-friendly concepts, driven by computational fluid dynamics simulations. With advances in computational fluid dynamics-driven optimization techniques, the continuous adjoint method has shown the ability to be a computationally efficient method of calculating design sensitivities. This research explores a two-pronged approach for applying the continuous adjoint method to automotive aerodynamic design, using the canonical *DrivAer Estateback*. This is achieved using computational fluid dynamics models validated against existing experimental data. The first method is using the continuous adjoint to define a Reynolds-Averaged Navier Stokes (RANS)-based Momentum Contribution Field, which informs the user how the momentum around a car contributes to its drag. Additionally, improvements to this RANS-based methodology are provided through Delayed Detached Eddy Simulations (DDES). The RANS Momentum Contribution Field information is then used to guide design modifications. From these modifications, an *improved* variant of the *Estateback* model was produced, with 8.8% less drag than the baseline model. The second approach makes use of unconstrained gradient-based optimization of the *base* and *improved Estateback* models, for multiple areas on the cars, using multiple optimization algorithms. Relative improvements in drag up to 7% were predicted with the optimizer, with limitations in optimizer efficacy noted in areas dominated by adverse pressure gradients. DDES post-optimality simulations confirmed the efficacy of the RANS-based optimizer in areas with weaker adverse pressure gradients, as well as the inaccuracy of the optimizer in areas of high separation. This two-pronged approach led to a unified design methodology, where expert-driven, followed by computer-driven continuous adjoint methods, can be used to yield drag reductions in automotive applications.

Keywords: automotive aerodynamics, aerodynamic shape optimization, continuous adjoints, computational fluid dynamics, open-source, applied aerodynamics

Résumé

Nayman, Maurice Noah. M.A.Sc. Collège militaire royal du Canada, Mai 2024. *Conception Aérodynamique Automobile à l'Aide de Méthodes Basées sur le Gradient*. Supervisé par Ruben E. Perez, B.Eng., M.A.Sc., Ph.D., P.Eng., Professeur.

Avec le passage à l'électrification dans l'industrie automobile, la traînée aérodynamique a une plus grande influence sur l'autonomie d'une voiture que sur les véhicules à combustion interne. La conception aérodynamique automobile est en grande partie axée sur l'esthétique, mais les manufacturiers commencent à s'orienter vers des concepts plus respectueux de la traînée, pilotés par des simulations numériques de dynamique des fluides. Avec les progrès des techniques d'optimisation numérique basées sur la dynamique des fluides, la méthode des adjoints continus a montré qu'elle pouvait être une méthode efficace de calculer des sensibilités de conception. Cette recherche explore une approche à deux volets pour l'application de la méthode de l'adjoint continu à la conception aérodynamique automobile, en utilisant le dos de domaine canonique *DrivAer*. Ceci est réalisé à l'aide de modèles numériques de dynamique des fluides validés par rapport aux données expérimentales existantes. La première méthode consiste à utiliser l'adjoint continu pour définir un champ de contribution de Navier Stokes basé sur la moyenne de Reynolds (RANS), qui informe l'utilisateur de la façon dont la vitesse autour d'une voiture contribue à sa traînée. De plus, des améliorations à cette méthodologie basée sur RANS sont apportées par le biais de simulations de tourbillons détachés retardés (DDES). Les informations du champ de contribution à la vitesse RANS sont ensuite utilisées pour guider les modifications de conception. À partir de ces modifications, une variante améliorée du modèle *Estateback* a été produite, avec 8,8 % de traînée moins que le modèle *baseline*. La deuxième approche fait appel à l'optimisation sans contrainte basée sur le gradient du modèle *baseline* et à des modèles de succession améliorés, pour plusieurs zones des voitures, à l'aide de plusieurs algorithmes d'optimisation. Traînée réduite de 7 % on été prédite, avec des limitations de le rendement de l'optimiseur notées dans les zones dominées par des gradients de pression défavorables. Les simulations post-optimales du DDES ont confirmé que le rendement de l'optimiseur basé sur RANS sont plus faibles dans les zones où les gradients de pression défavorable, ainsi que l'imprécision de l'optimiseur dans les zones de forte décollement. Cette approche à deux volets est rendu à une méthodologie de conception unifiée, où des méthodes d'adjoint continu pilotées par des experts, suivies d'une méthode d'adjoint continu pilotée par ordinateur, peuvent être utilisées pour réduire la traînée dans les applications automobiles.

Mots-clés: aérodynamique automobile, optimisation de la forme aérodynamique, adjoints continus, simulations numériques de dynamique des fluides, logiciels libres, aérodynamique appliquée

Contents

Acknowledgements	iii
Abstract	iv
Résumé	v
List of Tables	ix
List of Figures	x
Nomenclature	xv
1 Introduction	1
1.1 Research Objectives	2
1.2 Thesis Layout	3
1.3 Contributions	3
2 Literature Review	5
2.1 Automotive Aerodynamic Shape Optimization	5
2.1.1 <i>DrivAer</i> Geometry	8
2.2 Continuous Adjoints	10
2.2.1 Adjoint Turbulence Modelling	10
2.3 Concluding Remarks on the State-of-the-Art	13
3 CFD Model Validation Using the <i>DrivAer</i> Geometry	14
3.1 Canonical Automotive Geometry - <i>DrivAer</i>	14
3.2 Meshing	17
3.3 Reynolds-Averaged Navier Stokes Simulations	20
3.3.1 Spalart Allmaras Model	21
3.3.2 $k - \omega$ SST Model	22
3.4 Delayed Detached Eddy Simulations	24
3.4.1 $k - \omega$ SST Detached Eddy Simulation	24
3.4.2 $k - \omega$ SST Delayed Detached Eddy Simulation	25
3.4.3 Grid Refinement in DDES Models	26
3.5 DDES Mesh and Timestep Dependency Studies	27
3.6 DDES Validation Against Experimental Data	30
3.7 Comparison of DDES and RANS Simulations	33
4 Momentum Contribution Field Analysis and Design Modifications	41

4.1	Continuous Adjoint Model	41
4.2	Momentum Contribution Field	43
4.3	DDES-Based Improvements to Momentum Contribution Field Predictions	44
4.4	RANS MCF Evaluation	45
4.5	DDES-Based Improvements to the MCF	47
4.6	RANS MCF-Inspired Design Modifications	51
4.6.1	Front Splitter Modifications	51
4.6.2	Front Bumper Cheek Modifications	53
4.6.3	Rear Fender Modifications	56
4.6.4	Roof Spoiler Modifications	59
4.6.5	Final Design	62
4.7	Limitations of the RANS MCF	64
4.7.1	Hood Modifications	64
4.7.2	Diffuser Modifications	65
5	Drag Shape Optimization of the <i>DrivAer Estateback</i>	68
5.1	Half-Car Models	68
5.1.1	Drag Comparisons	69
5.1.2	Surface Pressure Drag Comparisons	69
5.1.3	Half-Car Modelling Conclusions	71
5.2	Gradient-Based Optimizers	71
5.2.1	Conjugate Gradient	72
5.2.2	BFGS	73
5.2.3	DBFGS	74
5.3	Freeform-Deformation Boxes	75
5.4	Numerical Setup	76
5.5	Optimization Problem	77
5.6	Roof Spoiler Optimization	80
5.6.1	<i>Base</i> Variant Results	80
5.6.2	<i>Improved</i> Variant Results	82
5.6.3	Comparing Optimization Results	83
5.7	Roof Spoiler Post-Optimality CFD	87
5.8	Front Bumper Cheek Optimization	90
5.8.1	<i>Base</i> Variant Results	90
5.8.2	<i>Improved</i> Variant Results	91
5.8.3	Comparing Optimization Results	92
5.9	Front Bumper Cheek Post-Optimality CFD	95
6	Conclusions and Recommendations	98
6.1	Conclusions	98
6.2	Recommendations for Future Work	100
	Bibliography	101
	Appendices	105
A	Supplemental Mesh Dependency Study Images	106
A.1	Sample <i>blockMesh</i> Output	106

A.2	<i>DrivAer Estateback</i> Surface Meshes	106
A.3	<i>DrivAer Estateback</i> Volume Meshes	109
B	Additional Post-Processing of <i>DrivAer Estateback</i> Variants	112
B.1	Wake Structures	112
B.2	Surface Pressure Coefficient	114
C	Additional Post-Processing of <i>DrivAer Estateback</i> Modifications	117
C.1	Additional Splitter Post-Processing	117
C.2	Additional Front Bumper Cheek Post-Processing	118
C.3	Additional Rear Fender Post-Processing	121
C.4	Additional Roof Spoiler Post-Processing	123
D	Mesh Deformation Images for the Spoiler Optimization	126
D.1	Conjugate Gradient Optimization	126
D.2	BFGS Optimization	130
D.3	DBFGS Optimization	133
E	Mesh Deformation Images for the Front Bumper Cheek Optimization	135
E.1	Conjugate Gradient Optimization	135
E.2	BFGS Optimization	136
E.3	DBFGS Optimization	138
F	Additional Post-Processing of Post-Optimality CFD	141
F.1	Additional Roof Spoiler Images	141
F.2	Additional Front Bumper Cheek Images	142

List of Tables

3.1	<i>DrivAer</i> configurations used in this study	16
3.2	<i>blockMesh</i> divisions used to generate E_D_wM_wW meshes	29
4.1	Splitter designs and their drag contribution	53
4.2	Front bumper cheek flick designs and their drag contribution	55
4.3	Rear fender infill designs and their drag contribution	58
4.4	Summary of rear spoiler modifications and their performance	61
5.1	Comparison of drag predictions for both <i>Estateback</i> variants for full-car and half-car simulation approaches	69
5.2	Number of points and degree of the NURBS radial basis function in the u, v, w directions for each FFD	79
5.3	Summary of optimization progress for each algorithm and variant studied, for the spoiler FFD	85
5.4	Summary of optimization progress for each algorithm and variant studied, for the front bumper cheek FFD box	94
A.1	Summary of surface cell level lengths in mm for each of the mesh dependency meshes	107

List of Figures

1.1	Existing emission trends across various transportation industries, as well as future projections [2].	1
2.1	Parametric deformation box defined and optimized surface results for a low-drag concept vehicle [11].	6
2.2	Differences in RANS <i>vs.</i> approximate DES surface sensitivities in the rear spoiler area, where red indicates moving the surface inwards to reduce drag and blue indicates moving the surface outwards to reduce drag [14].	7
2.3	<i>Fastback</i> (F), <i>Estateback</i> (E), and <i>Notchback</i> (N) configurations of the <i>DrivAer</i> model [18].	8
2.4	Existing canonical models, which did not feature as many complex flow structures and geometry as the <i>DrivAer</i> model	9
2.5	Examination of sensitivity derivatives for each control variable for flow over a NACA 0012 wing, with control point parameterization shown. Sensitivity calculations include finite differences, frozen turbulence adjoints, and adjoint $k - \omega$ SST models [29].	12
3.1	<i>Fastback</i> (F), <i>Estateback</i> (E), and <i>Notchback</i> (N) configurations of the <i>DrivAer</i> model [18].	15
3.2	<i>DrivAer Estateback</i> configurations used for model validation	16
3.3	Naming conventions used for key components of the <i>Estateback</i>	17
3.4	Full flow chart, with supporting dictionary files and utilities to run <i>snappyHexMesh</i> on an input triangulated geometry file: geometry.stl [32]	18
3.5	Sample two-dimensional meshing problem using <i>snappyHexMesh</i> [33]	19
3.6	Sample meshes of the <i>DrivAer Estateback</i> generated using <i>snappyHexMesh</i>	20
3.7	Effects of turbulence model and mesh resolution on DES flow field, a) Coarse SA-DES, b) Fine SA-DES, c) Fine SST-DES. Adapted from [43]	26
3.8	Boundary conditions for farfield and car surfaces	28
3.9	Drag coefficient as a function of cell count and time step	29
3.10	<i>DrivAer</i> configuration drag coefficient as a function of timestep	30
3.11	Drag coefficient comparisons between CFD and experimental results for the four tested <i>Estateback</i> configurations	31
3.12	Comparison of pressure coefficient between CFD and experimental results along the top of the car in the $y = 0$ mm plane. Absolute residuals between CFD and wind tunnel results in the right image	32
3.13	Comparison of pressure coefficient between CFD and experimental results along the bottom of the car in the $y = 0$ mm plane. Absolute residuals between CFD and wind tunnel results in the right image	32

3.14	Comparison of pressure coefficient between CFD and experimental results along the left side of the car in the $z = 150$ mm plane. Absolute residuals between CFD and wind tunnel results in the right image	33
3.15	Comparison of pressure coefficient between SST-DDES modeling and SA-RANS modeling on the front of the car	34
3.16	Comparison of pressure coefficient between SST-DDES modeling and SA-RANS modeling on the rear of the car	35
3.17	Comparison of pressure coefficient between SST-DDES modeling and SA-RANS modeling on the underside of the car	36
3.18	Normalized velocity contours on the $y = 0$ mm plane for the two primal solution methods	38
3.19	Normalized velocity contours on the $y = -700$ mm plane for the two primal solution methods	38
3.20	Normalized velocity contours on the $z = -90$ mm plane for the two primal solution methods	39
3.21	Normalized velocity contours on the $z = -15$ mm plane for the two primal solution methods	40
4.1	Drag MCF on the $y = 0$ mm plane. The top right view is zoomed in on the front fascia and splitter and the bottom right view is zoomed in on the rear fascia and diffuser	45
4.2	MCF for the drag of the car on the $z = -115$ mm plane	46
4.3	MCF for the drag of the car on the $z = -15$ mm plane	47
4.4	Momentum Contribution Field contours on the $y = 0$ mm plane for the two primal solution methods	48
4.5	Momentum Contribution Field contours on the $y = -700$ mm plane for the two primal solution methods	49
4.6	Momentum Contribution Field contours on the $z = -90$ mm plane for the two primal solution methods	50
4.7	Momentum Contribution Field contours on the $z = -15$ mm plane for the two primal solution methods	50
4.8	Splitter profiles, baseline and the three options	52
4.9	Comparison between the baseline splitter and Splitter Option 3 of pressure coefficient on the floor	53
4.10	Front bumper cheek geometry, baseline and three options	54
4.11	Comparison of pressure drag coefficient around the front tires for the baseline and Option 2 front bumper cheeks	56
4.12	Rear fender geometry, baseline and four infill options	57
4.13	Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 2	58
4.14	Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 2	59
4.15	Baseline roof spoiler geometry, along with the four proposed trailing edge modifications and the spoiler profile on car centreline (Spoiler Options 1 and 2 share the same profile)	60
4.16	Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 2	61
4.17	Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 2	62

4.18	Front isometric comparison of pressure drag coefficient between the baseline and <i>improved Estateback</i> models	63
4.19	Rear view comparison of pressure drag coefficient between the baseline and <i>improved Estateback</i> models	63
4.20	Hood geometry, baseline with the three trailing edge flick geometries tested	64
4.21	Pressure coefficient around the windshield for the baseline hood trailing edge, and the 3 mm tall trailing edge flick	65
4.22	Diffuser geometry, baseline with the three modifications to the diffuser profile tested	66
4.23	Normalized velocity contours on the $y = 0$ mm plane for the baseline and Diffuser Iteration 3	67
5.1	Surface pressure drag coefficient on the full-car and half-car (mirrored for clarity) <i>Estateback</i>	70
5.2	Example surface displacement that can be achieved with an FFD, from [54]	75
5.3	Side view of a non-global axis aligned FFD. The y/v directions are shared in this example, and pass into the page	76
5.4	Freeform deformation box placed around the rear spoiler. The car is shown in its symmetric state to highlight the FFD points on the symmetry plane	78
5.5	Freeform deformation box placed around the left front bumper cheek	78
5.6	Optimization results for the three optimization algorithms. The objective function is normalized to the drag of the <i>base</i> variant at the initial optimization iteration	80
5.7	Optimization results for the 10-line search iteration BFGS case, compared to the initial CG and BFGS cases	81
5.8	Optimization results for each of the three optimization algorithms	82
5.9	Optimization results for each of the three optimization algorithms, for both <i>Estateback</i> variants. Both variants' drag is normalized by the drag of the <i>base</i> variant at the initial optimization iteration ($J_{0,Base}$)	84
5.10	Deformed surfaces, subject to the spoiler FFD	86
5.11	Percent difference in J/J_0 from post-optimality CFD to optimizer predictions, for the spoiler design vectors chosen	88
5.12	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 9 design point of the <i>base</i> variant	89
5.13	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 6 design point of the <i>base</i> variant	89
5.14	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 12 design point of the <i>improved</i> variant	90
5.15	Optimization results for each of the three optimization algorithms. The objective function is normalized to the drag of the <i>base</i> variant at the initial optimization iteration	91
5.16	Optimization results for each of the three optimization algorithms	92
5.17	Optimization results for each of the three optimization algorithms, for both <i>Estateback</i> variants. Both variants' drag is normalized by the drag of the <i>base</i> variant at the initial optimization iteration ($J_{0,Base}$)	93
5.18	Deformed surfaces, subject to the front bumper cheek FFD	95
5.19	Percent difference in J/J_0 from post-optimality CFD to optimizer predictions, for the front bumper cheek design vectors chosen	96

5.20	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 3 design point of the <i>base</i> variant	97
5.21	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the DBFGS, Iteration 3 design point of the <i>improved</i> variant	97
A.1	Background <i>blockMesh</i> used for the 84 x 28 x 14 <i>Very Coarse</i> mesh	106
A.2	Surface mesh samples for the different background meshes	109
A.3	Volume mesh images resulting from the 84 x 28 x 14 <i>Very Coarse</i> mesh	109
A.4	Volume mesh images resulting from the 90 x 30 x 15 <i>Coarse</i> mesh	110
A.5	Volume mesh images resulting from the 96 x 32 x 16 <i>Medium</i> mesh	110
A.6	Volume mesh images resulting from the 102 x 34 x 17 <i>Fine</i> mesh	110
A.7	Volume mesh images resulting from the 108 x 34 x 18 <i>Very Fine</i> mesh	111
B.1	Wake structures of the E_D_wM_wW configuration. Surface contoured by pressure coefficient	112
B.2	Wake structures of the E_D_woM_wW configuration. Surface contoured by pressure coefficient	113
B.3	Wake structures of the E_S_wM_wW configuration. Surface contoured by pressure coefficient	113
B.4	Wake structures of the E_S_woM_wW configuration. Surface contoured by pressure coefficient	113
B.5	Front and rear views of the surface pressure coefficient of the E_D_wM_wW configuration	114
B.6	Front and rear views of the surface pressure coefficient of the E_D_woM_wW configuration	114
B.7	Front and rear views of the surface pressure coefficient of the E_S_wM_wW configuration	115
B.8	Underbody surface pressure coefficient of the E_S_woM_wW configuration	115
B.9	Underbody surface pressure coefficient of the E_D_wM_wW configuration	115
B.10	Underbody surface pressure coefficient of the E_D_woM_wW configuration	116
B.11	Underbody surface pressure coefficient of the E_S_wM_wW configuration	116
B.12	Underbody surface pressure coefficient of the E_S_woM_wW configuration	116
C.1	Comparison between the baseline splitter and Splitter Option 1 of pressure coefficient on the floor	117
C.2	Comparison between the baseline splitter and Splitter Option 2 of pressure coefficient on the floor	118
C.3	Comparison of pressure drag coefficient around the front tires for the baseline and Option 1 front bumper cheeks	119
C.4	Comparison of pressure drag coefficient around the front tires for the baseline and Option 2 front bumper cheeks	120
C.5	Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 1	121
C.6	Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 3	121
C.7	Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 4	122
C.8	Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 1	122
C.9	Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 3	122

C.10	Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 4	123
C.11	Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 1	123
C.12	Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 3	124
C.13	Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 4	124
C.14	Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 1	124
C.15	Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 3	125
C.16	Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 4	125
D.1	<i>Base</i> variant Conjugate Gradient spoiler deformation progression	128
D.2	<i>Improved</i> variant Conjugate Gradient spoiler deformation progression	129
D.3	<i>Base</i> variant BFGS spoiler deformation progression	131
D.4	<i>Improved</i> variant BFGS spoiler deformation progression	133
D.5	<i>Base</i> variant DBFGS spoiler deformation progression	134
D.6	<i>Improved</i> variant DBFGS spoiler deformation progression	134
E.1	<i>Base</i> variant Conjugate Gradient front bumper cheek deformation progression	136
E.2	<i>Improved</i> variant Conjugate Gradient front bumper cheek deformation progression	136
E.3	<i>Base</i> variant BFGS front bumper cheek deformation progression	138
E.4	<i>Improved</i> variant BFGS front bumper cheek deformation progression	138
E.5	<i>Base</i> variant DBFGS front bumper cheek deformation progression	139
E.6	<i>Improved</i> variant DBFGS front bumper cheek deformation progression	140
F.1	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 4 design point of the <i>base</i> variant	141
F.2	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 9 design point of the <i>improved</i> variant	142
F.3	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 3 design point of the <i>base</i> variant	142
F.4	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 5 design point of the <i>base</i> variant	143
F.5	Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 3 design point of the <i>improved</i> variant	143

Nomenclature

Roman Symbols

\mathbf{g}	Gravity vector field [m^2/s^2]
\mathbf{p}	Search direction
\mathbf{x}^*	Design vector at which the objective function is minimized
\mathbf{s}	Change in design vector when approximating the inverse hessian matrix
\mathbf{y}	Change in gradient vector when approximating the inverse hessian matrix
∇	Gradient of a function
A	Frontal Area [m^2]
a_1	k- ω SST model coefficient
b_1	k- ω SST model coefficient
$C_{\nu 1}$	Spalart Allmaras model coefficient
C_{b1}	Spalart Allmaras model coefficient
C_{b2}	Spalart Allmaras model coefficient
C_{w1}	Spalart Allmaras model coefficient
C_{w2}	Spalart Allmaras model coefficient
C_{w3}	Spalart Allmaras model coefficient
D	Effective diffusivity of some field
F_1	k- ω SST cross-diffusion function
G	Specific dissipation rate production due to mean velocity shear [s^{-1}]
$H(\mathbf{x})$	Hessian matrix
J	Adjoint objective function
k	Equivalent total number of CFD solutions
l	Car Wheelbase [m]

p	Primal kinematic pressure field [m^2/s^2]
Q	Conjugate Gradient Hessian matrix
U	Primal velocity field [m/s]
u	First cartesian direction of a freeform deformation box
v	Second cartesian direction of a freeform deformation box
$V(\mathbf{x})$	Approximate inverse hessian matrix
w	Third cartesian direction of a freeform deformation box
x	Streamwise location [m]
y	Spanwise location [m]
y^+	Normalized wall-normal distance
z	Vertical location [m]
D	Aerodynamic drag [N]
k	Turbulent kinetic energy [m^2/s^2]
L	Aerodynamic lift [N]
S	User-defined source term

Greek Symbols

α	Search direction step length
β	Conjugate Gradient search direction correction term
β^*	$k - \omega$ SST model coefficient
χ	Ratio of transformed Spalart Allmaras eddy viscosity to molecular kinematic viscosity
ϵ	Optimization tolerance
γ	Damped BFGS curvature correction term
κ	Von Karman constant
μ	Molecular dynamic viscosity [Pa s]
μ_1	<i>Armijo</i> Conditions scaling factor
ν	Molecular kinematic viscosity [m^2/s]
ω	Turbulent specific dissipation rate [s^{-1}]
ϕ	Wind tunnel blockage ratio
ρ	Fluid density [kg/m^3]

σ_{ν_t}	Spalart Allmaras model coefficient
τ	Primal wall shear stress [m^2/s^2]
θ	Damped BFGS damping term
$\tilde{\nu}$	Transformed Spalart Allmaras eddy viscosity [m^2/s]

Subscripts

∞	Freestream value of variable
a	Adjoint variant of a variable or field
g	Gradient
i	x component of a field
j	y component of a field
K	Value at major optimization iteration, K
k	z component of a field
t	Turbulent variant of variable or field
abs	Absolute
dev	Deviatoric component of a tensor
eff	Effective
rel	Relative

Superscripts

T	Transpose of the vector/matrix
-----	--------------------------------

Dimensionless Groups

C_D	Drag Coefficient
C_L	Lift Coefficient
Re	Reynolds Number

Acronyms

BFGS	Broyden-Fletcher-Goldfarb-Shanno
CFD	Computational Fluid Dynamics
CG	Conjugate Gradient
DBFGS	Damped Broyden-Fletcher-Goldfarb-Shanno

DDES	Delayed Detached Eddy Simulations
DES	Detached Eddy Simulations
DNS	Direct Numerical Simulations
DOE	Design of Experiments
LES	Large Eddy Simulations
MCF	Momentum Contribution Field
MRF	Moving Reference Frame
NTUA	National Technical University of Athens
OEM	Original equipment manufacturer
PCOpt	Parallel CFD & Optimization Unit
PDE	Partial Differential Equation
PISO	Pressure-Implicit with Splitting of Operators
RANS	Reynolds Averaged Navier Stokes
SA	Spalart Allmaras turbulence model
SIMPLE	Semi-Implicit Method for Pressure Linked Equations
SST	Shear Stress Transport

1 Introduction

With the rise of global emissions causing an increase in global temperatures, the United Nations Framework Convention on Climate Change set a goal of limiting the average global temperature increase to 1.5°C at the Paris Climate Accords in 2015. The Intergovernmental Panel on Climate Change then released a report from their 2018 conference with the assessment that emissions would have to be reduced by 45% by 2030, and be net-zero by 2050 [1].

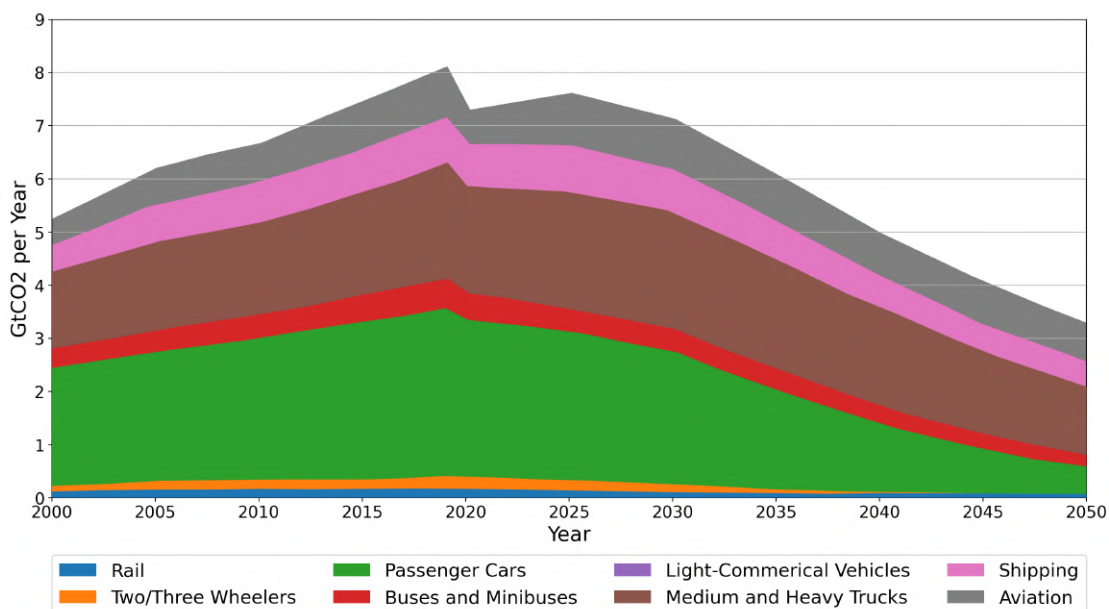


Figure 1.1: Existing emission trends across various transportation industries, as well as future projections [2].

As shown in Figure 1.1, passenger vehicles make up 41% of global emissions in the transportation sector, growing from approximately 2.2 billion metric tons of carbon dioxide in 2000, to 3.2 billion metric tons in 2019 [2]. As a result, many leading original equipment manufacturers (OEM's) have developed or are developing electric vehicles (EV) to reduce their fleet-averaged carbon dioxide emissions. This is being driven by new regulations and government requirements, such as the European Union's target for fleet-averaged carbon dioxide emissions of 95 g CO₂/km [3].

Current EV ranges are largely limited by battery technology. As of 2024, the longest range EV based on Environmental Protection Agency estimates is the *Lucid Air Dream Edition R* at 835 km. While this value rivals typical internal combustion figures, the average range

for the ten highest range EV, sits at approximately 575 km [4]. During the development of the *Porsche Taycan*, *Porsche* found that 29% of an EV total losses were due to aerodynamic drag, compared to 8% for traditional internal combustion vehicles [5]. As such, aerodynamics has become a more important factor in determining EV range, with manufacturers seeking to minimize drag. This has led to greater investments in aerodynamic simulation tools such as computational fluid dynamics (CFD), given its cost-advantage over wind tunnel testing, which can cost thousands of dollars per hour. With CFD methods with a focus on automotive applications being offered by more software vendors, more extensive optimization techniques are also becoming more prevalent. Leading software vendors such as *Engys*, *ANSYS*, *Siemens*, and more have all started offering optimization packages that can be paired with their CFD codes.

The nature of automotive design is largely aesthetic-driven, with each manufacturer trying to communicate a given design language, promoting brand recognition. This can be in direct competition with aerodynamic performance, especially for brands with traditionally sharp or angular car designs (*Lamborghini*, for instance). For internal combustion cars, this was less impactful, due to a smaller proportion of their losses from aerodynamic drag than that of EV. Therefore, manufacturers are beginning to shift their design language to be more in line with low aerodynamic drag concepts, such as more bluff and rounded bodies, removal of door handles, antennas, and more. Despite this, aerodynamics can still become compromised because of a desired styling theme.

In automotive design, aerodynamic shape optimization has been explored in CFD, with much of the literature focused on gradient-based methods, which will be explored in Chapter 2. Existing methods do not have a provision for styling-based constraints, as this can be largely subjective and unable to be analytically described. As such, traditional aerodynamic shape optimization can not always be implemented, as it could violate the styling theme of the car. It would therefore be more valuable to have a sensitivity method to extract information on the various design variables of the global car design, to be used in manual design routines. This would act as a supplement to traditional flow field information available from CFD computations, such as pressure and velocity fields. Optimization methods could then be used at later points in the design cycle, to improve local features of the car, extracting more performance from the platform.

1.1 Research Objectives

The aim of this research is to investigate the application of sensitivity approaches and optimization methods for improving automotive aerodynamic design practices. The first is the qualitative use of sensitivity information to allow for human-guided design modifications, that would be applicable to the early phases of a car's design cycle. The second is to apply traditional gradient-based optimization methods to a car geometry to evaluate the efficacy of different update methods within an automotive lens. The specific objectives are:

- Perform CFD analyses of a canonical automotive geometry to provide an understanding of the accuracy of CFD modelling methods, which will be subsequently applied to this research.
- Define a field that shows how local flow momentum contributes to a given objective function, called the Momentum Contribution Field (MCF). In this research, the MCF will

be computed using the continuous adjoint method to understand how flow momentum contributes to the aerodynamic drag of the canonical geometry.

- Using the results of the MCF, investigate design modifications on specific areas of the car, evaluated using the validated CFD models. Limitations of and improvements to the MCF method are also presented.
- Using the baseline and MCF-modified geometries, investigate the effect of different unconstrained gradient-based RANS optimizers with respect to aerodynamic drag, in key locations on the cars. This will characterize the efficacy of different algorithms, as well as the effect of the initial design vector on the results. Furthermore, this will highlight the efficacy of combined MCF/optimization approaches compared to gradient-based optimization alone.
- Perform higher-fidelity post-optimality CFD of key design vectors from the optimization models to evaluate the accuracy of the RANS optimizer. This will characterize regions in which the optimization fidelity is limited, as well as detail what magnitude of improvement is feasible from the optimizer.

1.2 Thesis Layout

This thesis is organized into six chapters. Following this Introduction, Chapter 1, Chapter 2 presents a literature review of the current state of automotive aerodynamic shape optimization research. This includes direct optimization of automotive geometries, methods pertaining to the evaluation and accuracy of continuous adjoint computations, and an exploration of the literature on the canonical automotive body studied, the *DrivAer* geometry. Chapter 3 presents CFD studies of the *DrivAer* geometry in various configurations, with mesh and time step dependency studies performed for a DDES approach. This is supplemented with an in-depth analysis of the simulation accuracy and differences between RANS and DDES results. Chapter 4 presents a novel approach using a proposed Momentum Contribution Field around the *DrivAer Estateback* to guide modifications shape modifications based on its information, to reduce drag. Limitations and improvements to the evaluated MCF methodology will also be discussed. Chapter 5 investigates the application of gradient-based optimization methods to the *DrivAer Estateback* base and improved variants. This will characterize the efficacy of different optimizers in regions with different flow physics, with post-optimality CFD presented to confirm the accuracy of optimization predictions. This chapter will also examine the effectiveness of combined MCF/optimization design methods. Finally, Chapter 6 summarizes key findings of this thesis, and provides recommendations for further investigations on this topic.

1.3 Contributions

The contributions of this thesis include:

- Validation of a DDES CFD approach for predicting the flow around largely separated automotive bodies.
- A thorough examination of an expert-guided methodology to develop aerodynamic design using a Momentum Contribution Field stemming from continuous adjoint solutions in a qualitative manner. This research investigates areas where the method is feasible and infeasible, along with proposed improvements to the sensitivity calculations that have not been previously presented.

- A study comparing the efficacy of different gradient-based optimizers when applied to a realistic automotive geometry, and a variant modified using the MCF method. Optimization is investigated for areas with small and large adverse pressure gradients, which will detail where optimization performance is most effective, and if prior expert-driven modifications affects this efficacy.
- Verification of the RANS optimizers' accuracy using higher-fidelity DDES CFD for key design vectors. This will verify the reliability of any predicted improvements in aerodynamic drag by the optimizers.
- A proposal for a unified, generic automotive aerodynamic design development map, employing both Momentum Contribution Fields and optimization methods to show how the drag of a car can be minimized more effectively when combining expert-driven design modifications with computer-driven optimization. This is based on the results of the MCF and optimization work.

2 Literature Review

This chapter examines the latest state of research surrounding automotive aerodynamic shape optimization. The first section covers current methodologies and results for automotive aerodynamic shape optimization. The second section summarizes existing research on continuous adjoints, as well as the state of adjoint turbulence modelling.

2.1 Automotive Aerodynamic Shape Optimization

Automotive design is largely styling driven, since each manufacturer has their own brand identity, conveyed through their products. As such, while aerodynamic performance is of concern, it can often be compromised by constraints from styling departments. As such, shape optimization is typically focused on the rear decklid/tailgate spoiler area and underfloor devices, since these pose the least impact on styling.

Current gradient-free automotive optimization methods largely focused on Design of Experiments (DOE) simulations, combined with response-surface methods, for optimization problems characterized by few design variables [6]. These methods tend to be useful in the automotive industry, since DOE solutions provide sensitivities to the design variables, which can then be communicated to the design studio for integration into an aesthetically-pleasing design. In certain studies, drag on a passenger hatchback was able to be reduced by approximately 10% using DOE and response surface methods, by modifying the front and rear fascias of the car, with the design variable ranges limited by styling desires [7]. These DOE/response-surface methods tend to be more costly than adjoint simulations, since a greater number of flow solutions are required to characterize parameter sensitivities. However, these methods provide more information on the effect of different geometric parameters, for use in later designs. In these DOE approaches, optimization fidelity is also limited by the number of design variables chosen, which also govern the computational cost of objective sensitivity calculations.

The use of gradient-based shape optimization routines in the aerospace industry has occurred since the 1980's [8]. These initial methods focused largely on simple airfoil optimization for potential flow solutions. In the 1990's, the practice evolved into the use of continuous adjoints for sensitivity derivative estimation for more complex geometries [9, 10]. These methods were compared against traditional finite-differences sensitivity derivative calculations to validate the accuracy of continuous adjoints. While the continuous adjoint method has existed for approximately three decades, it was not used widely in the automotive industry until the mid 2010's, largely led by the Parallel CFD & Optimization Unit (PCOpt) at the National Technical University of Athens [11, 12]. Additional work was carried out by Hare, where she investigated the applicability of adjoint-based shape optimization to the front wing of a mod-

ern Formula One car [13]. This work stands out as one of the few applications of the adjoint method in an effort to reduce the lift, or increase the downforce of a car.

PCOpt's initial work largely focused on rear decklid optimization using a steepest descent update method for directing surface deformation. When applied to a low-drag concept vehicle, drag was reduced by more than 2%, with corresponding reduction in lift of 30%. Drag improvements were realized in less than five equivalent flow solutions, with an equivalent flow solution being the time it took to generate a single, converged CFD model. This highlights that objective function minimization is possible for a relatively low equivalent cost. The parameterized deformation box and resultant optimized shape for the model are shown below in Figure 2.1.

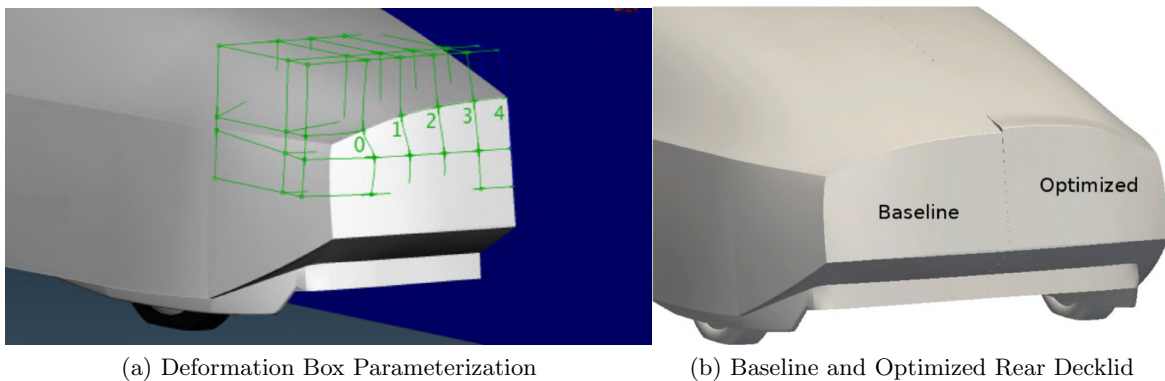


Figure 2.1: Parametric deformation box defined and optimized surface results for a low-drag concept vehicle [11].

Further use of the adjoint method employed pseudo-transient primal solutions to provide more accurate adjoint computations [11]. This used time-averaged velocity fields from transient CFD simulations to provide more accurate field information, used in the adjoint calculations. This resulted in a more accurate resolution of the surface sensitivities near the rear, more separated areas on an *Audi A7* [14]. The difference in surface sensitivities of drag with respect to the normal displacement were compared, which showed that the surface should be moved to create a rear spoiler-like device, as is seen on the *Audi A7*. These differences in surface sensitivities are shown below in Figure 2.2. *Audi* has performed additional research using sensitivity maps, validating the effect of outward displacements to minimize drag on their *Q5* model in the wind tunnel [15]. They also employed DDES simulations for the primal fields to improve the adjoint sensitivity maps.



Figure 2.2: Differences in RANS *vs.* approximate DES surface sensitivities in the rear spoiler area, where red indicates moving the surface inwards to reduce drag and blue indicates moving the surface outwards to reduce drag [14].

Papoutsis-Kiachagias *et al.* further employed the adjoint method to approximate wind-noise calculations to optimize side mirrors. This used a surrogate objective function derived from the square of the turbulent viscosity acting in a volume around the vehicle's driver-side window. In unconstrained optimization routines, the objective function was improved by approximately 25%, suggesting that adjoints could be used to minimize wind noise in automotive applications [12].

Through the use of commercial CFD codes, such as *Simcenter Star-CCM+*, shape optimization has been applied to passenger cars using the discrete adjoint method [16]. This work found similar results to Papoutsis-Kiachagias, in that drag could be reduced on the order of 5% for a passenger hatchback vehicle. In the case of that research, solution accuracy was determined using post-optimality transient CFD simulations.

In most of the existing literature, a single-objective optimization function was used. Papoutsis-Kiachagias *et al.* examined the effects of using multiple operating points for the primal solution, including the vehicle's yaw moment in the objective function [17]. Post-optimality transient CFD for select optimization results along the Pareto frontier showed that RANS-based optimization routines predicted the correct trends for improvement, but tended to overestimate moment improvements compared to the transient CFD results. This study was useful in showing the how stability metrics such as the yaw moment can be minimized, will also showing the effectiveness of drag optimization.

What that body of work did not fully examine is the role the optimization algorithm and optimized region have on the results. Most of the published work in this area focuses on the

optimization of the rear spoiler region of the car to show that optimization is an effective method for reducing aerodynamic drag. However, left unexamined were other regions of the car that are dominated by more complex flow features, such as around the front tires. This region of the front fascia is also of great importance, since better shrouding of the front tires leads to a smaller tire wake, and consequently reduced drag. Therefore, it is valuable to understand if gradient-based optimizers are robust in these regions. Furthermore, it is beneficial to determine which gradient-based optimizers are most favourable for automotive applications, since the more complex flow fields can result in poorer primal solution convergence, which can lead to less accurate sensitivity calculations.

2.1.1 *DrivAer* Geometry

Most automotive research tends to rely on partnerships with OEM's, whereby computer-aided design (CAD) geometry of actual, or concept vehicles is provided. This has two effects on the literature. The first being that reported results are limited to maintain confidentiality and protect the OEM's intellectual property, and the second is the limited public availability of research, since few OEM's are willing to publish their own data. As such, in collaboration with *Audi* and *BMW*, TU Munich's Institute of Aerodynamics and Fluid Mechanics created the *DrivAer* model in 2011 [18]. This car is meant to be a modular platform, allowing for different body configurations, such as the *Notchback*, *Fastback*, and *Estateback* configurations, shown below in Figure 2.3. Given the modularity and suite of existing experimental data, the *DrivAer* platform has been used previously by automotive aerodynamic optimization researchers to examine the efficacy of optimization algorithms [19, 20].




Figure 2.3: *Fastback* (F), *Estateback* (E), and *Notchback* (N) configurations of the *DrivAer* model [18].

The *DrivAer* model has advantages over existing canonical models such as the Ahmed body [21] and the SAE body [22], since it can investigate more complex flow patterns, including along the underfloor, and wheelwell regions, as well as the tire-body interactions. Different configurations of the *DrivAer* model additionally allow for engine bay flow and cooling mod-

elling to occur, allowing detailed investigations. Furthermore, the *DrivAer* model allows for more robust aerodynamic investigations and optimization, given the readily available configurations, each with different flow characteristics. The Ahmed and SAE notchback models are shown below in Figure 2.4, highlighting the simplicity of these models.




Figure 2.4: Existing canonical models, which did not feature as many complex flow structures and geometry as the *DrivAer* model

Given the greater availability of this modular platform, the *DrivAer* geometry has been used in shape optimization routines. Using a simplified fastback model, He *et al.* implemented a discrete adjoint solver within *OpenFOAM* [19], which predicted a drag improvement of approximately 12% compared to baseline. It is important to note that while this is valuable, exempting key flow structures such as tire and mirror wakes can lead to unrealistic results. The use of a generic model had similar results when compared to OEM models such as the *Audi A7*, in that the development of a rear spoiler to reduce the low-pressure region behind the car was generated by the optimization algorithm. This shows that a generic model that is representative of similar vehicle platforms can be used to develop effective optimization methods, without requiring direct OEM support.

Karpouzas *et al.* found further use of the *DrivAer* model for aerodynamic shape optimization routines, relying on continuous adjoints [23]. Drag improvements were predicted when deforming the spoiler area of a more detailed fastback variant, using a volumetric B-spline based mesh displacement tool.

Based on the existing literature, the *DrivAer* model has proven that it can be an informative geometry substitute for vehicle investigations, when real vehicle models are not available. To extend the use of the *DrivAer* geometry and to reflect the move towards more SUV's in the

passenger sector, the *AeroSUV* model was released in 2019, derived from the *DrivAer* car [24]. This model has not been used for much optimization research, but it has the potential to provide a better indication of RANS-based optimization accuracy for flows with larger regions of separation.

2.2 Continuous Adjoints

Since gradient-based optimizers are explored in this research since they tend to be more efficient than gradient-free methods. As such, the CFD code of choice must have a method with which sensitivities can be computed, the most common method being the adjoint method (continuous or discrete). The adjoint method is favourable over other methods such as finite-differencing, since it calculates the Jacobian (i.e. gradient) vector in fewer function evaluations for a large number of design variables than does finite-differencing. In CFD, where flow evaluations are costly, such efficiency is required. Analytic methods are not feasible for automotive applications due to the absence of an exact solution for automotive flows.

In the vast majority of existing literature, automotive aerodynamic shape optimization makes use of the continuous adjoint method. Continuous adjoints are derived from the continuous governing equations for a given problem: the Navier-Stokes equations for fluid flow. This is contrasted to discrete adjoints, where the adjoint equations are derived from the discretized flow equations, typically using an automatic or algorithmic differentiation approach [25]. Aerospace research has shown that discrete adjoints tend to be more robust than continuous adjoints for difficult problems, since they can handle discontinuities in the flow solution, due to the discretized approach [25,26]. This tends not to be an issue for automotive applications, due to lower freestream velocities, reducing the magnitude of velocity and pressure gradients. The existing research landscape confirms that continuous adjoints are sufficiently robust to be suitable for automotive applications. There is however a lack of openly-available research regarding how robust continuous adjoint solutions are for different optimization algorithms (Conjugate Gradient, Quasi-Newton methods) when applied to automotive applications.

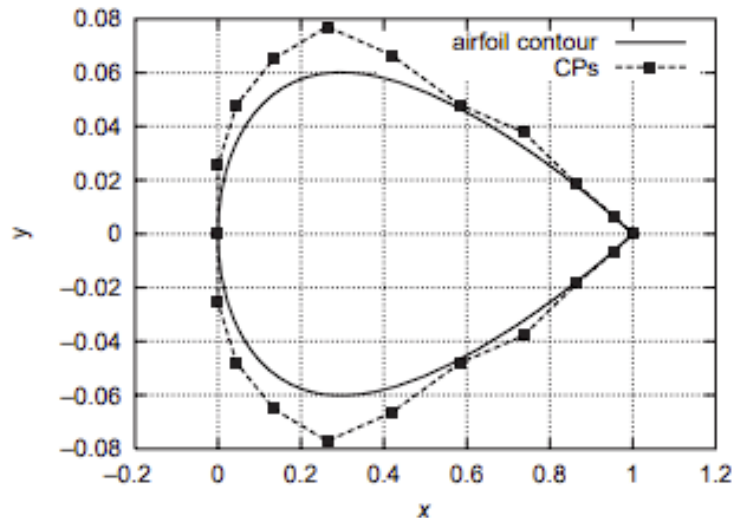
Discrete adjoints have also been used to explore how flow momentum around a body contributes to an objective function of interest (aerodynamic lift, drag, moments), which can be extended to the continuous adjoint method [27]. This was used to study what areas around a Formula Student car contribute most to its downforce. In that presentation, the qualitative use of adjoints was only to understand the flow around the car. This research will take the method one step further, by using the information contained within that field to develop modifications to the *DrivAer* platform. This allows for a qualitative, expert-guided use of the adjoint method, where design variable information can be synthesized during a car’s development. In early phases of a car’s design, this design variable information is important to communicate to the design studio, such that certain design features can be integrated into a car’s styling. This tends to be more valuable than finding an optimal solution for different areas on a car, when styling constraints could invalidate that optimization activity.

2.2.1 Adjoint Turbulence Modelling

The majority of literature regarding the use of continuous adjoints revolves around turbulence closure models to the adjoint equations. In the existing literature, there are two approaches


to account for turbulence in adjoint models. This first is to assume that the turbulence is "frozen" and does not affect any of the sensitivity derivatives, and the second is to have a fully-differentiated adjoint model for a given turbulence model. The second option is currently limited to the adjoint Spalart Allmaras model, which is offered in all major CFD codes with adjoint capabilities. While the $k - \omega$ Shear Stress Transport (SST) model is more frequently used in automotive applications, due to its ability to better predict adverse pressure gradients [28], it is not available in many CFD codes. Furthermore, the adjoint variant of the $k - \omega$ SST model can be less robust than the adjoint Spalart Allmaras (SA) model, due to poorer convergence of the primal solution.

Kavvadias *et al.* developed an adjoint $k - \omega$ turbulence model, that was used in optimization routines and compared to the frozen turbulence assumption [29]. They found that using the frozen turbulence assumption resulted in incorrectly signed sensitivity derivatives compared to finite differences results and the use of an adjoint turbulence model. The magnitude of the sensitivity derivatives tended to be up to an order of magnitude larger than those calculated using finite differences or an adjoint turbulence model. An example of these differences is presented in Figure 2.5 for flow over a NACA 0012 airfoil at 3° angle of attack, parameterized by Bézier polynomials with 12 control points on the suction and pressure surfaces, each. Control points are numbered from the leading edge to the trailing edge for the suction side, followed by the pressure side. The first 24 design points represent the x coordinate of each control point, and the second 48 design points are for the y coordinate. Additionally, the model that assumed frozen turbulence predicted less improvement in the objective function in an internal flow test case.



(a) Control Points and Airfoil Profile

Figure 2.5: Examination of sensitivity derivatives for each control variable for flow over a NACA 0012 wing, with control point parameterization shown. Sensitivity calculations include finite differences, frozen turbulence adjoints, and adjoint $k - \omega$ SST models [29].



(b) Calculated Sensitivities

Figure 2.5: Examination of sensitivity derivatives for each control variable for flow over a NACA 0012 wing, with control point parameterization shown. Sensitivity calculations include finite differences, frozen turbulence adjoints, and adjoint $k - \omega$ SST models [29].

Similar results were found by Zymaris *et al.* during the development of the adjoint SA turbulence model [30]. It was found that an adjoint turbulence model better matched the sensitivity derivatives calculated from direct differentiation methods, when compared to frozen turbulence models. This was exacerbated in cases whose objective functions were heavily dependent on shear stresses within the fluid, such as the minimization of total pressure losses in a duct.

When developing the surrogate objective function to minimize noise, Papoutsis-Kiachagias *et al.* noted that an adjoint turbulence model was necessary to compute sensitivity derivatives [12]. Since this objective function was dependent on primal turbulent viscosity, a frozen turbulence approach would suggest zero sensitivity to turbulent viscosity. This is a case where neglecting turbulent sensitivities would result in no change of a model's geometry.

From the existing body of research, it is made evident that including the effects of turbulence in adjoint sensitivities is imperative for accurate optimization studies. This is most important for models where the results are driven by shear stresses, such as minimization of pressure loss, noise, and drag. As such, all use of continuous adjoints will employ the SA turbulence model for primal computations, and its corresponding adjoint turbulence model for continuous adjoint computations. This would help improve the accuracy of sensitivity analyses, while still be as robust as using the $k - \omega$ SST model for primal calculations, with frozen turbulence for the adjoint calculation. It is important to note that reliance on the SA turbulence model for primal computations would limit the accuracy of the CFD solution. A one-equation SA

approach was less favourable than the two-equation approach of the $k - \omega$ SST model for separation predictions, but the desire to have a robust adjoint turbulence model necessitated the use of the primal SA model.

2.3 Concluding Remarks on the State-of-the-Art

While automotive design is largely aesthetics driven, more focus is being placed on aerodynamic performance to achieve desired range targets since OEM's seek to meet efficiency targets. Automated shape optimization routines have as a result been more popular, to extend the effectiveness of existing designs. It should be noted that these methods are performed at later stages in a design to enhance late-stage performance from the car. Published literature also emphasizes the importance of including turbulence contributions in sensitivity derivatives, for shear-driven objectives, such as drag. Save for some work by Papoutsis-Kiachagias, automotive aerodynamic shape optimization has been focused on minimizing drag in a straight line. Furthermore, the majority of existing work has focused on the optimization of the spoiler region of a car, where the flow structures tend to be relatively well-behaved. The robustness of optimizers in regions dominated by strong adverse pressure gradients has not been fully explored, such as around the front fascia of a car, where such flow has a significant effect on the front tire wakes. Additionally, the efficacy of different gradient-based optimizers such as Conjugate Gradient methods and Quasi-Newton methods have not been examined for automotive applications. This is an important gap in the literature, since traditionally efficient optimization algorithms may suffer from inaccuracies that can arise in the adjoint equations due to insufficient convergence of the primal flow solution. These are often driven by structures such as the front tire wakes, evidencing the lack of optimization performed in these areas. The accuracy of predictions in these areas would be impacted by the reliance on the SA turbulence model, but this was a limit of existing methodologies.

Finally, any RANS-based optimization routines influenced by separated wakes and should use post-optimality CFD models with greater fidelity to determine more realistic changes in aerodynamic performance; this is under-documented in the automotive landscape. The use of DDES has been a favourable high-fidelity method for the automotive industry, resolving turbulent structures at a fraction of the cost of a Large Eddy Simulation (LES).

3 CFD Model Validation Using the *DrivAer* Geometry

This chapter presents an evaluation of two modelling approaches on a canonical automotive geometry using the open-source computational fluid dynamics (CFD) package, OpenFOAM [31]. The simulation approach described in this chapter was used to evaluate different designs inspired from the Momentum Contribution Field, discussed in Chapter 4. The two methods are the use of a $k - \omega$ DDES unsteady simulation approach, and a SA RANS steady-state simulation approach. The DDES model is used to perform the mesh and timestep dependency studies, that will govern the final automated meshing approach, and the desired simulation timestep. The RANS approach is evaluated, since this is the solver employed in the computation of an adjoint’s primal solution, that will be discussed in Chapter 4. This chapter will highlight the accuracy in CFD modelling approaches in terms of absolute force differences to the existing experimental data of the canonical geometry [18], as well as the accuracy of changes in the design of this geometry. It is for this reason that the *DrivAer* geometry is chosen for the canonical model, since it is modular, with a range of vehicle configurations, and it has experimental data against which results can be validated.

3.1 Canonical Automotive Geometry - *DrivAer*

The canonical *DrivAer* geometry is used in this research as a representative test case on which the proposed gradient-based methods can be applied. This geometry has the advantage of being publicly available, and has well-documented wind tunnel results for CFD validation.

The naming conventions for the *DrivAer* model is as follows: the first letter denotes the body style (F, E, N) shown in Figure 3.1, the second letter denotes the floor being simple or detailed (S, D), the third portion denotes if mirrors are used or not (wM, woM, respectively), and finally, if wheels are used or not (wW, woW, respectively). In this work, the *DrivAer Estateback* model with mirrors and a simple floor (E_S_wM_wW) will be the base geometry. These models are shown in Figure 3.2. The E_S_wM_wW model is more representative of current electric vehicles, since they tend to have flat underfloors due to the battery tray. This means that different design modifications would be needed to reduce drag, compared to the underbody of internal combustion vehicles, which contain more complex details, such as exhaust pipes, driveshafts, transmission tunnels, and more. Additional validation will be performed on the *Estateback* model with and without mirrors, as well as with a simple and detailed underfloor. Naming conventions for different areas of interest in this work are presented in Table 3.1 and Figure 3.3.




Figure 3.1: *Fastback* (F), *Estateback* (E), and *Notchback* (N) configurations of the *DrivAer* model [18].





Figure 3.2: *DrivAer Estateback* configurations used for model validation



(c) Detailed Floor


(d) Simple Floor

Figure 3.2: *DrivAer Estateback* configurations used for model validation

The configurations used in the validation study will be designated as E_D_wM_wW, E_D_woM_wW, E_S_wM_wW, E_S_woM_wW, with the E_D_wM_wW model used for the mesh and time step dependency studies. Table 3.1 shows the components used in each of these four models.


Table 3.1: *DrivAer* configurations used in this study

Configuration	Roof Style	Floor	Mirrors	Wheels
E_D_wM_wW	<i>Estateback</i> (E)	Detailed (D)	With (wM)	With (wW)
E_D_woM_wW	<i>Estateback</i> (E)	Detailed (D)	Without (woM)	With (wW)
E_S_wM_wW	<i>Estateback</i> (E)	Simple (S)	With (wM)	With (wW)
E_S_woM_wW	<i>Estateback</i> (E)	Simple (S)	Without (woM)	With (wW)



(a) Top View

Figure 3.3: Naming conventions used for key components of the *Estateback*



(b) Underbody View

Figure 3.3: Naming conventions used for key components of the *Estateback*

3.2 Meshing

To generate the meshes used for all CFD analyses, *OpenFOAM's* built-in, hex-dominant, automated mesher *snappyHexMesh* was employed. This mesher was an automated, fully-parallel utility, that generated axis-aligned, hex-dominant meshes to model triangulated input geometry [32]. The mesher first relied on the creation of a background mesh, which was typically made using the *OpenFOAM* utility *blockMesh*. This background mesh would be comprised solely of hexahedral elements with aspect ratios as close to unity as possible (cubic elements are preferred).

Meshing then proceeded with the castellation phase, where the background mesh is subdivided by feature edges, which can be sharp edges generated using the *OpenFOAM* utility, *surfaceFeatureExtract*, and by the input triangulated surfaces. When the mesh was refined, it was broken into cell levels, with an initial mesh size from the block mesh called *level 0*. If marked for refinement, each cell was divided into four, meaning the length of each edge of the hexahedral element halved. This then increased the cell level by one every time a cell was divided. This also meant that mesh refinement occurred in factors of two, rather than having a continuous growth rate defined for the surface and volume meshes. Once the mesh had been refined based on the defined feature edges and the input surfaces, cells were removed based on a defined point which should have lived within the finite-volume mesh. Then any additional refinement regions, such as wake boxes, volumes created by offsetting the input geometry, etc. are used to provide additional refinement around the input geometry.

Next, the refined mesh was snapped to the input geometry, to match the input surface. The mesher attempted to snap the cell faces to the input geometry within a given quality criteria. In the presence of particularly complex or low-quality geometry, *snappyHexMesh* could not capture the input geometry perfectly.

The final step for *snappyHexMesh* was to generate layers on the surface, used in viscous calculations to resolve the wall-normal boundary layer gradients. This stage worked by projecting the existing mesh away from the surface based on the specified layer thicknesses. If quality criteria were met, layers will be added in this area. If not, the outward projection would be retracted, and layers were not grown on the surface. It was important to note that this retraction of layers could happen on each mesh face along the surface, meaning total layer coverage is not guaranteed by *snappyHexMesh*. A flowchart of this process, along with the different *OpenFOAM* dictionaries and supporting utilities is presented in Figure 3.4. Figure 3.5 presents the intermediate meshes for a sample two-dimensional case that would be generated by *snappyHexMesh*.

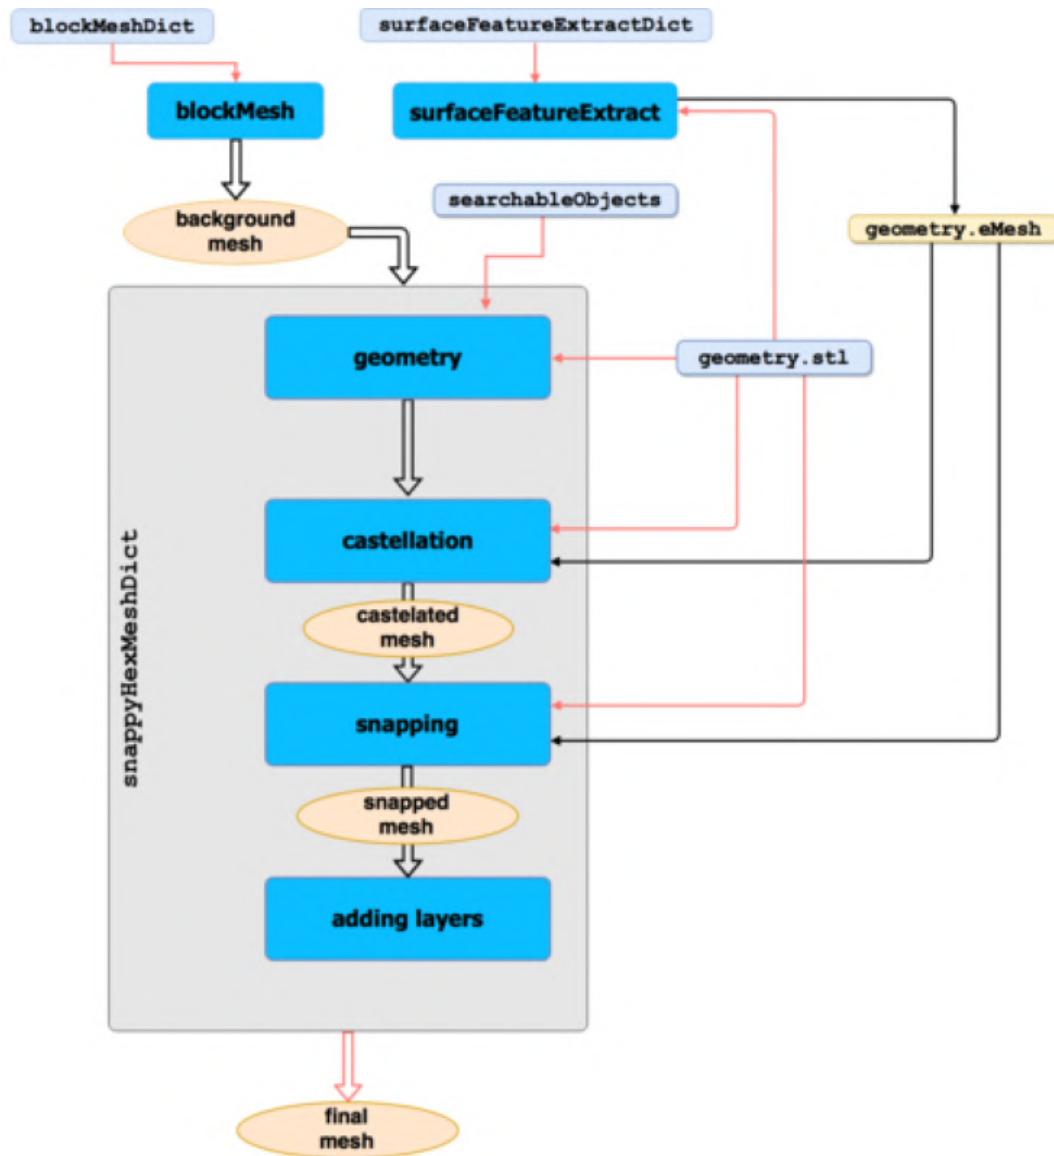


Figure 3.4: Full flow chart, with supporting dictionary files and utilities to run *snappyHexMesh* on an input triangulated geometry file: `geometry.stl` [32]

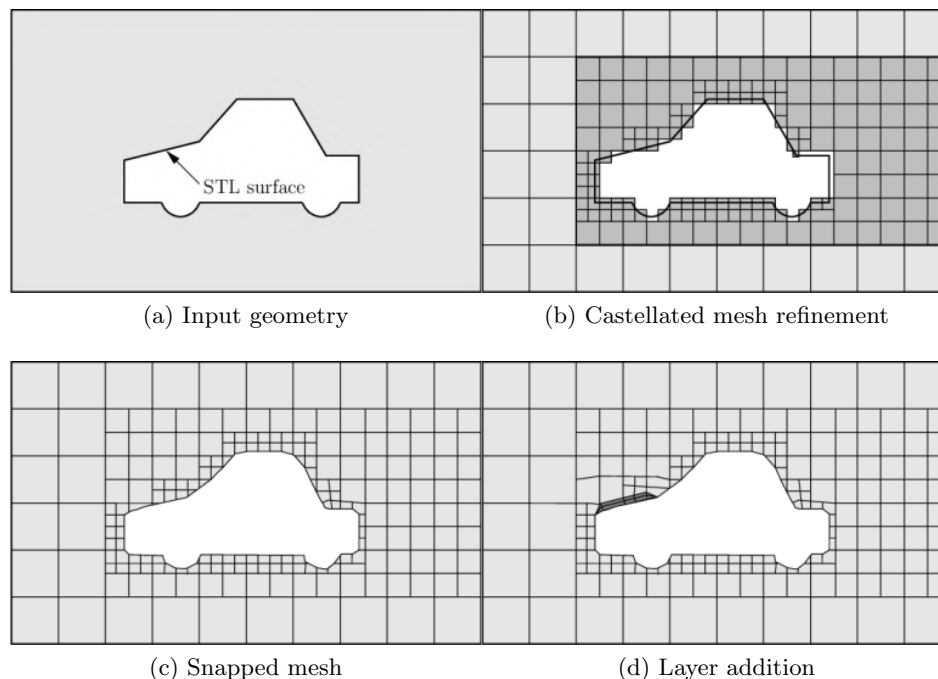



Figure 3.5: Sample two-dimensional meshing problem using `snappyHexMesh` [33]


Since *snappyHexMesh* generated the local mesh sizing by progressively dividing the cells of the background mesh in powers of two, the final cell count could be modified easily by changing the size of the cells in the background mesh. For the *DrivAer* geometry, the background mesh size was chosen to be a 153.6 m long (51.2 m upstream, 102.4 m downstream), 51.2 m wide, and 25.6 m high computational domain, resulting in a blockage ratio of 0.16%, based on industry practices and experience. This meant that to maintain an aspect ratio of unity in all cells in the *blockMesh*, increasing or decreasing the height breakup by one cell resulted in the addition of two cells in the width, and 6 cells in the streamwise direction.

Using *snappyHexMesh*, specific regions (*cellZones* in *OpenFOAM*) could be defined for volume sources to be added, such as porosity models for heat exchangers or moving reference frames for rim spokes. This was done for each of the four wheels, just around the rim spokes, to avoid adding artificial rotation to the flow that was not caused by the rims' motion. Layer addition was controlled using the local aspect ratio of the mesh faces for the height of the anisotropic cells. This results in a high- y^+ meshing approach, for which three layers are placed on all vehicle surfaces, and five layers placed on all tire and rim surfaces. The goal of the high- y^+ meshing approach was to have the first cell centre across the mesh at a non-dimensional wall-normal spacing between 30 and 300. This meant wall functions were applied for modelling viscous walls, rather than integrating the boundary layer to the wall when computing wall shear stresses. A sample mesh along vehicle centreline is shown, along with the overall refinement on the car in Figure 3.6.

The *level 0* size of the mesh was computed by dividing the length of the domain in each direction by the number of subdivisions in each direction. The *blockMesh* generated made use of the *simpleGrading* division scheme in the *blockMesh* utility, resulting in each element



(a) Mesh along car centreline, with a closeup of the rear liftgate



(b) Surface mesh on the vehicle, with a closeup of the side mirror

Figure 3.6: Sample meshes of the *DrivAer Estateback* generated using *snappyHexMesh*

in a direction being of equal length. The specific division schemes that were employed in the mesh dependency study will be covered in Section 3.5.

3.3 Reynolds-Averaged Navier Stokes Simulations

The RANS equations are derived by decomposing the velocity field of the Navier Stokes equations into a mean and fluctuating component:

$$U = \bar{U} + u \quad (3.1)$$

Where U is the total velocity field, \bar{U} is the mean value of that field as a function of time, and u is the fluctuating component of the velocity. This decomposition has the property that the average of the mean value is the mean value, and the average of the fluctuating component is zero. Using this decomposition for the field and taking the time-average of the Navier Stokes equations leads to the following equations for continuity and momentum [34]:

$$\frac{\partial}{\partial t} (\rho) + \nabla \cdot (\rho \bar{\mathbf{U}}) = 0 \quad (3.2)$$

$$\frac{\partial}{\partial t} (\rho \bar{\mathbf{U}}) + \nabla \cdot (\rho \bar{\mathbf{U}} \otimes \bar{\mathbf{U}}) = \mathbf{g} + \nabla \cdot \bar{\boldsymbol{\tau}} - \nabla \cdot (\rho \mathbf{R}) \quad (3.3)$$

In Equation 3.3, \mathbf{g} is the gravity vector, $\bar{\boldsymbol{\tau}}$ is the averaged stress tensor (pressure and viscous stresses), and \mathbf{R} is the Reynolds stress tensor, which arises due to fluctuating velocity fields. These equations as shown are the Unsteady-RANS (URANS) equations, with the steady-state RANS equations forgoing the partial derivatives with respect to time in the continuity and momentum equations. Additionally, for incompressible simulations, it is convenient to divide both equations by the density, ρ , leading to the dynamic viscosity, ν being used in the average stress tensor. This also causes the calculated pressure to actually be p/ρ , which is the calculated pressure variable for incompressible simulations in *OpenFOAM*.

When representing the Reynolds stress tensor, it can be broken into its isotropic and deviatoric components. In *OpenFOAM*'s convention, the isotropic component: $\frac{2}{3}k\mathbf{I}$ is added to the mean pressure term, yielding a \bar{p}' term. The deviatoric term, \mathbf{R}_{dev} is left on its own, due to the additional terms created by these Reynolds stresses. This yields a final momentum equation of:

$$\begin{aligned} \frac{\partial}{\partial t} (\rho \bar{\mathbf{U}}) + \nabla \cdot (\rho \bar{\mathbf{U}} \otimes \bar{\mathbf{U}}) = & \mathbf{g} - \nabla \bar{p}' + \nabla \cdot (\mu \nabla \bar{\mathbf{U}}) \\ & + \nabla \cdot \left[\mu dev2 \left((\nabla \bar{\mathbf{U}})^T \right) \right] - \nabla \cdot (\rho \mathbf{R}_{dev}) \end{aligned} \quad (3.4)$$

The *dev2* operator in Equation 3.4 calculates the deviatoric component of a tensor, in this case being the mean velocity gradient tensor.

The deviatoric components of the Reynolds stress tensor leads to the closure problem of the RANS equations, where there are an insufficient number of equations for the number of unknown variables. This led to the creation of different turbulence models, which seek to model these stresses using different, semi-empirical models. Two eddy-viscosity models will be discussed, both of which model these stresses through a turbulent viscosity term, ν_t (for incompressible solutions). The first of the two turbulence models to be discussed is the SA model [35], which is employed in the continuous adjoint solutions due to the available adjoint SA turbulence model. The second is the $k - \omega$ SSST model [28, 36], which will be developed into the DES variant in Section 3.4.

3.3.1 Spalart Allmaras Model

The SA turbulence model was initially developed for attached turbulent flows, with applications in the aerospace industry. The model itself is not based on a physical quantity, but it attempts to solve for a transport equation of a modified eddy (turbulent) viscosity. This model was developed empirically and using dimensional analysis, rather than being based on a known quantity. The SA model as used in *OpenFOAM* (excluding the non-implemented f_{t2} terms) is given by Equation 3.5 [37]:

$$\begin{aligned} \frac{D}{Dt}(\rho\tilde{\nu}) = & \nabla \cdot (\rho D_{\tilde{\nu}}\tilde{\nu}) + \frac{C_{b2}}{\sigma_{\nu_t}}\rho|\nabla\tilde{\nu}|^2 + C_{b1}\rho\tilde{S}\tilde{\nu} \\ & - C_{w1}f_w\rho\frac{\tilde{\nu}^2}{d^2} + S_{\tilde{\nu}} \end{aligned} \quad (3.5)$$

Equation 3.5 represents the material derivative of the transformed eddy viscosity (left side) being equal to the turbulent diffusion of the eddy viscosity (first two right side terms), plus the production of eddy viscosity (third right side term) minus the dissipation of turbulent viscosity (fourth right side term). The final term of the right side, $S_{\tilde{\nu}}$, represents any additional user-defined turbulence sources to the simulation. The eddy viscosity, ν_t is calculated by Equation 3.6:

$$\nu_t = \tilde{\nu}f_{v1} \quad (3.6)$$

where f_{v1} is given by Equation 3.7:

$$f_{v1} = \frac{\chi^3}{\chi^3 + C_{v1}^3} \quad (3.7)$$

The χ term in this equation is the ratio of turbulent viscosity to the local molecular viscosity. This model has a number of coefficients, which were determined empirically based on different basic turbulent flows. All model coefficients were set to their default values in *OpenFOAM v2206*. In general, the SA model is at most applicable to mildly separated flows, given its calibration for the aerospace industry. As such, it can often underpredict points of separation, which mostly affects drag predictions in highly-separated flows. These types of flows are more common around the wakes of cars, which typically contain a great deal of separation. Despite this, the SA equation is favourable for adjoint computations given the stability of the adjoint SA model. Therefore, the SA model allows for more accurate resolution of turbulence contributions to the adjoint equations compared to a frozen turbulence assumption, while still being robust. Furthermore, unsteady adjoint computations were far too computationally expensive to allow the Detached Eddy Simulation variant of the SA model to be used, so RANS models are typically used for adjoint solutions.

3.3.2 $k - \omega$ SST Model

The $k - \omega$ SST model is a modification to the original two-equation $k - \omega$ model, first proposed by Menter [28]. This turbulence model solves for the turbulent kinetic energy, k , and the specific dissipation rate of turbulent kinetic energy ω . The specific dissipation rate is given by the following equation, as used in *OpenFOAM* [38]:

$$\begin{aligned} \frac{D}{Dt}(\rho\omega) = & \nabla \cdot (\rho D_{\omega}\nabla\omega) + \rho\gamma\frac{G}{\nu} - \frac{2}{3}\rho\gamma\omega(\nabla \cdot \bar{\mathbf{U}}) \\ & - \rho\beta^*\omega^2 - \rho(F_1 - 1)CD_{k\omega} + S_{\omega} \end{aligned} \quad (3.8)$$

Equation 3.8 represents the material derivative of the dissipation on the left side. On the right side, the term involving the gradient of the dissipation is the viscous dissipation diffusion, with the G term being the dissipation production due to the mean velocity shear along with the divergence of the velocity field term. The ω^2 term is the dissipation of the dissipation

(generates k), the $D_{k\omega}$ term being the cross-diffusion. The SST model correction is through the $(F_1 - 1)$ term [36]. This behaves like the traditional $k - \omega$ model near walls, where $F_1 = 1$, and transitions to a $k - \epsilon$ model away from the wall, when $F_1 = 0$, meaning the cross-diffusion term will be included. This modification was made to reduce the standard $k - \omega$ model's sensitivity to freestream turbulence, leading to a more stable model. The S_ω term allows for any user-defined turbulence sources to be added to the simulation. The other terms are model coefficients, which have been derived semi-empirically based on the calibration of the turbulence model with different test cases. All model coefficients were set to their default values in *OpenFOAM v2206*. The k equation is given by Equation 3.9 [36]:

$$\begin{aligned} \frac{D}{Dt}(\rho k) = & \nabla \cdot (\rho D_k \nabla k) + \rho G - \frac{2}{3} \rho k (\nabla \cdot \bar{\mathbf{U}}) \\ & - \rho \beta^* \omega k + S_k \end{aligned} \quad (3.9)$$

As with the dissipation equation, the left side of Equation 3.9 represents the material derivative of the turbulent kinetic energy. On the right side, there is turbulent diffusion of turbulent kinetic energy in its gradient term. Additionally, kinetic energy production due to the mean velocity shear is the G term and additional production is governed by the divergence of the mean velocity term, as with the dissipation equation. The ωk term is a model for the dissipation of the turbulent kinetic energy, and the S_k term are any additional user-defined turbulence sources in the simulation.

To model the Reynolds stresses, an eddy viscosity is calculated, as with the SA model, which is given by Equation 3.10 [28]:

$$\nu_t = a_1 \frac{k}{\max\{a_1 \omega, b_1 F_{23} \mathbf{S}\}} \quad (3.10)$$

The F_{23} term is given by Equation 3.11 [38]:

$$\begin{aligned} F_{23} = & \tanh \left(\min \left\{ \max \left\{ \frac{2\sqrt{k}}{\beta^* \omega}, 500 \frac{\nu}{\omega y^2} \right\}, 100 \right\}^2 \right) \times \\ & \left(1 - \tanh \left(\min \left\{ 150 \frac{\nu}{\omega y^2}, 10 \right\}^4 \right) \right) \end{aligned} \quad (3.11)$$

The reason for the maximum operation in the denominator of the eddy viscosity equation is to limit the production of cross fluctuations in the velocity field in adverse-pressure gradient boundary layers. This is one of the modifications that makes the SST model particularly effective for more separated flows than the SA model. Like the SA model, the SST model is still overly-dissipative, suggesting there is less turbulent kinetic energy than is realistic. This can result in delayed reattachment of separated boundary layers. This makes RANS models less effective for highly separated automotive flows, compared to Detached Eddy Simulations. However, for automotive RANS modelling, the SST model tended to outperform the SA model in areas of separated flow. Furthermore, full understanding of the RANS SST model is required in the derivation of the associated DES variant.

3.4 Delayed Detached Eddy Simulations

As noted, RANS models can perform poorly in automotive applications, since the majority of the aerodynamics involve thin shear layers, with large regions of separation. Using LES would therefore prove to be an advantage, whereby the eddy content in the highly separated wakes would be resolved by the solver. However, LES faces a large grid requirement within the boundary layer, even for Wall-Modelled LES (WMLES) computations. As such, hybrid RANS/LES models are extremely attractive for automotive purposes, where the thin boundary layers can be modelled using a RANS approach, and the large separation regions can be modelled using a LES approach. This is where DES is an extremely useful purpose to the automotive industry, and also the reason why these models are used in this research.

Since the Reynolds-Averaged and Filtered (used in LES) Navier Stokes equations take the same form, it is possible to solve a simulation containing both RANS and LES regions, using a unified turbulence model that can switch between RANS solutions and LES solutions [39]. Depending on the turbulence model chosen, the behaviour of the model will have different characteristics. For example, the SA-DES model was calibrated to perform like a traditional WMLES solution, whereas $k - \omega$ SST-DES models behave more like a hybrid RANS/LES solution. For this research, only the $k - \omega$ SST DES and DDES models will be discussed.

3.4.1 $k - \omega$ SST Detached Eddy Simulation

The SST-DES equations are generated by modifying the turbulent kinetic energy equation of the RANS SST model, shown in Equation 3.9. This DES turbulence model is then described by Equation 3.12 for the turbulent kinetic energy [40]:

$$\begin{aligned} \frac{D}{Dt}(\rho k) = & \nabla \cdot (\rho D_k \nabla k) + \min(\rho G, c_1 \beta^* \rho k \omega) - \frac{2}{3} \rho k (\nabla \cdot \bar{\mathbf{U}}) \\ & - \rho \frac{k^{1.5}}{\tilde{d}} + S_k \end{aligned} \quad (3.12)$$

In Equation 3.12, the important term is the $k^{1.5}/\tilde{d}$ term, where the RANS or LES behaviour is controlled by the \tilde{d} parameter. This is a modified length scale given by Equation 3.13:

$$\tilde{d} = \min \left\{ C_{DES} \Delta_{LES}, \frac{\sqrt{k}}{\beta^* \omega} \right\} \quad (3.13)$$

With a RANS length scale l_{RANS} being given by Equation 3.14:

$$l_{RANS} = \frac{\sqrt{k}}{\beta^* \omega} \quad (3.14)$$

The Δ_{LES} term is a measure of the local grid sizing, or the LES length scale. The calculation of this term in *OpenFOAM* simulations is the *cubeRootVol* method, which computes the grid filter size as the cube root of each cell's volume. This means that the spacing requirement is calculated once per simulation, and is not dependent on flow conditions. While this is a computationally inexpensive approach, it is a model for the actual flow spacing, and can become skewed for highly anisotropic cells. In the case where the local grid sizing is larger than the RANS length scale, the turbulence model will operate as a RANS model. In the

case where the DES length scale is smaller than the RANS scale, the solution will act as an LES computation, where the sub-grid (or sub-filter) viscosity is calculated using a turbulent kinetic energy equation. In the regions where the solution uses an LES computation, the resolved turbulent kinetic energy is still subject to the local grid spacing, which determines the size of the spatial filter.

3.4.2 $k - \omega$ SST Delayed Detached Eddy Simulation

Despite DES solutions providing a computational cost advantage over pure LES models, this initial model was highly susceptible to a phenomenon known as *Grid-Induced Separation* (GIS) [41]. There are three types of grids that can be used in a DES simulation. The first type are grids with a thin boundary layer, where the local spacing tangent to the wall will exceed the wall-normal spacing. This causes the boundary layer to be solved with a RANS formulation, with the transition to an LES solution occurring outside of the boundary layer, which is the optimal operation regime for DES. The second type is one for which the grid spacing within the boundary layer is sufficiently small, such that most of the boundary layer is solved using a turbulent kinetic energy sub-grid scale model, with a RANS model very close to the wall, and a small transition region in between. This type of solution is a form of WMLES, but is not inline with the initial purpose of DES solutions, which sought to reduce the computational cost of scale-resolving simulations. The third type of grid, which was most problematic was one for which the application of a RANS or LES model was ambiguous through much of the boundary layer. This tends to occur when the transition to the LES model is in the upper two-thirds of the boundary layer. This results in a dissipation of the RANS eddy viscosity, due to the model transition, but the grid is not fine enough to fully resolve the boundary layer fluctuations that the LES model seeks to capture. As such, there is a rapid decrease in shear stress due to the loss of eddy-viscosity, which creates non-physical GIS [41].

With this overarching problem noted in initial DES models, a modification to the calculation of the DES length scale was proposed, to delay the point at which the simulation would transition from a RANS to a LES model. Equation 3.15 is a modification to Equation 3.13 proposed by Menter and Kuntz to delay this transition [42]:

$$\tilde{d} = l_{RANS} - f_d \max\{0, l_{RANS} - C_{DES} \Delta_{LES}\} \quad (3.15)$$

This equation makes use of a delaying function, f_d , to delay the onset of the activation of the LES model. The full activation of the LES model would therefore only occur in areas where the grid spacing was sufficiently smaller than the RANS length scale. The delaying function used in this computation is given by Equation 3.16 [42]:

$$f_d = 1 - \tanh\left([8r_d]^3\right) \quad (3.16)$$

With the independent variable of this function being given by Equation 3.17 [42]:

$$r_d = \frac{\nu_t + \nu}{\sqrt{\nabla U \nabla U} \kappa^2 l_{RANS}^2} \quad (3.17)$$

The inclusion of the molecular dynamic viscosity helps ensure that near the wall, this term remains non-zero reducing the likelihood of full activation of the LES model. This change to the switch between the RANS and LES models was the key differentiation in the DDES

approach, while also making it more robust to traditionally ambiguous grids. The other benefit of the DDES model is the transition from RANS to LES is more abrupt, resulting in a smaller ambiguous interpolation region between the two modelling approaches [41].

3.4.3 Grid Refinement in DDES Models

Since DDES models resolve turbulent content when in their LES mode, the amount of resolved content is dependent on the local grid size (the spatial filter). As such, not all DDES simulations are "created equal", with finer grids resolving more of the eddies in the flow. A sample of the effect of different turbulence models and the grid resolution on DES modelling is shown in Figure 3.7. As the grid is refined from coarse to fine, the smallest turbulent structures in the wake decrease in size due to the spatial filter size decreasing.




Figure 3.7: Effects of turbulence model and mesh resolution on DES flow field, a) Coarse SA-DES, b) Fine SA-DES, c) Fine SST-DES. Adapted from [43]

As the mesh resolution increases, turbulent structures become finer, due to fewer length scales being smaller than the spatial filter. In automotive applications, additional resolution of these eddies will yield a better representation of the car wakes, primarily around the wheels, mirror, and rear fascia. This is most influential on drag predictions, which are highly dependent on the separation points and pressure within these separated regions. The typical industrial approach of automotive CFD is to generate the grids to be as coarse as reasonably possible, to maintain the maximum amount of computational throughput. As such, the upper limit

of a DES simulation approaching a full LES simulation is infeasible for industrial automotive applications.

As a final remark on DDES turbulence modelling: for automotive flows with large regions of separation, they will outperform comparable RANS models in terms of wake predictions, and turbulent content resolution. However, the time-dependent nature of the LES equations makes DES unfeasible for most adjoint computations. Therefore, the RANS turbulence models are still of value to the automotive industry.

3.5 DDES Mesh and Timestep Dependency Studies

In the experimental data first generated by Heft, a 40% scale model of the *DrivAer Estateback* was used. The force and pressure data presented in that research was for a Reynolds number of 4.87 million. The velocity at this Reynolds number was determined to be 40 m/s based on freestream conditions for a dry environment at 20° C. [18]. Since the full-scale *DrivAer Estateback* was used for validation, the freestream velocity was set to 16 m/s (40% of 40 m/s), to preserve the Reynolds number when comparing data.

The *DrivAer* model was found to have a blockage ratio, ϕ , of 8% in a wind tunnel with an effective cross-section of 1.74 m x 2.4 m. Therefore, the frontal area of the scale model was found with Equation 3.18:


$$A = \phi A_{eff} \quad (3.18)$$

Therefore, the frontal area of the 40% scale model was 0.334 m², which translates to a frontal area of 2.088 m² for the full-scale car. This value was used for normalizing drag and lift.

For the mesh and timestep dependency studies, the DDES equations were solved using *OpenFOAM's pimpleFoam* solver, which is a combination of the Pressure-Implicit with Splitting of Operators (*PISO*) [45] and Semi-Implicit Method for Pressure Linked Equations (*SIMPLE*) [46] algorithms, for time-stepping solutions. The combined *PISO-SIMPLE* algorithm is abbreviated as the *PIMPLE* algorithm [47]. Both the *PISO* and *SIMPLE* algorithms solve the discretized RANS equations through an initial momentum equation solution, that is then corrected through a pressure equation. These classes of solvers are typically known as a *segregated* solver. These algorithms are in contrast to *block-coupled* solvers, where pressure and velocity would be solved simultaneously in a single matrix.

The benefit of the *PIMPLE* algorithm over the standard *PISO* algorithm is it does not require the maximum Courant number to not exceed one, as it uses *SIMPLE*-based inner pressure correction iterations to allow for higher Courant numbers at each time step. Second-order spatial and temporal schemes are employed with the *PIMPLE* algorithm to ensure simulation accuracy. The simulation is allowed to run for five seconds, with averaging of the flow fields and forces occurring over the last two and a half seconds. While some of the *Estateback* configurations are symmetric, an averaging window sufficiently long to remove time-averaged asymmetries due to wake vortex shedding would require too long an averaging window to be feasible. The $k - \omega$ SST DDES turbulence model is used for modelling the eddy and sub-grid scale eddy viscosities. Boundary conditions used for both the farfield and the car

are presented in Figure 3.8. The inlet-outlet scheme used employs a velocity specification on the inlet, with a static pressure on the outlet. The ground is modelled using a tangential wall velocity. Further details on the specific boundary conditions used are shown in Figure 3.8.



(a) Farfield boundary conditions



(b) Car boundary conditions

Figure 3.8: Boundary conditions for farfield and car surfaces

To understand the sensitivity of the results with respect to the mesh and unsteady time step, two studies were performed. The first varied the mesh sizing and time step simultaneously on the *E_D_wM_wW* configuration, and the second varied just the time step using a constant mesh sizing, for all four *DrivAer Estateback* configurations presented in Table 3.1. This was done to first understand if the drag trends were consistent for different time steps as a function of cell count. This study allowed for a mesh sizing to be chosen, which was used to analyze the time step sensitivity for the four *Estateback* configurations of interest.

To vary the mesh refinement, five background mesh sizes were chosen, such that all cells in the initial mesh had an aspect ratio of unity, with characteristic lengths of 1.4222 m, 1.5059 m, 1.6000 m, 1.7067 m, 1.8286 m. A larger characteristic length resulted in a coarser mesh. The specific division of the background mesh, which was 153.6 m long, 51.2 m wide, and 25.6 m high are shown in Table 3.2. It was noted that modifying the *level 0* size to change the mesh resolution was a limitation of the *snappyHexMesh* process. Further modifications

to this could have involved changing the size of the domain to get larger changes in *level 0* size, but this was not desired, due to the impact this would have on the blockage ratio. This would have had additional affects on the results that could not be isolated from changes in the mesh. The results of the initial mesh and time step sweep are shown in Figure 3.9. Supplemental images showing the effect of the divisions on the surface volume meshes are available in Appendix A.

Table 3.2: *blockMesh* divisions used to generate E_D_wM_wW meshes

Mesh	Divisions ($x \times y \times z$)	Base Length [m]	Cells (M)
Very Coarse	84 x 28 x 14	1.8286	33.0
Coarse	90 x 30 x 15	1.7067	38.3
Medium	96 x 32 x 16	1.6000	43.6
Fine	102 x 34 x 17	1.5059	49.9
Extra Fine	108 x 36 x 18	1.4222	56.2

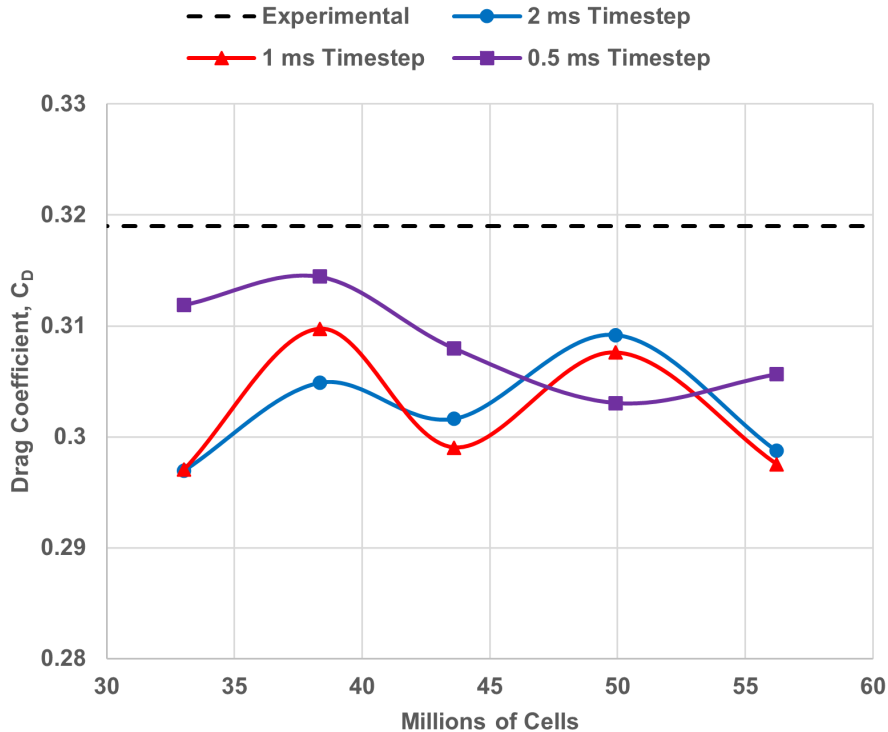


Figure 3.9: Drag coefficient as a function of cell count and time step

Figure 3.9 shows that as the cell count increases (decrease in the base length), the drag coefficient oscillated for all three time step sizes. This oscillation was most similar for the 2 ms and 1 ms time steps, with the 0.5 ms time step oscillating more slowly. All three curves tended to converge however, around 50 million cells in the mesh. This mesh sizing has the best agreement among the three time steps, as well as the best overall agreement with the experimental data for the E_D_wM_wW configuration. In general, it is expected that the drag coefficient should decrease with a smaller timestep, due to better resolution of higher

frequency eddy content. This also helps to reduce temporal numerical dissipation.

As such, the Fine mesh sizing was chosen for subsequent analyses, since it had the best agreement in the 2 ms and 1 ms time steps, and the 0.5 ms time step followed a consistent trend. The three timesteps tested were then used with this mesh sizing to evaluate the remaining three *DrivAer Estateback* configurations. These results are shown in Figure 3.10.

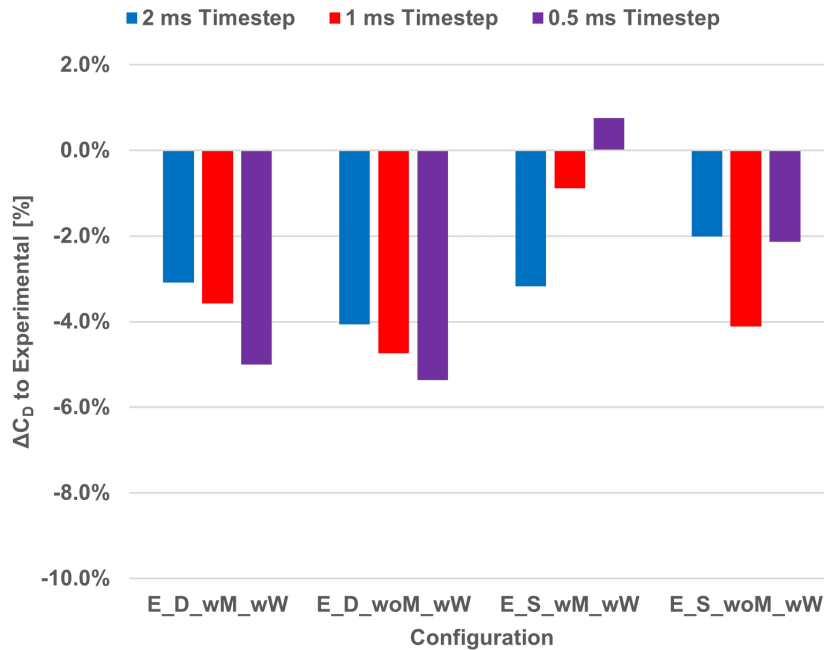


Figure 3.10: *DrivAer* configuration drag coefficient as a function of timestep

From Figure 3.10, the 2 ms timestep had the lowest average absolute percent difference to the experimental data at 3.1%. The 1 ms and 0.5 ms time steps have a slightly greater average absolute error of 3.3%. The 2 ms time step data had the lowest standard deviation in the absolute error across the four configurations at 0.8%, compared to 1.7% and 2.2% for 1 ms and 0.5 ms, respectively. Despite the 0.5 ms timestep having the lowest absolute error for the E_S_wM_wW configuration, it was valuable to have consistent trends for different, but similar geometries. The smaller standard deviation across the different designs for the 2 ms timestep, means it will be more consistent across different design modifications, as discussed in Chapter 4.

3.6 DDES Validation Against Experimental Data

Using the results presented in Figure 3.10, this is re-plotted to examine the residuals in the CFD simulations with respect to the experimental drag of the four tested configurations, shown in Figure 3.11. It was found that the simulations underpredicted the drag of the vehicle by around 4%. While detailed flow surveys of the tire wakes were not performed in Heft's work [18], it is expected that the drag discrepancies stem from inaccuracies within the tire wake regions. Since the tires are within close ground proximity and rotate, they have

highly turbulent and complex wakes, which can be difficult to predict accurately with CFD. However, since CFD drag predictions were within 5% of the experimental results, the DDES models were deemed to be sufficiently accurate for the purpose of this research.

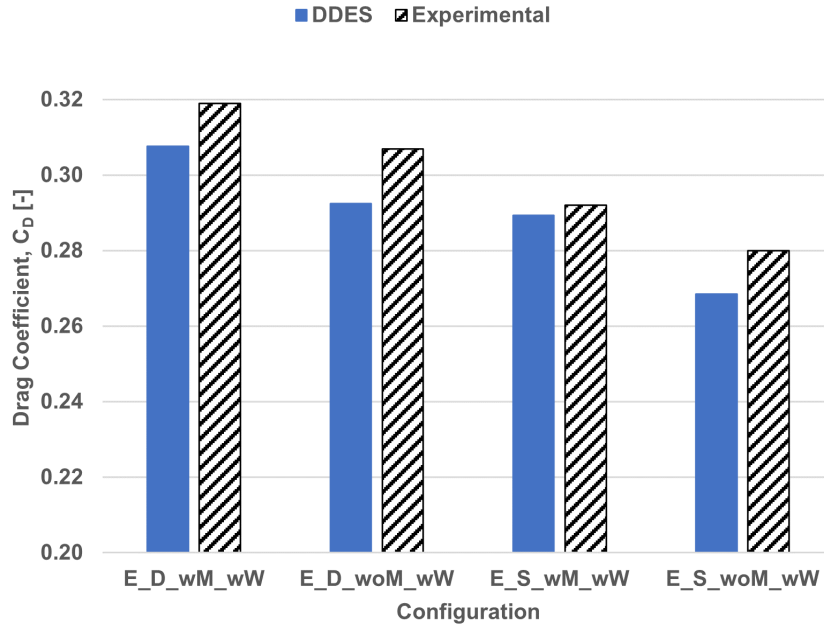


Figure 3.11: Drag coefficient comparisons between CFD and experimental results for the four tested *Estateback* configurations

Using the *Fine* mesh approach outlined in Table 3.2 and a 2 ms timestep, the DDES simulations of the E_D_wM_wW configuration could be compared to the existing experimental data. For the scale model wind tunnel results, pressure probes were placed along car centreline on the top and bottom surfaces, as well as along the $z = 60$ mm plane [18]. For the full-scale model, this plane was sampled at $z = 150$ mm. These pressure results as well as the errors (residuals) between CFD and experimental data are presented in Figures 3.12 - 3.14.

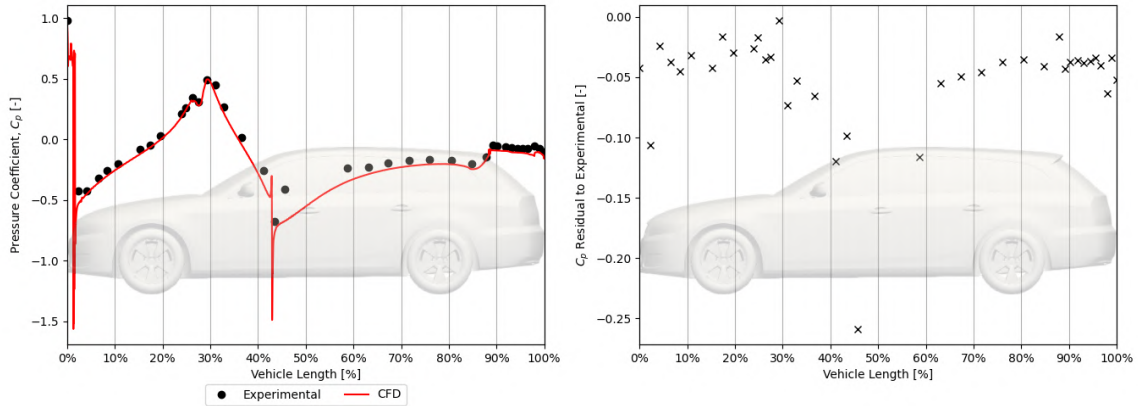


Figure 3.12: Comparison of pressure coefficient between CFD and experimental results along the top of the car in the $y = 0$ mm plane. Absolute residuals between CFD and wind tunnel results in the right image

Figure 3.12 shows that there is very good agreement between CFD and experimental results along the top of the car on the centreline, with the error mostly in the range of $|\Delta C_p| \leq 0.05$. There is an over-prediction of suction at the junction between the seam at the windshield and body as it transitions to the roof line, between 40% and 50% of the car's length. This was not of concern, since the geometry in this region has particularly tight curvature. It was expected that CFD would overpredict the suction in this region. Furthermore, the scale model likely had simpler features in this area that would result in a smaller suction peak. There was particularly good agreement along the rear window and liftgate, which is one of the most important contributors to vehicle drag.

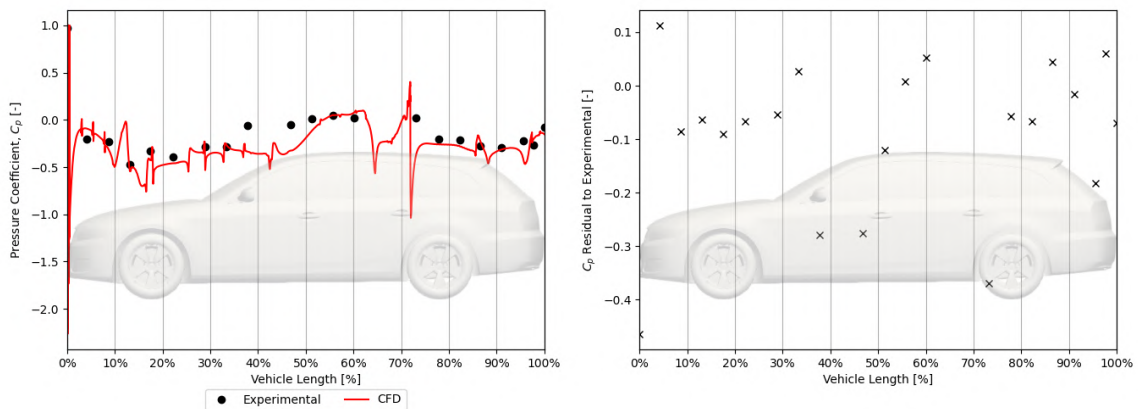


Figure 3.13: Comparison of pressure coefficient between CFD and experimental results along the bottom of the car in the $y = 0$ mm plane. Absolute residuals between CFD and wind tunnel results in the right image

Figure 3.13 shows that agreement deteriorated in the underfloor results between CFD and experimental data. Since the test article was a 40% scale model, it is possible some of the detail and divots in the vehicle's underside were not well captured, explaining many of the differences between the 30% and 60% vehicle length. There was however good overall resolu-

tion of the underfloor pressure trends.

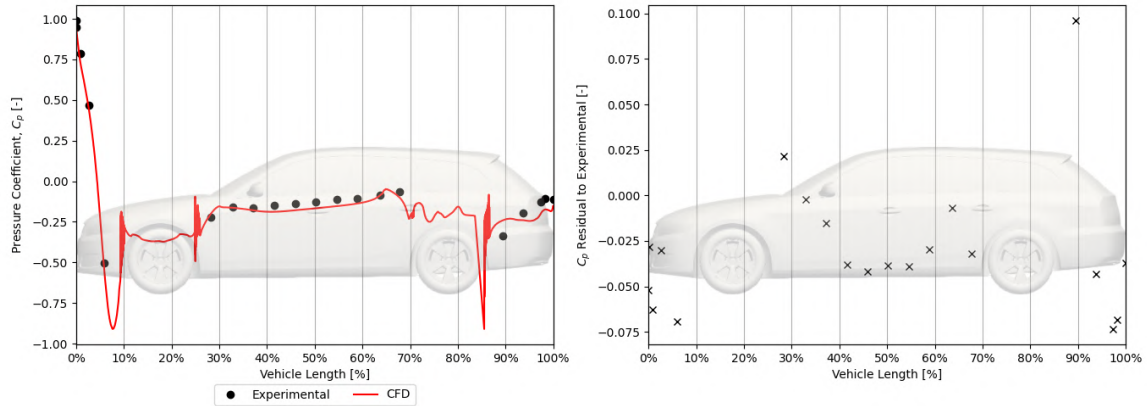


Figure 3.14: Comparison of pressure coefficient between CFD and experimental results along the left side of the car in the $z = 150$ mm plane. Absolute residuals between CFD and wind tunnel results in the right image

Examining the surface pressures along the $z = 150$ mm plane for the full-scale vehicle in Figure 3.14 shows there was very good agreement between CFD and experimental results. There was a slight under-prediction of surface pressures between 30% and 60% vehicle length. This suggests that the tire wake is slightly inaccurate, hence the under-prediction of pressure along the body. This could be due to the rims not being modelled using mesh motion. When applying source terms to model rim motion, there is some error associated with the results due to rotation being added artificially to more of the flow than is realistic. However, the additional computing and memory requirements for solid body rotation for rotating rim simulations were too expensive for the computing resources available.

Furthermore, the CFD simulations tended to underpredict the pressure around the separated region of the rear fenders and fascia. This area is difficult to obtain accurate predictions, due to the rear tire wake interactions with the geometry and the overall amount of separation. This again, was a result of the mesh resolution, rim modelling approach, and the eddy content that is resolved by the LES model, that is not able to be captured fully due to computational limits.

All three pressure slices show that there was acceptable agreement between the CFD results and experimental results. As such, this meshing and modelling approach was deemed appropriate in evaluating the performance the *DrivAer* model. Despite CFD results tending to under-predict vehicle drag, there were no suspension arms modeled, which will also contribute to this error between CFD and experimental results. There is strong overall agreement between numerical and experimental results and supplemental images examining the surface pressures on the different *Estateback* configurations tested are presented in Appendix B.

3.7 Comparison of DDES and RANS Simulations

When comparing the DDES (denoted SST-DDES) simulations to the RANS modelling approach, to be representative of highway speeds, the freestream and ground speeds shown in

Figure 3.8 were increased to 120 km/h, with a corresponding increase in the tire rotational speed. Since the freestream velocity roughly doubled, the timestep for DDES simulations was decreased to 1 ms to match the average Courant number from the validation simulations. Mesh settings were not changed from the validation study, since wake sizes did not become larger at higher Reynolds number. Furthermore, the meshing approach chosen was shown to predict consistent drag changes for different designs. This behaviour was assumed to be consistent at the higher freestream speed as best practices were followed.

These comparisons were performed on the E_S_wM_wW configuration, since this configuration had the adjoint solver applied to it, which makes use of the RANS equations. As with the DDES simulations, RANS models employed second-order spatial schemes to compute gradients and divergence terms. The RANS model employed the SA turbulence model (denoted SA-RANS) to close the RANS equations. Comparisons of the surface pressures on the car are presented in Figures 3.15 - 3.17.





Figure 3.15: Comparison of pressure coefficient between SST-DDES modeling and SA-RANS modeling on the front of the car

In Figure 3.15, it is clear the CFD predictions in areas with little separation were quite similar. In fact, the pressure distributions over the front fascia and fronts of the side mirrors are almost identical. This is expected in these areas, as the flow will be mostly attached. As such, the SA model should be accurate in these areas. The main area that was different was the pressure distribution at the base of the windshield. The SA results predicted lower pressures, which suggests that it is not able to resolve the separation off the edge of the hood as well as the SST-DDES model was. Based on the results shown in Figure 3.12, the SST-DDES model slightly underpredicted the pressure in this area, which meant the SA-RANS model would under predict the pressure in this area more.



(a) Surface Pressure Coefficient

Figure 3.16: Comparison of pressure coefficient between SST-DDES modeling and SA-RANS modeling on the rear of the car



(b) ΔC_p SA-RANS - SST-DDES

Figure 3.16: Comparison of pressure coefficient between SST-DDES modeling and SA-RANS modeling on the rear of the car


Moving towards the rear of the car, shown in Figure 3.16, it was evident there are greater differences in the highly separated wake. Since this was a separated region, it was expected that the SA model would not have as accurate a prediction. As such, it tended to predict higher pressures on the rear fascia, which meant it had a lower drag prediction than did the SST-DDES model. This effect is also seen on the back of the side mirrors, which also had higher pressures for the SA model than the SST-DDES model. This suggested that the RANS modeling overpredicted the flow attachment before separating off the side mirrors. A quantitative investigation of the differences between these two methods was also shown in Figure 3.16. This shows direct differences in surface pressure between the RANS and DDES models, with areas of low magnitude having the greatest agreement.

At the rear of the car, the main region where the SA model overpredicted the pressure coefficient compared to the SST-DDES model tended to be a $|\Delta C_p| = 0.05 - 0.125$ overprediction.

A similar margin was seen on the back of the rear tires. This meant the SA model had a lower drag prediction, since the suction on the rear fascia and the backs of the rear tires contributes significantly to the drag. It was found that the SA model predicted the drag of the car to be 6.8% lower than the SST-DDES model did, most of which stemmed from the rear fascia.



(a) Surface Pressure Coefficient



(b) ΔC_p SA-RANS - SST-DDES

Figure 3.17: Comparison of pressure coefficient between SST-DDES modeling and SA-RANS modeling on the underside of the car

For the underside of the car, highlighted in Figure 3.17, the main areas that differed between the two models were the splitter region and front wheelwell region. The SST-DDES model predicted a larger suction peak than the SA model did at the splitter leading edge, suggesting that the SA model was unable to capture the leading edge separation as effectively as the unsteady solution. Since there tended to be lower pressure over more of the splitter region for the SA model, as more of the flow is attached in this region. In the front wheelwell area, the SST-DDES results tended to predict less suction than the SA model did. This was likely a result of the different upstream conditions between the two models. Floor devices on a vehicle are highly sensitive to the energy in the airflow, meaning that the differences in the splitter predictions will affect the wheelwell predictions. Some asymmetries were noted in the

surface pressures on the floor of the SST-DDES model. This was attributed to cross-body velocity fluctuations due to the front tire wakes. Despite the large averaging window used, these oscillations had a low frequency, meaning the averaging period was not long enough to fully average these fluctuations. Existence of features like this contributed to the error in the DDES solutions. A further look at the differences is also presented in Figure 3.17.

For the direct differences between the SA model and the SST-DDES model shown in Figure 3.17, the wheelwell region tended to have a pressure coefficient that was 0.05 - 0.125 lower in the SA model. Around the edge of the wheelwells, this rose to about 0.2, where the separation at this edge was not predicted as well by the SA model. At the front splitter, the leading edge had a pressure coefficient more than 0.3 larger for the SA prediction than the SST-DDES prediction. This then extends to the SA model tending to predict a pressure coefficient more than 0.3 lower than the SST-DDES model. This is characterized by the SA model not fully capturing the leading-edge separation on the splitter, which cause the high suction peak and loss of suction seen in the SST-DDES model. There are some additional discrepancies in the diffuser region, with the tendency to be slight over-prediction of the pressure coefficient around the edges of the geometry, where there will be more separation. This is on the order of a difference of $|\Delta C_p| = 0.05$. Further asymmetries were noted on the rear diffuser, which are attributed to insufficient time-averaging of the rear tire wakes.

In summary, the surface pressure fields between the SST-DDES and SA-RANS modeling were similar, primarily in the regions without much flow separation. In the areas of higher flow separation such as the rear fascia, the results began to differ more, typically by a $|\Delta C_p|$ of 0.05-0.125.

The velocity fields were also compared for SST-DDES and SA-RANS simulations. Four planes were chosen for this purpose: the $y = 0$ mm and -700 mm planes, and the $z = -90$ mm and -15 mm planes. The $y = 0$ mm plane highlighted the flow features along the centreline of the car: the flow coming off the front splitter, spoiler, and rear diffuser. The $y = -700$ mm plane highlighted flow features around the left-hand side wheels and their wakes. The $z = -90$ mm and -15 mm planes highlight not only the flow downstream of the tires, but also that from the front bumper and around the rear fenders.

Figure 3.18 presents flow field differences on the centreline plane of the car, for the two primal computation methods. The key differences occurred in the areas around which separation was expected: the hood, splitter, and roofline. At the hood's leading edge, there were lower velocities for the DDES case, which meant there was less flow attached to the hood than in the RANS case, which was accelerated more by the hood's curvature. Under the car, the DDES model predicted more separation around the splitter, since there were lower velocities underneath the car, indicative of increased losses at the leading edge. The overall wake size behind the car was smaller for the RANS results, suggesting that there was a smaller separation region behind the car, largely being dominated by better-attached diffuser flow, reducing the size of the counter-rotating vortex system behind the car. As such, in all areas of great separation, the DDES model predicted there would be larger wakes for this vehicle, which will have a noticeable impact on subsequent adjoint flow sensitivity solutions.

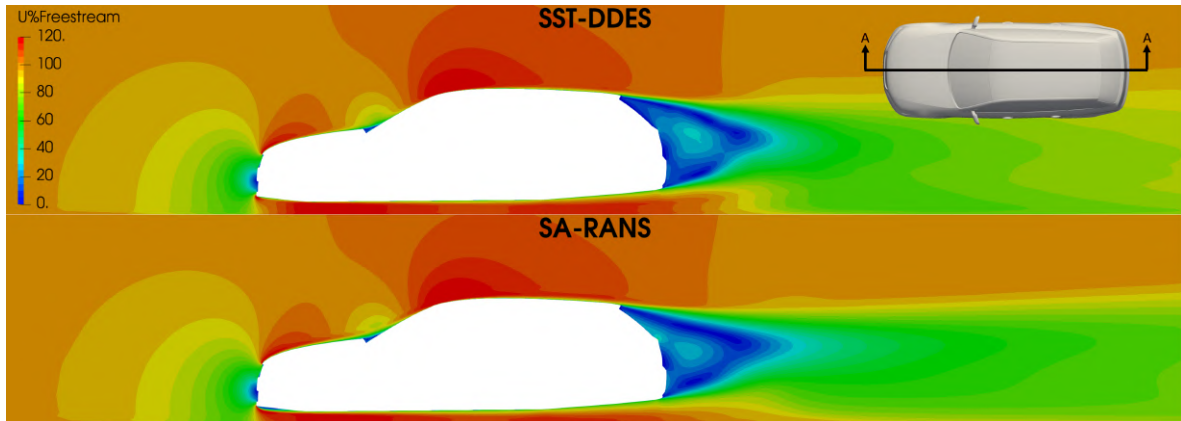


Figure 3.18: Normalized velocity contours on the $y = 0$ mm plane for the two primal solution methods

The flow on the $y = -700$ mm plane for the two primal computation methods is shown in Figure 3.19. At the front of the vehicle, the velocity fields were quite similar, since this flow was largely dominated by the attached regions around the front bumper, which DDES and RANS would predict to be similar. The main differences occurred downstream of the front and rear tires: downstream of the front tires, the DDES model predicted higher velocities in the wake of the front tire than did the RANS model. In the DDES case, there was no noticeable attachment of the flow to the front wheelwells. It can be seen there was some partial attachment of the flow to the front wheelwell in the RANS model, before separating, reducing downstream velocities. Behind the rear tires, the DDES model predicted more loss and separation than the RANS model did, which was expected. The DDES model had a larger low velocity region, which would further impact an adjoint computation for the areas interacting with the front and rear tire wakes.




Figure 3.19: Normalized velocity contours on the $y = -700$ mm plane for the two primal solution methods

The flow on the $z = -90$ mm provides useful insight into the limitations of the RANS model, which is shown in Figure 3.20. Around the front tires, the DDES model predicted a smaller wake than the RANS model did. Since the DDES fields used the $k - \omega$ turbulence model,

there would be a more accurate prediction of the pressure gradients in this region. It is noteworthy that around the rear tires, the DDES model predicted less flow attachment than the RANS model, since there was greater flow acceleration and higher velocities. This suggested that while there was early separation in the RANS model around the front tires, the turbulent content dissipated too quickly, whereas the DDES model propagated that turbulence farther downstream. This reduced the flow energy around the rear wheels, indicated by the wider wake in the DDES model. While there were lower velocities downstream of the rear wheels for the RANS model, the wake structures were narrower than the DDES structures, suggesting the RANS model was predicting the car to be more streamlined than it was.




Figure 3.20: Normalized velocity contours on the $z = -90$ mm plane for the two primal solution methods

The $z = -15$ mm plane, shown in Figure 3.21 further shows that the DDES model was predicting a less streamlined wake than did the RANS model, along with some asymmetry in the wake, driven by periodic vortex shedding. As with the $z = -90$ mm plane, there were higher velocities within the rear tire wake predicted by the DDES model. However, the wake of the RANS model was narrower and smoother than the DDES predicted. This meant there was more turbulent content within the wake of the DDES model, while the RANS model was dissipating the turbulent content much faster. There was also less attachment along the rear fenders predicted by the DDES model. The flow in this region was less aligned with the fender geometry than in the RANS model, meaning there was less separation off the rear tires in the latter case. This reduced the effect that this wake had on the rear fender flow, which the RANS model could not predict as accurately as the DDES model could.




Figure 3.21: Normalized velocity contours on the $z = -15$ mm plane for the two primal solution methods

The SA-RANS model predicted a higher total pressure in the wake of the car, which did not extend as close to the ground plane as the SST-DDES model. This suggested that there was less loss within the wake, explaining the higher pressures on the rear of the car in the SA-RANS approach. The smaller overall wake also suggests that there was less loss from the rear tire wake oscillating across car centreline. Both of these effects are consistent with the SA-RANS model being less suited to computing regions of separation than SST-DDES. This would affect the ability for adjoint solutions to provide accurate sensitivities in regions with great adverse pressure gradients due to modelling errors in both the primal and adjoint solutions. The effect of primal turbulence modelling approaches on adjoint computations will be discussed in Chapter 4.

4 Momentum Contribution Field Analysis and Design Modifications

This chapter examines a method with which a continuous adjoint solution can be used to provide the user with a measure of how the momentum around a body contributes to a given objective function. This field, dubbed the Momentum Contribution Field, will be examined for the aerodynamic drag around the *DrivAer Estateback*. This chapter extends the work presented in Chapter 3, by using the validated RANS and DDES models and meshing approach for the continuous adjoint solutions.

The information from the MCF will then be used to develop design modifications to the car, to reduce its drag. This is a novel demonstration of the ability of a RANS-based MCF to provide sufficient information to yield design modifications that reduce the *Estateback* drag, and introduces the use of DDES simulations to improve the accuracy of the RANS-based MCF. Furthermore, discussions on the limitations of the RANS-based MCF are presented. The application to the test geometry is then used to outline how the MCF method can be applied to industrial design routines in its most effective form, along with the areas where MCF information alone may not be sufficient in yielding design improvements. This is novel to the automotive field, with previous qualitative applications of continuous adjoint solutions being focused on surface sensitivity maps. These do not show how the volumetric flow solution can cause interactions on the car's surface. A sensitivity map only provides the user with a relative measure of how the surface should be deformed along or against the local surface normal to reduce the objective function. While this information is valuable, it does not contain the same physical phenomenological information that a MCF contains.

4.1 Continuous Adjoint Model

In order to compute the MCF around a body for a given objective function, a method with which the sensitivity of the objective function with respect to the flow momentum/velocity at each grid point can be computed. For traditional numerical sensitivity analysis, such as the finite difference method, this would require $n + 1$ CFD evaluations, where n is the number of points in the mesh. Since this would be prohibitively expensive for a mesh of practical size, a cost-effective method with which the sensitivity analysis can be performed is the use of the continuous adjoint method. The adjoint method computes the sensitivity of the main flow variable fields (pressure, velocity, etc.), providing a measurement of the objective function's sensitivity with respect to those fields, at each grid point. This allows a full sensitivity analysis to be performed for the cost of two flow solutions, for any number of grid points, n . As such, for a large number of design variables (mesh points in this instance), the continuous

adjoint method is preferred to finite differences.

The evaluation of continuous adjoints within *OpenFOAM* is performed using the *adjointOptimisationFoam* solver, which was implemented by the National Technical University of Athens' (NTUA) Parallel Computing and Optimization Unit (PCOpt) [48]. The solution is computed using the *SIMPLE* algorithm, to solve the governing equations iteratively. This acts similarly to the *simpleFoam* solver for incompressible RANS flows within *OpenFOAM*.

In their derivation, the continuous adjoint equations stem from the differentiation of the continuous RANS residual equations for some objective function, J (drag, lift, moment, etc.), which has been fully described in the literature [23, 29, 48]. They have also provided differentiation of the SA and $k - \omega$ SST equations [29], with the differentiated SA model shown in Equation 4.3 below. These give rise to the following equations:

$$R^{p_a} = -\frac{\partial U_{a,j}}{\partial x_j} = 0 \quad (4.1)$$

$$R_i^{U_a} = U_{a,j} \frac{\partial U_j}{\partial x_i} - \frac{\partial (U_j U_{a,i})}{\partial x_j} + \frac{\partial \tau_{a,ij}}{\partial x_j} + \frac{\partial p_a}{\partial x_i} + \tilde{\nu}_a \frac{\partial \tilde{\nu}}{\partial x_i} - \frac{\partial}{\partial x_l} \left(\tilde{\nu}_a \tilde{\nu} \frac{C_Y}{Y} \epsilon_{mjk} \frac{\partial U_k}{\partial x_j} \epsilon_{mli} \right) = 0, \quad i = 1, 2, 3 \quad (4.2)$$

$$R^{\tilde{\nu}_a} = -\frac{\partial (U_j \tilde{\nu}_a)}{\partial x_j} - \frac{\partial}{\partial x_j} \left[\left(\nu + \frac{\tilde{\nu}}{\sigma} \right) \frac{\partial \tilde{\nu}_a}{\partial x_j} \right] + \frac{1}{\sigma} \frac{\partial \tilde{\nu}_a}{\partial x_j} \frac{\partial \tilde{\nu}}{\partial x_j} + 2 \frac{c_{b2}}{\sigma} \frac{\partial}{\partial x_j} \left(\tilde{\nu}_a \frac{\tilde{\nu}}{\partial x_j} \right) + \tilde{\nu}_a \tilde{\nu} C_{\tilde{\nu}} + \frac{\partial \nu_t}{\partial \tilde{\nu}} \frac{\partial U_{a,i}}{\partial x_j} \left(\frac{\partial_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i} \right) + (-P(\tilde{\nu}) + D(\tilde{\nu})) \tilde{\nu}_a = 0 \quad (4.3)$$

$$R^{\Delta_a} = -2 \frac{\partial}{\partial x_j} \left(\Delta_a \frac{\partial \Delta}{\partial x_j} \right) + \tilde{\nu} \tilde{\nu}_a C_{\Delta} = 0 \quad (4.4)$$

In these equations, the residual, R , of the main adjoint variables p_a , $U_{a,i}$, $\tilde{\nu}_a$, and Δ_a are used to generate a system of partial differential equations (PDE's) that can be solved using the *SIMPLE* algorithm. In these equations, q refers to the adjoint pressure, u_i is the i^{th} component of the adjoint velocity vector, $\tilde{\nu}_a$ is the adjoint to the SA turbulence quantity, and Δ_a is the adjoint to the wall distance. The P and D terms in Equation 4.3 are the production and dissipation of the primal modified eddy viscosity. All other unlisted variables are model coefficients. A more detailed derivation of these equations has been previously presented [23, 48]. These residual equations then form the governing equations for the continuous adjoint model, which are discretized for use by a CFD solver.

The solution of the continuous adjoint equations was performed using *OpenFOAM's adjointOptimisationFoam* solver [48]. Since the SA turbulence model was used for the primal solution, the corresponding adjoint SA model was included to evaluate the differentiated turbulence equations. The adjoint SA model was included because of the improved sensitivity derivative approximation accuracy when including the differentiated turbulence model. Using a frozen turbulence approach for the adjoint solution has been shown to produce inaccurate and sometimes incorrect sensitivities [49]. The primal equations are solved using second-order

spatial schemes for the momentum and turbulence variables. In contrast, the adjoint equations are solved using a first-order upwind spatial scheme, which was found to provide better solution stability.

4.2 Momentum Contribution Field

The adjoint field variable of particular interest to this research is the adjoint velocity, which is defined in Equation 4.5, given by the symbol U_a . It is comprised of the components $U_{a,x}$, $U_{a,y}$, and $U_{a,z}$. Each component of the adjoint velocity can be thought of as the sensitivity of an objective function, J , with respect to the local momentum, meaning the physical phenomenon described by the adjoint velocity is:

$$\begin{aligned} U_a &= \nabla_U J \\ &= \frac{\partial J}{\partial U_x} \hat{i} + \frac{\partial J}{\partial U_y} \hat{j} + \frac{\partial J}{\partial U_z} \hat{k} \end{aligned} \quad (4.5)$$

Equation 4.5 shows that contours of the adjoint velocity field, U_a will provide the user with information as to how the primal velocity, U , contributes to the objective function, J . Based on *OpenFOAM*'s implementation of continuous adjoints, each component takes on a positive value when momentum at some point in space reduces the cost function and a negative value when it increases the cost function.

This feature of the adjoint velocity field has given rise to the definition of a MCF that turns the adjoint velocity into a scalar measure of the local flow momentum's contribution to the adjoint's objective function [27]. This MCF for a given objective function, J , can be defined by:

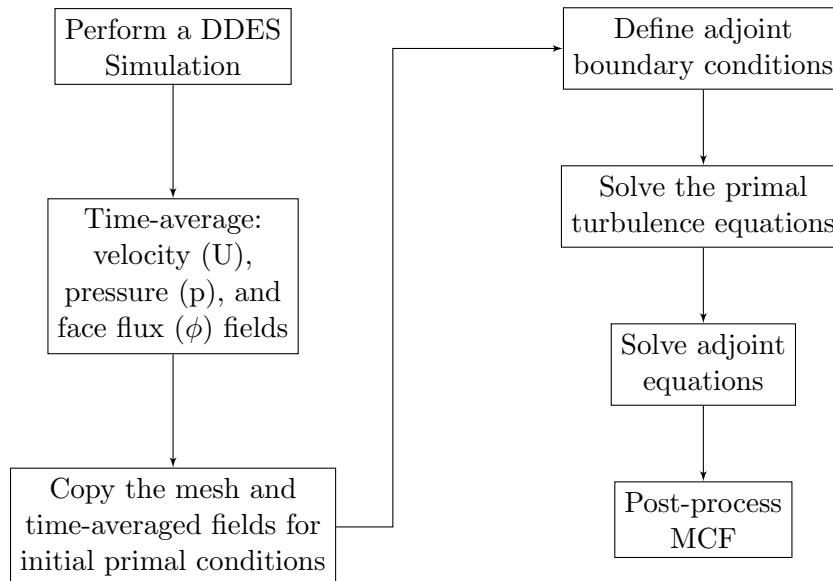
$$\begin{aligned} MCF &= -U \cdot U_a \\ &= - \left(U_x \frac{\partial J}{\partial U_x} + U_y \frac{\partial J}{\partial U_y} + U_z \frac{\partial J}{\partial U_z} \right) \end{aligned} \quad (4.6)$$

In Equation 4.6, the negative dot product of the primal velocity, U , and the adjoint velocity, U_a , is taken to scale the adjoint sensitivities by the local primal velocity. This means that areas of low primal magnitude, will show less of an overall MCF contribution. Furthermore, the dot product turns these two fields into a scalar product that represents the total flow velocity's contribution to the objective function. The negative sign in Equation 4.6 is incorporated such that the following behaviour is achieved: positive areas in the flow increase the cost function, and negative areas decrease the cost function. In problems where the objective function is to be minimized, such as for aerodynamic drag, momentum should be decreased within positive MCF regions, and increased in negative MCF regions. The magnitude of the MCF is of little physical meaning, but the relative scale of different areas around a body provide insight as to which regions contribute the most to the objective function (aerodynamic drag in this research). This allows for key areas where the drag can be minimized to be determined, and provides guidance on the types of modifications to the flow field that would reduce aerodynamic drag.

4.3 DDES-Based Improvements to Momentum Contribution Field Predictions

In an industrial setting, the MCF of a vehicle would likely not be calculated without first performing a full CFD evaluation of the geometry using the modelling approach chosen for the program. As such, prior to performing the adjoint solution, it would be expected that if the program was using DDES simulations (as is common in passenger car development), time-averaged velocity and pressure fields would already be available.

Within *adjointOptimisationFoam*, users can compute the primal solution with constant velocity, face flux, and pressure fields, from which the turbulence fields are calculated. This allows the user to make use of time-averaged velocity and pressure fields calculated from a DDES simulation as the constant velocity, pressure, and face flux fields, which would be more accurate than their RANS counterparts. Therefore, the MCF calculation can be sped up, and made more accurate by employing the *RASTurbulenceModel* keyword for the primal solution, instead of the *simple* keyword in *system/optimisationDict*, which solves the turbulence model's PDE's on the constant velocity, face flux, and pressure fields. This yields the following simulation approach:



Since a DDES simulation better resolves the large-scale turbulent structures, this modification to the RANS-based MCF approach would yield a more accurate prediction of the sensitivities around the body. Additionally, since the primal calculation only computes turbulence quantities it is on average, two to three times faster than computing the full RANS primal. Applying this methodology to the *DrivAer Estateback* studied would highlight where inaccuracies in the RANS solution affects the MCF. When computing the MCF, it is important to note that the DDES simulations only used the SST model to calculate time-averaged velocity, face flux, and pressure fields. The primal computation for both MCF approaches relied on the SA turbulence model, so the adjoint SA equations may be solved.

Computing the MCF using a RANS-based method is approximately 50% of the cost of a DDES simulation. Using time-averaged DDES fields, computing the MCF takes approximately 10% of the time it would take to run a DDES simulation, being five times faster than the RANS-based method, assuming the DDES fields were already computed. In the case that a geometry has not yet had a DDES simulation run, it would cost 110% that of a DDES simulation to compute the MCF using this improved method. This means that using this improved method is on the order of the cost of a RANS-based MCF, if not faster, depending on the workflow in which the method is employed.

4.4 RANS MCF Evaluation

Using the continuous adjoint formulation described in Section 4.1 and the MCF described in Section 4.2, *OpenFOAM* was used to compute a drag MCF for the *DrivAer Estateback*, using a RANS primal solution. This provided a number of areas that were isolated for having a high positive or negative MCF value around the E_S_wM_wW configuration. From this solution, the key areas of drag contribution can be visualized along three distinct planes: the $y = 0$ mm plane (car centreline), the $z = -115$ mm plane, and the $z = -15$ mm plane. The $y = 0$ mm plane will highlight how the different stagnation regions contribute to the vehicle's drag as well as how the wake behind the car could be modified to reduce drag. The $z = -115$ mm plane will show how the flow around the front fascia and into the front tires contributes to drag. Finally, the $z = -15$ mm plane will highlight the interaction between the rear tires and the rear fascia.




Figure 4.1: Drag MCF on the $y = 0$ mm plane. The top right view is zoomed in on the front fascia and splitter and the bottom right view is zoomed in on the rear fascia and diffuser

The MCF images in Figure 4.1 highlight how the momentum is contributing to drag. There is a positive region around the stagnation streamline that leads to the front fascia. Furthermore, there is a negative region underneath the vehicle. This suggests that moving mass flow from the front fascia underneath the car will reduce drag. It is important to note the positive zone near the leading edge of the splitter means that momentum in this separation region is increasing the drag of the vehicle. As such, it will be important to reshape the front splitter to reduce this separation. This will require a larger radius and gentler curvature at the leading

edge region to reduce adverse pressure gradients.

The top right image of Figure 4.1 shows that there is a positive region (drag increasing) arising as the flow goes over the hood. This stems from the pressure buildup on the windshield to which this flow contributes. As such, it could be beneficial to add a small raised flick here to get the incoming flow better aligned to the curvature of the windshield. This would reduce pressure buildup in this region, thereby reducing drag. The flow in this area is one that has a high amount of downstream interactions, since the windshield flow helps dictate the amount of air that spills over the A-pillars down the side of the car, which subsequently affects the pressure behaviour there.

The bottom right image of Figure 4.1 shows that filling in the wake behind the car will be beneficial in reducing drag. By increasing the energy of the flow behind the car, the base pressure there can be increased, reducing pressure drag. There is a strong negative MCF region along the rear spoiler of the car, suggesting that extending this feature, or even curving it downwards would increase the local flow velocity. This would contribute to filling in the rear wake more effectively, reducing car drag by increasing base pressure. The positive MCF region around the diffuser also suggests that a more gradual expansion angle would reduce drag, by reducing suction buildup, and leaving more attached flow.




Figure 4.2: MCF for the drag of the car on the $z = -115$ mm plane

The MCF field shown in Figure 4.2 highlights the drag contribution stemming from the flow that moves alongside the front fascia towards the front wheels, over what is known as the *front bumper cheek*. There is a clear area of positive MCF, since the cheek flow not only hits the front of the tires, but also gets entrained in the wheelwell, generating drag. As such, there would be a benefit from adding a small flick feature to this portion of the cheek to help outwash the flow in front of the front tires. This would help reduce the amount of air that hits the front of the front tires, reducing front tire drag, but will also help reduce the amount of flow that is drawn inwards towards the wheelwells. This can also help inject more

momentum alongside the front tires, which is shown to be a negative MCF region. This additional momentum would help decrease the size of the front tire wake, decreasing loss. In the stagnation region, there were some asymmetries in the MCF solution. These were numerical artifacts, rather than physical phenomena. The primal computation for both the SST-DDES and SA-RANS approaches were similar and showed symmetry in Figure 3.21. This highlighted that there were some numerical instabilities in the underlying adjoint solver, regardless of the primal solution method.




Figure 4.3: MCF for the drag of the car on the $z = -15$ mm plane

Figure 4.3 shows how the momentum from the rear tires' wakes contributes to drag, on the $z = -15$ mm plane. Behind the rear tires, there is a low pressure region, due to their wake losses. Since the rear bumper does not extend out towards the edge of the rear tires, it is unable to mitigate this wake by encouraging flow to remain attached to the bodywork. This leads to separation and a suction buildup on these surfaces, increasing the drag of the car. The positive MCF region does not extend all the way around the car, which is important to note. This means that the rear fender should likely be filled in within this positive region. As such, a sharper edge will be placed at the trailing end of this infill region, leading to a defined separation region. In automotive designs, these trip features can be beneficial to cause the rear wake to become fully turbulent earlier, rather than generating Von Karman vortex streets behind the car. This specific region will be of interest in determining the efficacy of the MCF, since this is a region of greater separation, which will suffer in accuracy due to the RANS turbulence closure model being used.

4.5 DDES-Based Improvements to the MCF

To analyze the changes in the MCF due to the primal solution method, the same four planes used in the primal velocity field comparisons from Chapter 3 will be used. This will highlight the effect the noted changes in the primal velocity fields had on the adjoint solution.

The differences in MCF predictions for the $y = 0$ mm plane are shown in Figure 4.4, for the DDES and RANS primal methods. The key differences occur around the splitter, diffuser, and hood. In the splitter region, the DDES predicted that most of the flow was creating drag on the vehicle, whereas the RANS model predicts a near-zero region just perpendicular to the surface (annotated in Figure 4.4). This is representative of the additional separation predicted by the DDES solution, which would lead to additional drag on the car, something the RANS model did not fully capture. Additionally, there was more positive drag contribution predicted by the DDES model on the underfloor. Part of this was likely due to the interactions of the tire wakes with the underfloor, which are quite different between the two models. This difference in MCF was propagated to the diffuser, with the DDES model predicting a much smaller positive MCF region than the RANS model did. This stems from the differences in the wakes between the DDES and RANS models. While the diffuser does create drag due to the local suction and its orientation, it helps to reduce the size of the wake. Since the RANS model predicted a smaller overall wake, it is expected that the diffuser contributes a larger proportion of the wake drag.

Along the hood, the DDES model predicted a larger negative MCF region. This can be attributed to the lower primal velocities in this region due to greater leading edge separation of the hood in the DDES model, compared to that of the RANS model. This negative MCF region also extended along the hood in the DDES-based solution, while it did not in the RANS solution. Therefore, a hood flick that would reduce the velocities in this area, as will be shown in Section 4.7 would not necessarily reduce drag. This effect was predicted by the design evaluation CFD models. This area further highlights the importance of employing more accurate simulation methodologies to compute the velocity, face flux, and pressure fields for an adjoint computation.

Despite the differences in the roofline flow between the primal solutions, the two resultant MCFs predict very similar negative MCF regions there. Since this area is particularly well-conditioned, attached flow, it was expected that the downstream wake differences would have less of an effect on the adjoint, and would create more similarity in the solutions.




Figure 4.4: Momentum Contribution Field contours on the $y = 0$ mm plane for the two primal solution methods

The MCF contours on the $y = -700$ mm plane shown in Figure 4.5 show the differences largely occurred around the tires. Within the barrels of the wheels, the MCF results were completely different, indicative of the extremely turbulent nature of this flow, which was resolved very differently using the DDES approach compared to the RANS approach. Upstream of both the front and rear tires, the MCF tended to be more positive for the DDES primal than it did for the RANS model. For the front tires, this difference was caused by the interaction of the splitter flow and the wheelwell. This was a significant region of separation, that is better resolved by the DDES approach. This was supported by the contribution within the top of the wheelwell, being nearly zero for the DDES model compared to the RANS model. Resolution of this turbulent content is imperative to effective sensitivity prediction in these areas where numerous components interact.

For the rear tires, there was a more consistent trend in the MCF, with a distinct positive region downstream of the tires and a distinct negative region downstream of the positive region in the DDES model, in contrast to the inconsistent results for the RANS model. Since the DDES model would resolve the rear tire wake more effectively, it is expected that it would be better suited to capturing the drag contribution from the adjoint solution.




Figure 4.5: Momentum Contribution Field contours on the $y = -700$ mm plane for the two primal solution methods

Due to the differences in the front tire wakes on the $z = -90$ mm plane shown in Figure 3.20, the resultant MCF was quite different around the rear tires. These contours are shown in Figure 4.6. Since the primal DDES solution predicted more separation off the front tires, it is expected that more momentum would contribute to the drag of the vehicle in this region, propagating farther downstream. While there was a thinner positive region between the axles in the DDES model, the positive MCF region stems from the front bumper cheek, all along the car. This suggests that the turbulent content getting resolved in this region was having a greater effect on the sensitivities than the dissipated turbulent content in the RANS model. From Figure 3.20, there were higher velocities along the rear tire sidewalls in the RANS model, suggesting better flow attachment in this area, than the larger wake predicted by the DDES model. This results in a further increase in the positive MCF region around the rear tires in the DDES model than in the RANS model.




Figure 4.6: Momentum Contribution Field contours on the $z = -90$ mm plane for the two primal solution methods

Figure 4.7 compares the MCFs for the two primal solutions on the $z = -15$ mm plane. The key differences between the two solvers occurred around the front bumper cheeks and along the rear fenders, behind the rear tires. The DDES solution predicted there was a positive MCF region from the momentum along the front bumper cheeks, whereas the RANS model did not. The positive MCF region stems from the air that would impinge on the tires and separate off the tires' sidewalls. Therefore, if flow velocity were to be increased as the RANS solution suggests, it would impinge more air on the tires, increasing drag. It has been shown that the effective modification in this area is to outwash this tire flow, reducing the impingement and the tire wake. Along the rear fenders, the overall trend is similar between the two solutions, with a positive region perpendicular to the car's surface, then a negative region, then another positive region further outboard. However, the positive region in the DDES solution is larger perpendicular to the wall, suggesting that a trip tab that encouraged the fender flow to separate would be beneficial to the design. The RANS solution also suggests a trip tab, but the position of maximum contribution is more vague, meaning the DDES solution carries more information regarding the placement of a trip tab.




Figure 4.7: Momentum Contribution Field contours on the $z = -15$ mm plane for the two primal solution methods


In summary, these results highlight the importance of using primal calculation methods that better capture the turbulent eddies in complex automotive flows, primarily around the tires, and behind the car. RANS primal solutions can be valid in some areas of the car, primarily in regions with little to no separation, such as along the upper door region and along the roof. Furthermore, in industrial settings, the MCF would likely be computed for a geometry that has already been evaluated. As such, if this program used DDES simulations for design evaluations, the desired geometry would already have the time-averaged fields available. As such, computing the MCF from time-averaged fields was roughly five times faster than computing a full RANS-based MCF based on testing for the *DrivAer Estateback*. This makes the method not only more accurate, but faster to compute than the initial RANS-based methodology.

4.6 RANS MCF-Inspired Design Modifications

From the RANS MCF, a number of areas on the vehicle were isolated on which design modifications were created. Four areas were chosen to be modified, the front splitter/lower front bumper leading edge, the front bumper cheeks, the rear fender, and the roof spoiler. For each area on the car, post-processing images will only be presented for the best design in each region. All post-processing images can be found in Appendix C.


4.6.1 Front Splitter Modifications

Modifications to the front splitter were largely focused on increasing the curvature between the front bumper and floor, helping to reduce the adverse pressure gradient in this region. Three geometries were developed to improve the flow in this region, presented in Figure 4.8. Not only was the radius between the front bumper and underbody increased, but the panel was then blended into the flat floor, yielding a single, tangent profile. Between the different splitter options, the size of the leading-edge splitter radius, and the blending profile of the splitter into the flat floor was varied to create the different designs.




(a) Splitter Region of Interest


Figure 4.8: Splitter profiles, baseline and the three options




(b) Baseline Splitter



(c) Splitter Option 1



(d) Splitter Option 2



(e) Splitter Option 3

Figure 4.8: Splitter profiles, baseline and the three options

From the splitter modifications, the overall design goal was not only to increase the leading-edge radius, but to also flatten the geometry. This was done without modifying the approach angle of the vehicle. On the baseline splitter, a sharp transition existed between the splitter and the flat floor. This was made more gentle in its curvature, to reduce any separation and loss of energy being fed to the floor. The effect these modifications had on the drag of the car is shown in Table 4.1.

Table 4.1: Splitter designs and their drag contribution

Splitter Option	Description	ΔC_D [%]
1	Larger Splitter Radius	+0.2
2	Option 1 with Downstream Blending	-1.6
3	Option 2 with Outboard Blending	-1.9

Of the three models, Splitter Option 3 performed the best, reducing drag by 1.9%, compared to a 1.6% reduction for Option 2 and no improvement from Option 1. The majority of improvement stemmed from the splitter having more attached flow, which helped reduce losses to the front tires. Figure 4.9 shows how the pressure coefficient on the underside of the car changed between the baseline splitter to Option 3. For the baseline splitter, there was a region just downstream of the splitter's leading-edge where the pressure coefficient was close to zero, while the upstream pressure coefficient was lower than -1. This suggests that there was a sharp suction peak at the leading edge due to the sharp curvature, that caused separation over the majority of the splitter surface. For Splitter Option 3, this region produced suction, which suggests that the majority of separation has been mitigated. As such, the MCF has correctly informed drag contribution stemming from this separation, which was able to be successfully mitigated, thereby reducing the drag of the car by 2.1%.





Figure 4.9: Comparison between the baseline splitter and Splitter Option 3 of pressure coefficient on the floor


4.6.2 Front Bumper Cheek Modifications

As was shown in Figure 4.2, there was an area of positive drag contribution from the air moving around the bumper cheek, impinging on the front tires, as well as getting trapped

within the wheelwell. The focus in this region was the addition of a 5 mm flick, perpendicular to the wheel arch surfaces outwashing the bumper cheek flow around the front tires. Three different heights of the 5 mm flick were tested, as presented in Figure 4.10.




(a) Front Bumper Cheek Region of Interest



(b) Baseline Front Bumper Cheek

(c) Cheek Flick 1



(d) Cheek Flick 2

(e) Cheek Flick 3

Figure 4.10: Front bumper cheek geometry, baseline and three options

The three modifications are described, and their performance changes are highlighted in Table 4.2. This table shows that there is a balance between the height of the cheek flick and the

drag reduction. Too small a flick, such as Option 1, and the drag saved on the front tires is gained on the flick itself and within the wheelwell. Option 3 had the tallest flick, meaning there was the largest drag gain on the flick’s surface and the front wheelwell, greater than the drag savings on the front wheels. Option 2 had intermediate drag gain on the flick itself, while having the greatest drag reduction on the front tires. Option 2 reduced the front wheel drag by 33.2% compared to 17.5% for Option 1, and 16.6% for Option 3.


Table 4.2: Front bumper cheek flick designs and their drag contribution

Flick Option	Height	ΔC_D [%]
1	Smallest	+0.2
2	Medium	-2.1
3	Tallest	+0.5


Figure 4.11 highlights the pressure drag coefficient around the front tires, which was the area of interest for the baseline vehicle, and Front Bumper Cheek Option 2. Around the front of the tires, Front Bumper Cheek Option 2 helped generate more suction around the sidewall of the tire, which created a local thrust effect, due to the sidewall’s surface normal. This additional suction on the sidewall also reduced the pressure buildup on the forwards-facing portion of the tires. However, this had a negative effect of increasing drag within the front wheelwells, which was noticed for this model. However, this gain in drag did not outweigh reduced tire drag.

Considering the rear of the front tires in Figure 4.11, the improved outwash from Option 2 helped reduce suction on the aft surfaces of the tires, reducing pressure drag in this region. Like the change on the forwards surfaces, the reduction of suction in this area reduced the suction in the rear faces of the wheelwell, further increasing drag in this area. As such, there was a balance between the amount of drag saved on the wheels themselves, and the amount of drag gained within the wheelwells.

The MCF around the front bumper cheek was able to provide additional information regarding how flow should be manipulated. Specifically, it showed that the key area was the trailing edge condition of the front bumper cheek, which governed how the air would interact with the front tires. The outwash generated by Front Bumper Cheek Option 2 helped increase momentum around the sidewalls of the front tires, which is shown in Figure 4.2 to be an area of negative contribution. It is valuable to note that the front bumper cheek is a region of low separation, so for RANS-based sensitivities it was expected to provide a more accurate prediction of the primal and adjoint flow fields.



(a) Front Three Quarter View



(b) Rear Three Quarter View

Figure 4.11: Comparison of pressure drag coefficient around the front tires for the baseline and Option 2 front bumper cheeks

4.6.3 Rear Fender Modifications

In the slices of the MCF field in the z direction, there was a prediction of positive drag contribution arising from the rear fender, behind the rear tires. In automotive applications, trip features tend to be added to these areas to ensure that the wake behind the car is made turbulent. A more laminar, oscillating wake tends to induce more drag behind the vehicle than a turbulent wake due to the pressure fluctuations. As such, a number of rear fender infill components were designed to create this trip effect. Two widths were tested, with two different infill heights per width, for a total of four different concepts. The taller of the two heights was designed to be taller than the region of positive drag contribution to determine the accuracy of the MCF prediction. The narrow width infills were designed to be closer to the fender than the MCF suggested an infill should be. These geometries are presented in Figure 4.12.




Figure 4.12: Rear fender geometry, baseline and four infill options

The results of the rear fender infills are presented in Table 4.3. The trend from these data was that the best-performing designs were those that were wider and shorter, matching the prediction from the MCF results. For both the short and tall infills, the narrow width concepts tended to generate more than 4% more drag relative to their baseline, when compared to their counterparts. The drag improvements predicted for Infill Option 2 stemmed from a reduction in drag on the rear fascia and the rear wheels. All four models tended to increase

the drag on the front wheels, however. In the case of Option 4, a reduction in body and rear wheel drag was predicted, but the gain in drag on the front wheels offset this reduction in drag.

Table 4.3: Rear fender infill designs and their drag contribution

Infill Option	Description	ΔC_D [%]
1	Narrow Width, Half Height	+3.8
2	Wide Width, Half Height	-1.6
3	Narrow Width, Full Height	+9.9
4	Wide Width, Full Height	+1.6

Figure 4.13 presents a comparison in the pressure drag coefficient on the rear fascia of the car, for the baseline and Infill Option 2 designs. A noticeable reduction in pressure drag occurred. Additionally, the pressure drag on the rear fascia was more symmetric, suggesting that there was weaker vortex shedding within the wake behind the car. This effect is shown in Figure 4.14 by examining the total pressure coefficient in the wake of the car. In the left image, the wake of the baseline fender infill is more asymmetric than the right image. This suggests that there are larger vortices being shed by the car without the trip features, generating larger pressure fluctuations. These fluctuations are then reduced with the introduction of the fender modifications, generating a more uniform wake structure, with equal counter-rotating vortices stemming from the separation of the rear fenders. The rotation direction can be seen by the line integral convolution on the $z = 310$ mm plane for Infill Option 2.




Figure 4.13: Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 2

The predictions by the MCF in the rear fender region of the car generated a reduction of drag, with the highest performing concepts being those that were placed most consistently within the positive region predicted by the MCF. The narrower and taller options that were designed to deviate from the MCF prediction, performed worse in CFD, showing the efficacy of these MCF predictions to drive design modifications, even in areas of separated flow.




Figure 4.14: Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 2

4.6.4 Roof Spoiler Modifications

The MCF around the rear spoiler suggested that the momentum around the spoiler was an area of negative drag contribution (this momentum decreases the drag). As such, increasing the momentum as the flow comes off the spoiler should be an appropriate method to reduce drag. To test this hypothesis, four spoiler variations were generated. Two were focused on increasing the curvature of the spoiler at its trailing edge, thus accelerating the flow prior to departing the roof. The second of these two modifications also featured a fairing around the D-pillar of the car to help provide a cleaner trip surface for the flow in this area. The increased curvature was achieved by translating the trailing edge rearwards and downwards from its original position. The other two designs were a simple tangent extension of the roofline along car centreline in two increments of 50 mm. These tested if a longer roofline would help increase the energy injected into the car's wake, reducing the base pressure behind the car. The spoiler modifications examined are presented below in Figure 4.15.

A summary of the performance changes from the spoiler modifications is presented in Table 4.4. The only design that predicted an improvement in drag was the second spoiler option, which increased the roofline curvature with an extension on the D-pillar. Spoiler Options 1 and 2 both predicted a decrease in drag on the body of the car due to the additional curvature of the spoiler. Notably, Spoiler Option 1 predicted an increase in drag on the front wheels, which was the result in the net drag gain. The simple tangent extension of the roof spoiler in Spoiler Options 3 and 4 both predicted an increase in the drag on the body of the vehicle. This suggests that the increased curvature was necessary to accelerate the roof flow before departing from the spoiler's trailing edge.

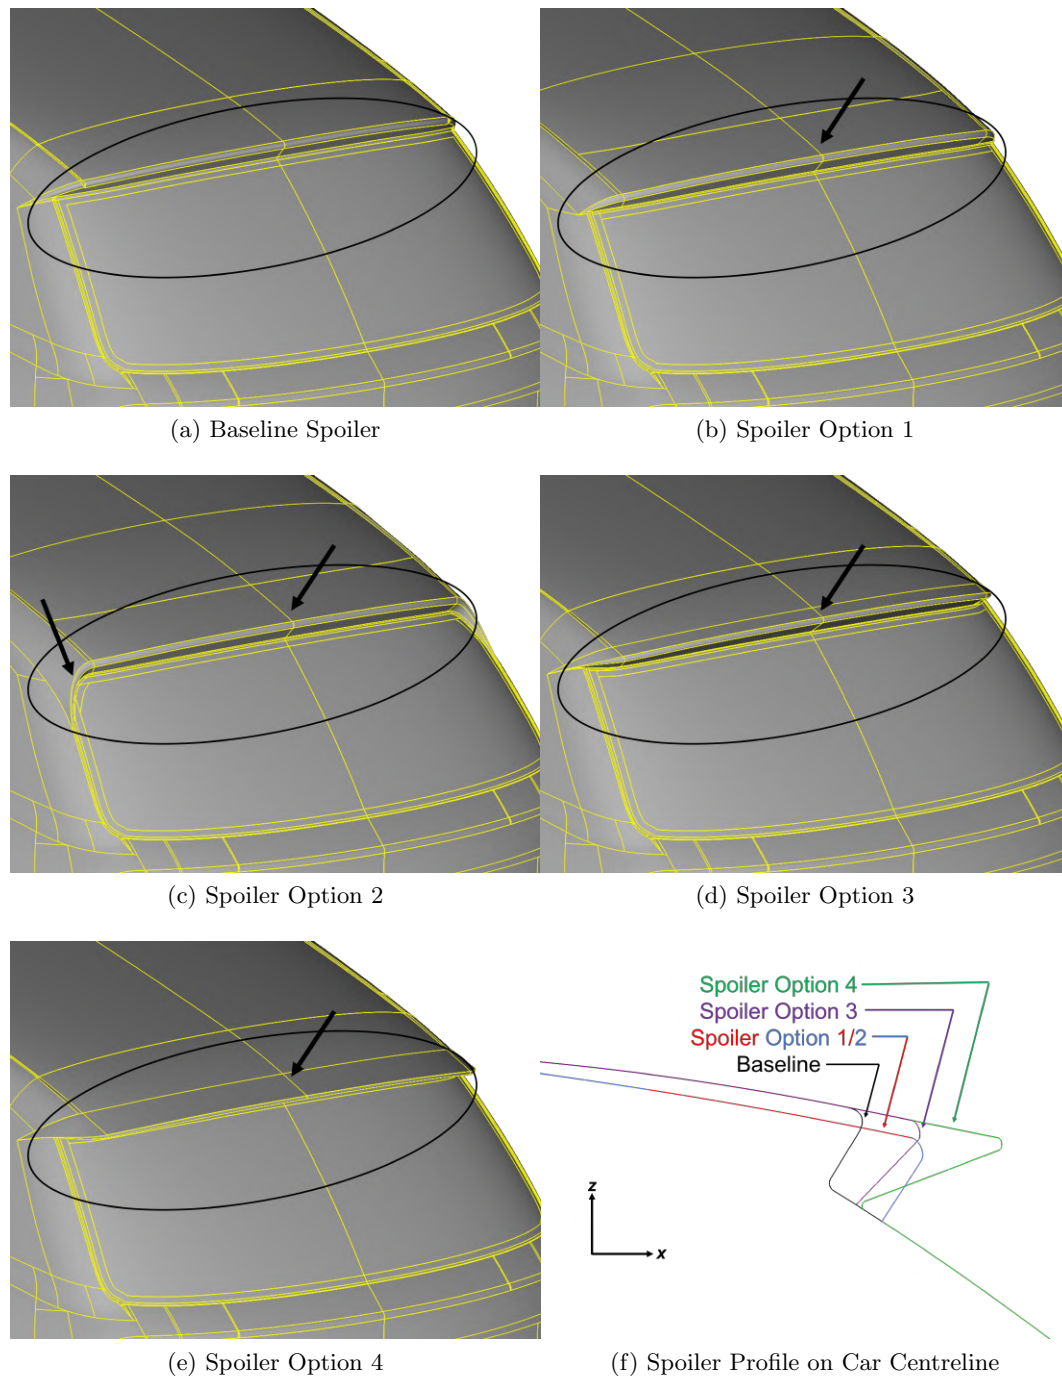


Figure 4.15: Baseline roof spoiler geometry, along with the four proposed trailing edge modifications and the spoiler profile on car centreline (Spoiler Options 1 and 2 share the same profile)

Table 4.4: Summary of rear spoiler modifications and their performance

Spoiler Option	Description	ΔC_D [%]
1	Downwards Spoiler Extension	+4.0
2	Downwards Spoiler and D-Pillar Extensions	-1.0
3	50 mm Tangent Spoiler Extension	+3.3
4	100 mm Tangent Spoiler Extension	+6.1

Investigating the pressure drag coefficient on the rear fascia of the car showed the effect that Spoiler Option 2 had on the flow field, presented in Figure 4.16. As with the rear fender infills, modifying the roof spoiler had the effect of reducing pressure drag. This was most noticeable on the rear window of the car, which required the greatest decrease in suction. This means the spoiler was effectively adding momentum into the wake behind the car, increasing the base pressure as the MCF had suggested. Looking at the normalized velocity along the $y = 0$ mm slice (car centreline) in Figure 4.17 shows that the local airspeed increased over the roofline. Therefore, the lowered roof spoiler is accelerating the roof flow as desired, more effectively filling in the void behind the car.




Figure 4.16: Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 2

For the roofline area, the MCF provided further correct guidance to reduce drag. When designing modifications that improved drag, it was important to guide the design based on the information of the high negative MCF area, rather than the area downstream that was around zero contribution. The MCF results were most effective in areas where the MCF had a large magnitude, either positive or negative. This could arise due to modeling error in the CFD code, which could result in the near-zero contributions not necessarily having the correct sign. Despite this, the roof spoiler was another area where the most effective change came from exploiting the regions of high MCF magnitude.




Figure 4.17: Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 2

4.6.5 Final Design

Through modifications of the splitter, front bumper cheeks, rear fender, and roof spoiler of the *DrivAer Estateback*, an 8.8% improvement in drag was achieved using the expert-driven MCF method, compared to the baseline model. C_L went from -0.231 to -0.236 (downforce produced in both cases), representing a 2.2% increase in *downforce*. This modified *Estateback* geometry will now be referred to as the *improved* variant, which will be one of the subjects of the gradient-based optimization presented in Chapter 5. Using the *improved* variant for the gradient-based optimization will allow for comparisons of the effect of optimization alone, and optimization paired with the MCF method. This will highlight the importance of expert-driven modifications in automotive design routines.

Figures 4.18 and 4.19 present the comparisons of pressure drag coefficient between the *base* model and the *improved* design. The key improvements in this model stem from the improved suppression of the tire wake, reducing tire drag, as well as the reduction in suction on the rear fascia. It was found that the most effective designs were those that directly followed the guidance provided by the MCF. This was particularly seen in the rear fender infill and roof spoiler.

The centre of pressure shifted from 75.1% of the downforce on the front axle to 92.2% from the *base* to the *improved* variant. This was due to the additional suction generated at the front of the underbody for the *improved* variant, shown in Figure 4.9. With a centre of pressure like this, neither the *base* nor the *improved* variants would be stable dynamically. However, should these variants be made more realistic with underbody panel gaps and tire air deflectors, there would be less suction at the front of the underbody. Therefore, these variants of the *DrivAer Estateback* were insufficiently effective for determining the stability of the platform based on centre of pressure metrics.

Nonetheless, this work highlighted that despite employing a RANS-based MCF, it was still able to yield improvements in aerodynamic drag. However, the improvements proposed in Section 4.5 could have yielded information that allowed the car's drag to be decreased further. This was not explored however, and is considered outside the scope of this research.




Figure 4.18: Front isometric comparison of pressure drag coefficient between the baseline and *improved Estateback* models




Figure 4.19: Rear view comparison of pressure drag coefficient between the baseline and *improved Estateback* models

4.7 Limitations of the RANS MCF

Section 4.5 highlighted two areas where the MCF suggested improvement could be made, the trailing edge of the hood, and the diffuser profile. These two areas had a number of complex interactions that can be difficult to resolve using a RANS simulation approach, as employed in the MCF methodology. As such, the initial examination of the MCF-guided designs that did not reduce the car's drag. Select designs in each area will be presented to discuss how the MCF information could be reinterpreted to produce designs that did reduce the drag of the *Estateback*.

4.7.1 Hood Modifications

From the drag MCF presented in Figure 4.1, there was a suggestion of drag reduction where the flow on the hood separated at its trailing edge, and reattached on the windshield, causing a pressure buildup. As such, a flick of various heights (3 mm, 6 mm, 10 mm) at the trailing edge was added in an attempt to divert this air further up the windshield, reducing the size of this pressure buildup region. The geometries for these flicks are shown below in Figure 4.20.




Figure 4.20: Hood geometry, baseline with the three trailing edge flick geometries tested

CFD computations then predicted that none of the designs reduced drag, since the flick led

to higher velocities downstream of the hood. It resulted in less flow acceleration along the hood, which had the effect of reducing the thickness of the boundary layer at the hood's trailing edge. This in turn led to an increase in flow speed at the trailing edge of the hood. This increased impingement on the windshield, yielding an increase in pressure on the front windshield. A sample of this pressure increase is shown below in Figure 4.21 for the 3 mm tall trailing edge flick.




Figure 4.21: Pressure coefficient around the windshield for the baseline hood trailing edge, and the 3 mm tall trailing edge flick

The guidance from the MCF in this location was that the flow velocity around trailing edge of the hood increased the drag of the car, and moving the flow velocity upwards would have yielded less drag. While gentle flicks tangent to the hood surface were tested, it may have been more beneficial to use a tab perpendicular to the hood's surface. This would have allowed for the airflow to have been given more upwash, while still permitting boundary layer growth that reduced the velocities at the hood's trailing edge. Since this region has a number of downstream interactions: the windshield and A-pillars, accurate sensitivity information is paramount. As such, using the improvements detailed in Section 4.5 will help improve guidance in this area.

4.7.2 Diffuser Modifications

The MCF information around the rear diffuser, shown in Figure 4.1 shows there is a region of negative MCF along the diffuser just off the surface, ahead of the trailing edge. Downstream of this is a large positive MCF region, caused by the diffuser flow creating a shear layer off the rear bumper. This MCF information was therefore taken to advise that reducing the velocities in this area could be beneficial, which could be achieved with a larger radius diffuser profile. Having less curvature yields less flow acceleration, reducing the strength of the shear layer off the rear bumper, reducing the local flow velocities around this region. The diffuser concepts tested are presented in Figure 4.22.




Figure 4.22: Diffuser geometry, baseline with the three modifications to the diffuser profile tested

All of the diffusers tested increased drag on the car. Examining the third diffuser iteration in Figure 4.23 helps explain why. In general, the diffusers resulted in a higher peak velocity underneath the car, but a smaller region of high velocity. As such, the underbody shear layer's strength was reduced. However, this imbalanced the counter-rotating vortex system between the roof shear layer and the diffuser shear layer. When these vortex systems are imbalanced, there tends to be a greater fluctuation in pressure on the rear fascia, which yields a lower time-averaged base pressure, increasing drag.

While the MCF did inform that the flow from the diffuser contributed to the drag on the car, it was not able to correctly inform how modifying the velocity in this area should be




Figure 4.23: Normalized velocity contours on the $y = 0$ mm plane for the baseline and Diffuser Iteration 3

balanced with the rest of the structures in the rear wake. This area was particularly sensitive to the interactions with neighbouring regions around the rear fascia, and the shear layers they generated. What could have been a more effective modification based on the MCF would be to have a trip feature on the diffuser generating early separation. This would have decreased the flow velocity in this region due to the separation, while still creating a strong vortex system to balance the flow coming off the roof and the rear fenders. While this information was available within the MCF, it was not as immediately apparent as the MCF guidance was for the roof spoiler or the rear fender infills. As such, in areas of high interaction, a larger number of iterations may be required to use the MCF information effectively. Additionally, it may be more effective in some areas to incorporate trip features that will decrease the downstream flow velocity and energy, rather than just diverting air away from these regions.

In general, the MCF tended to perform better in regions where a negative sensitivity was being exploited to reduce the objective function. This resulted in a need to increase flow velocities, rather than divert air away from the region. In the case of the rear fender infills, this came from adding in material towards the edge of the negative MCF region along the rear fenders. In areas where high positive MCF needed mitigation, design modifications were less effective. It becomes evident that areas of high positive MCF tend to be involved in more flow interactions with other components on the car, where simply slowing down the local flow velocity or moving air away from the positive MCF region may not improve the objective function. These are areas where supplementing the MCF with traditional methods, such as examining the velocity, pressure, and total pressure fields will be more effective than the MCF alone.

5 Drag Shape Optimization of the *DrivAer Estateback*

Results from the MCF-derived modifications to the *DrivAer Estateback* showed that drag could be reduced. However, tested modifications were relatively simple geometrically, and focused on a few key design changes in each area that were isolated from the MCF. Therefore, making use of freeform deformation boxes within a gradient-based optimization loop could yield more detailed deformation of the surface, which could reduce drag further.

This chapter makes use of the baseline *DrivAer Estateback* model (referred to as the *base* variant), as well as the *improved Estateback* variant, as presented in Chapter 4. Using these two variants, gradient-based optimization will be performed on the roof spoiler and front bumper cheek areas, using three different unconstrained gradient-based optimizers. These two locations were targeted because of the variation in flow physics between the roof spoiler and front bumper cheek. Both the *base* and *improved* variants are used to examine the effect of the initial design vector on the optimization results, as well as to understand theoretical limits of improvements to the *Estateback*. This is to evaluate the potential efficacy of a unified design approach, combining initial MCF-driven design modifications, followed by computer-driven optimization.

To expedite the optimization simulations, a half-car model was employed, since the *base* and *improved* variants of the *Estateback* were symmetric. The effect of a half-car on the CFD predictions will be discussed first in this chapter. Then optimization methods, problems, and results will be discussed, alongside post-optimality CFD of the optimized designs. This chapter is a detailed investigation into the effect of initial design vector and optimizer algorithm on optimization results for automotive applications, with design improvements verified by post-optimality CFD.

5.1 Half-Car Models

This section investigates the effect modelling the *Estateback* as a half-car symmetric DDES model has on the drag and flow field predictions. These results are compared to a full-car DDES modelling approach for both the *base* and *improved* variants. This will detail areas in which the post-optimality CFD predictions will be affected by the numerics enforced by the symmetry condition.

5.1.1 Drag Comparisons

The symmetry plane was placed at $y = 0$, with a symmetry condition imposed on all fields, which forced gradients perpendicular to the wall to be zero. Therefore, this limited the effects of periodic vortex shedding, since these vortices can no longer migrate across the symmetry plane. This could have the greatest significance on two of the modifications that were made to the *improved* variant: the rear fender infills and roof spoiler, which both sought to reduce asymmetries in the wake. Since this effect was artificially removed due to the symmetry plane, it is expected that the *improved* variant will not reduce as much drag as was predicted by the full car models. Table 5.1 highlights the effect the half-car simulations had on the predicted drag of each variant.

Table 5.1: Comparison of drag predictions for both *Estateback* variants for full-car and half-car simulation approaches

Variant	$C_{D,Full-Car}$	$C_{D,Half-Car}$	ΔC_D [Counts ($C_D = 0.001$)]
<i>Base</i>	0.295	0.293	-2
<i>Improved</i>	0.269	0.285	+16

For the *base* variant, the difference between the full-car and half-car models is small, with the drag being 0.7% lower for the half-car model. For the *improved* variant, this trend is greatly different, with the half-car model predicting 5.9% more drag than the full-car model. The higher drag prediction for the half-car model on the *improved* variant stemmed from higher drag being predicted on the upper and lower body surfaces, as well as on the front tires. Since the *improved* model had the front bumper cheek flicks, this would alter the tire wake region, and work to create a larger area of separation upstream of the tire. This wake structure had some cross-body oscillations, which could not be resolved by the half-car model. Therefore, it is expected that some of the advantage predicted for the front bumper cheek flick was not realized by the half-car model.

5.1.2 Surface Pressure Drag Comparisons

Investigating pressure drag for the full- and half-car models helps highlight where discrepancies in the flow field lay, as well as why the *improved* half-car variant is predicted to have considerably higher drag than the full-car model. Figure 5.1 shows the pressure drag coefficient for the two modelling approaches, for the both variants of the *Estateback*.




Figure 5.1: Surface pressure drag coefficient on the full-car and half-car (mirrored for clarity) *Estateback*

The effect of the symmetry plane can be clearly seen on the rear fascia, where there is less pressure drag on car centreline. This supports the idea that there were cross-body fluctuations in the wake that are not getting captured by the half-car model. These oscillations would create a lower base pressure, increasing pressure drag on the rear fascia for the full-car model. For the *base* variant, this had the effect of reducing drag on the rear fascia, which explains the corresponding decrease in predicted drag by the half-car model. Considering at the results for the *improved* variant in Figure 5.1 show a different behaviour.

As with the *base* variant, the effect of the symmetry boundary condition on the *improved* variant can be clearly seen on the rear fascia, denoted by a reduction in pressure drag coefficient. However, the *improved* variant predicted there would be more pressure drag on the rear fascia for the half-car model, than for the full-car model. Furthermore, there was increased pressure drag on the rear fender infill surfaces. The rear fender infill and spoiler modifications were both shown to help reduce asymmetries in the rear wake, but that did not mean there were no cross-body oscillations present in these designs. Therefore, not having the full flow field modelled resulted in these modifications not contributing as much to drag reduction. Furthermore, there is more pressure drag on the rear of the tires, likely due to a phenomenon

not being captured at the front of the car, where the front tires generate cross-body flow fluctuations.

5.1.3 Half-Car Modelling Conclusions

Despite the inaccuracies from the half-car modelling approach, it was still deemed acceptable to use these models for optimization. The regions chosen for optimization should be unaffected by the symmetry plane, since their flow fields are not driven by upstream, cross-body fluctuations. For the front bumper cheeks, the initial flow comes from the stagnation region on the car, which contains mostly steady flow features. For the roof spoiler, the roof flow is also mostly steady, with cross-body fluctuations occurring in the wake. While the half-car models would not resolve these features, drag changes across different designs would be resolved. This would not have prevented the optimizer from reducing drag for both variants of the *Estateback*. Therefore, since both regions upstream flow conditions were mostly symmetric, a half-car approach was deemed suitable.

For applications that have large regions of cross-body flow, it would be more accurate to use a full-car model to capture any oscillations. However, given computational costs, this was not feasible for this research. Nevertheless half-car methodology, the optimization algorithms and problems could still be defined.

5.2 Gradient-Based Optimizers

For all gradient-based optimizers, the derivative of the objective function is used to define a search direction to either maximize or minimize the objective function, J . For CFD solutions, an efficient method with which the Jacobian (∇J) and the Hessian ($\nabla^2 J$) could be approximated was the adjoint method, as previously described in Chapter 4.

When developing the aerodynamic surfaces of a car to reduce its drag, there tend to not be any aerodynamic constraints. Some aerodynamic considerations such as lift balance play a role in the dynamic behaviour of a vehicle, this was noted in Chapter 4 that the *Estateback* with a simple floor would be unsuitable for effectively evaluating these types of metrics.. As such, unconstrained gradient-based optimization is sufficient for automotive applications, and they are easier to converge than constrained methods. In contrast, aerospace applications require constrained optimizers for aerodynamic development, as there are a number of performance requirements, such as lift and pitching moment, which constrain drag optimization. In the case of the unconstrained gradient-based optimizers available in *OpenFOAM*, freeform deformation box constraints and total displacement constraints are possible, through the choice of active design variables.

Should stability metrics be of interest, these can be incorporated into the objective function by considering lift balance or yaw moment about the car's centre of gravity when experiencing some side slip at an operating point with a small sideslip angle, which is not compatible with a half-car model. These could have been implemented as a multi-objective function, where a weighted sum of drag and these stability metrics should be minimized, or they could have been considered as constraints. Since these stability metrics would have required a full-car

model for the optimization cases, these stability metrics were out of the scope of this research.

Unconstrained gradient-based optimizers in *adjointOptimisationFoam* are applied in the following steps, with an initial major iteration, $K = 0$, and an initial design vector: \mathbf{x}_K :

1. Convergence is tested, by evaluating the objective function value at the current design vector, \mathbf{x}_K , and comparing it to the previous major iteration's objective function value (if available). If convergence criteria are met, the solution is stopped.
2. A search direction, \mathbf{p}_K , is computed using gradient information from the objective function. A number of search direction methods can be chosen to compute this vector.
3. Compute the step length to find a positive scalar, α_K , such that the *Armijo* conditions are met as implemented in *adjointOptimisationFoam* [50].
4. The design vector is updated such that, $\mathbf{x}_{K+1} = \mathbf{x}_K + \alpha_K \mathbf{p}_K$.
 $K = K + 1$, and steps 1 - 4 are repeated if K is less than the maximum number of desired iterations.

As mentioned, *adjointOptimisationFoam* employs the *Armijo* condition for the line search, which is an inexact line search method. It is formulated such that the line search converges when Condition 5.1 is satisfied:

$$J(\mathbf{x}_K + \alpha_K \mathbf{p}_K) \leq \mu_1 \alpha_K \nabla J(\mathbf{x}_K)^T \mathbf{p}_K \quad (5.1)$$

In Condition 5.1, μ_1 is a small scaling factor, which is 10^{-4} by default in *adjointOptimisationFoam* [50]. This ensured that some minimum improvement is required for the line search to have converged. It should be noted that the line search in *adjointOptimisationFoam* can be limited to a certain number of iterations, which could result in line search termination before an improvement is found. As such, this implementation within *adjointOptimisationFoam* was not guaranteed to result in an improvement in the objective function at each major optimization iteration. For all optimization simulations in this work, the line search was allowed to run for up to five function evaluations before terminating, to balance computational expense with finding an improved solution.

Since the line search method implemented does not make use of any curvature conditions, each subsequent iteration of the line search is computed through backtracking by the following method:

1. The initial step length $\bar{\alpha}$ is set to one, and some reduction ratio, $0 < \rho < 1$ is chosen
2. If the *Armijo* condition outlined in Equation 5.1 is met, $\alpha_K = \bar{\alpha}$, and the line search is terminated
3. If the condition is not met, set $\bar{\alpha} = \rho \bar{\alpha}$ and repeat step two until convergence is met, or the maximum number of line search iterations is reached

In this paper, three search direction computation methods are explored: Conjugate Gradient (CG), the Broyden-Fletcher-Goldfarb-Shanno (BFGS), and the Damped-BFGS (DBFGS) algorithm.

5.2.1 Conjugate Gradient

The Conjugate Gradient method is a modification to a steepest descent method, whereby the gradient search direction is *pushed* towards the solution more rapidly by using previous

optimization iterations' gradient information [51]. It can be formulated by assuming that the objective function space can be represented by a quadratic function:

$$J(x) = \frac{1}{2}x^T Qx - x^T b, Q = Q^T \quad (5.2)$$

Therefore, the gradient of the objective function becomes:

$$\nabla J = Qx - b \quad (5.3)$$

Using this function and its gradient, when applying the typical steepest descent search direction, which is the negative gradient, the search direction for a Conjugate Gradient problem can be formulated as:

$$\mathbf{p}_{K+1} = -\nabla J(\mathbf{x}_{K+1}) + \beta_K \mathbf{p}_K \quad (5.4)$$

The computation of β_K requires it to be Q-conjugate at every optimization iteration, K . To negate the need to compute Q , so the search direction is solely dependent on the gradient of the objective function, the Fletcher-Reeves approximation is applied. This approximation is given by [52]:

$$\beta_K = -\frac{\nabla J(\mathbf{x}_{K+1})^T \nabla J(\mathbf{x}_{K+1})}{\nabla J(\mathbf{x}_K)^T \nabla J(\mathbf{x}_K)} \quad (5.5)$$

This allows for curvature information to be approximated, without the need for computing the Q matrix.

5.2.2 BFGS

The BFGS method is a quasi-Newton method, whereby the Hessian matrix is approximated, to try and represent the curvature of the design space [52]. Since the inverse Hessian matrix is what will actually be needed, instead of solving for $H_K(\mathbf{x})$, the inverse Hessian, $V_K(\mathbf{x})$ is directly solved for at each major iteration, K . Assuming that a Taylor-series expansion of the gradient of a function can be written as:

$$\nabla J(\mathbf{x}_K + \mathbf{s}_K) = \nabla J(\mathbf{x}_K) + H_K \mathbf{s}_K + \dots \quad (5.6)$$

In Equation 5.6, $\mathbf{s}_K = \alpha_K \mathbf{p}_K$, and $\mathbf{y}_K = \nabla J(\mathbf{x}_K + \mathbf{s}_K) - \nabla J(\mathbf{x}_K)$. The Hessian can then be given by:

$$H_K \mathbf{s}_K = \mathbf{y}_K \quad (5.7)$$

Which then places a requirement for the inverse Hessian to meet the following criterion:

$$V_{K+1} \mathbf{y}_K = \mathbf{s}_K \quad (5.8)$$

Using this information, the inverse Hessian matrix is solved for by attempting to minimize the norm of the difference between the inverse Hessian at iteration $K + 1$ and the inverse Hessian at iteration K . The inverse Hessian is then solved using the following equation:

$$V_{K+1} = \left[I - \frac{\mathbf{s}_K \mathbf{y}_K^T}{\mathbf{s}_K^T \mathbf{y}_K} \right] V_K \left[I - \frac{\mathbf{y}_K \mathbf{s}_K^T}{\mathbf{s}_K^T \mathbf{y}_K} \right] + \frac{\mathbf{s}_K \mathbf{s}_K^T}{\mathbf{y}_K^T \mathbf{s}_K} \quad (5.9)$$

Using this approximation for the inverse Hessian matrix, the search direction for a BFGS problem can then be defined by:

$$\mathbf{p}_K = -V_K \nabla J(\mathbf{x}_K) \quad (5.10)$$

By using the approximated Hessian matrix, the BFGS algorithm leverages curvature information about the design space, leading to larger steps across the objective function topology. This makes the accuracy of BFGS dependent on how well-behaved the objective function's topology is, as well as how accurately the sensitivity derivatives can be computed, since it directly impacts the inverse Hessian calculation.

5.2.3 DBFGS

The Damped BFGS algorithm was developed to act as a more stable version of the BFGS algorithm [53]. It is similar to the presented BFGS algorithm, except it is derived from a BFGS algorithm that computes the Hessian directly and takes its inverse, rather than solving for the inverse Hessian problem that BFGS does. This makes the DBFGS algorithm more akin to the Davidon-Fletcher-Powell (DFP) algorithm, which also directly computes the Hessian, then inverts it to compute the new search direction. The damping of the Hessian is performed to maintain the positive-definiteness of the matrix, and occurs by examining the net curvature at the current design vector. For some user-defined scalar γ , if the following inequality is met, the solution will proceed in its damped form:

$$\mathbf{s}_K \mathbf{y}_K^T < \gamma \mathbf{s}_K^T H_K \mathbf{s}_K \quad (5.11)$$

Whether the solution is in its damped form or not, the solution of the Hessian matrix takes the same form as the DFP algorithm, being:

$$H_{K+1} = H_K - \frac{H_K \mathbf{s}_K \mathbf{s}_K^T H_K}{\mathbf{s}_K^T H_K \mathbf{s}_K} + \frac{\mathbf{r}_K \mathbf{r}_K^T}{\mathbf{s}_K^T \mathbf{r}_K} \quad (5.12)$$

The \mathbf{r}_K in Equation 5.12 is a damped version of the change in the gradient, \mathbf{y}_K . This damped term is computed by the following equation:

$$\mathbf{r}_K = \theta \mathbf{y}_K + (1 - \theta) \mathbf{s}_K H_K \quad (5.13)$$

When the inequality shown in Equation 5.11 is not met, θ is equal to unity, returning a typical DFP solution. When the inequality is met, θ is given by the following:

$$\theta = \frac{(1 - \gamma) \mathbf{s}_K^T H_K \mathbf{s}_K}{\mathbf{s}_K^T H_K \mathbf{s}_K - \mathbf{y}_K^T \mathbf{s}_K} \quad (5.14)$$

Through this process, the DBFGS algorithm helps maintain positive-definiteness of the Hessian matrix, but it still requires the Hessian to be inverted. This can become numerically ill-conditioned, resulting in inaccurate or poorly conditioned search directions. When applied to *adjointOptimisationFoam*, this could lead to extremely large deformations of the geometry, resulting in mesh quality violations. For the DBFGS method, the search direction is calculated as:

$$\mathbf{p}_K = -H_K^{-1} \nabla J(\mathbf{x}_K) \quad (5.15)$$

5.3 Freeform-Deformation Boxes

For shape optimization problems, freeform deformation boxes (FFD) are an efficient means of defining the design variables for an optimization problem. For a three-dimensional problem, each control point of an FFD has three design variables, for each of its three axes. Therefore, for an FFD with n points, it will comprise $3n$ design variables in three-dimensional space. An example of how an FFD can be used to displace a triangulated surface is shown in Figure 5.2. A similar approach is implemented in *OpenFOAM*.




Figure 5.2: Example surface displacement that can be achieved with an FFD, from [54]

Within *OpenFOAM*, the FFD method is derived from NURBS B-spline radial basis functions [50]. Rather than using a global coordinate system-aligned parameterization for an FFD, a local cartesian coordinate system is created, aligned to be orthogonal to the rectangular box, such that a new coordinate system with axes u , v , w is defined. The FFD is therefore comprised of Cartesian points in x , y , z that can then be parameterized in u , v , w space, with radial basis functions U , V , W . For *OpenFOAM*'s implementation of the FFD boxes, the box must be an orthogonal rectangular prism. An example of the u , v , w space for a non-global axis aligned freeform deformation box is shown in Figure 5.3.

The degree of each radial basis function U , V , W can then be described as pu , pv , and pw . For an FFD with a number of points in the u , v , w directions of: $I + 1$, $J + 1$, and $K + 1$, the m^{th} component of a mesh point, x , can be calculated by Equation 5.16.

$$x_m(u, v, w) = \sum_{i=0}^I \sum_{j=0}^J \sum_{k=0}^K U_{i,pu}(u) V_{j,pv}(v) W_{k,pw}(w) b_m^{ijk} \quad (5.16)$$

Using the initial mesh, a mapping is found for each mesh point $\mathbf{x}(x, y, z)$ in u , v , w space. Using the known knot vectors $\mathbf{r} = (x_r, y_r, z_r)$, each internal and boundary point of the finite volume mesh that lives within the FFD can be parameterized in u , v , w space by iteratively solving the following system of linear equations:




Figure 5.3: Side view of a non-global axis aligned FFD. The y/v directions are shared in this example, and pass into the page

$$\mathbf{R}(u, v, w) = \begin{bmatrix} x(u, v, w) - x_r = 0 \\ y(u, v, w) - y_r = 0 \\ z(u, v, w) - z_r = 0 \end{bmatrix} \quad (5.17)$$

This mapping is independent for each mesh point and as such, the B-spline implementation within *OpenFOAM* solves for these parameterization using parallel computations for decomposed problems [50]. At the end of each optimization iteration, displacements are computed for each point within the FFD(s). Once the updated points b_m^{ijk} are known, Equation 5.16 is used to update all of the mesh points within the FFD. When using *adjointOptimisationFoam* in double-precision, these parameterizations can be calculated to machine precision, allowing for accurate deformations of complex geometries.

When FFD points are used as the design variables of an optimization problem, a search direction will be computed at each point, in each of its three directions. Once a step length is then computed from the line search, each control point is displaced. This displacement is then mapped back to the mesh points within the FFD, which are displaced according to the NURBS radial basis functions.

5.4 Numerical Setup

This section describes the specific setup operations employed within *OpenFOAM's* optimization solver. As with the MCF evaluation, all optimization models make use of the same meshing approach and boundary conditions, just modified to yield a half-car model.

When performing an automated optimization loop in *adjointOptimisationFoam*, a RANS primal solution is performed to compute the primal velocity, pressure, and turbulence fields. For bluff bodies with significant areas of areas of separation, RANS models tend to underpredict the amount of separation along the body, but give reliable changes in results between different designs. Both the primal and adjoint RANS equations are solved using the *SIMPLE* algorithm, with the RANS equations closed using the SA turbulence model. Averaging of the primal fields and objective function occurs over the final 1000 iterations of each objective function evaluation.

Since *adjointOptimisationFoam* stores the previous optimization iteration’s fields to initialize the current optimization iteration (other than the initial function evaluation), fewer *SIMPLE* iterations are needed to find a suitable solution for the flow field. As such, the initial function evaluation solves for 2500 iterations, with each subsequent function evaluation solving for 1500 iterations. This allows for accelerated computation of each objective function evaluation, reducing the total cost of the simulation. The adjoint equations are solved for 1000 iterations every time the continuous adjoint fields are computed.

RANS models can be insufficient for the prediction of automotive aerodynamics. As such, DDES are used on select geometries from the various optimization runs to perform post-optimality CFD simulations. The deformed mesh from the optimization runs are used for these simulations, reducing the grid-dependence on the DDES simulations. The DDES setup was described in Section 3.5.

5.5 Optimization Problem

For a given FFD around a portion of the geometry, the optimization problem will be formulated as:

$$\begin{aligned} \text{minimize} \quad & C_D \\ \text{w.r.t.} \quad & \mathbf{x} = (x_{i1}, x_{i2}, x_{i3}, \dots, x_{n3})^T \end{aligned} \tag{5.18}$$

In Equation 5.18, x_{i1} represents the u -component of the i^{th} active design variable, x_{i2} represents the v -component of the i^{th} active design variable, and x_{i3} represents the w -component of the i^{th} active design variable. While minimizing C_D^2 may improved the conditioning of the optimization problem by enforcing additional convexity in the design space, it would have required the derivation and coding of a unique force squared objective function within *adjointOptimisationFoam*. Deriving this objective function and its sensitivity derivatives was outside the scope of this research.

When defining a deformation box within *OpenFOAM*, the user has the ability to deactivate certain design variables within the box. It is typically required to constrain all of the design variables on the boundary of an FFD, which prevents the deformed mesh from moving through points that are outside the FFD. Since a half-car model is used for the optimization, any FFD points that pass through the symmetry plane are only constrained in the parameterized v/y direction. This allows for motion in the symmetry plane, while constraining plane-normal displacements, representative of a symmetry deformation constraint.

For this problem, two different FFD's were evaluated, one around the front bumper cheeks, and one around the rear spoiler. The positions of which are highlighted in Figures 5.4 - 5.5. The FFD's are coloured by their activation of the u and w degrees of freedom (red for active, blue for inactive). This shows that on the symmetry plane of the vehicle, these points of the FFD are only constrained in the y/v direction. As previously described, this means these points cannot move through the symmetry plane, but they can move on the plane, ensuring a proper symmetry displacement constraint.

The front bumper cheek FFD was made to be relatively wider than the surface geometry to allow for sufficient deformation of the surface, without artificially constraining the results. These boxes were evaluated for each of the two *Estateback* variants, for each of the three optimization algorithms described in Section 5.2.




Figure 5.4: Freeform deformation box placed around the rear spoiler. The car is shown in its symmetric state to highlight the FFD points on the symmetry plane




Figure 5.5: Freeform deformation box placed around the left front bumper cheek

The number of points in the u , v , w directions as well as the degree of the NURBS radial basis function for each FFD is shown below in Table 5.2. All directions were defined with a third degree NURBS radial basis function. This ensured that there can be curvature continuity within the deformed surface, while preventing artificial damping of the deformations due to

a high-degree radial basis function.

Table 5.2: Number of points and degree of the NURBS radial basis function in the u , v , w directions for each FFD

FFD	u Points	v Points	w Points	Degree U, V, W
Rear Spoiler	6	9	7	3
Front Bumper Cheek	6	7	10	3

Convergence of the optimization problem is monitored for both the L2-norm of the gradient, as well as the change in objective function value for successive iterations. These criteria are defined as follows:

$$\|\nabla J(\mathbf{x}_{\mathbf{K}})\| \leq \epsilon_g \quad (5.19)$$

$$|J(\mathbf{x}_{\mathbf{K}+1}) - J(\mathbf{x}_{\mathbf{K}})| \leq \epsilon_{abs} + \epsilon_{rel}|J(\mathbf{x}_{\mathbf{K}})| \quad (5.20)$$

The convergence criterion in Equation 5.19 can be met at each iterations, whereas the convergence criterion in Equation 5.20 must be met for two successive major iterations before it is considered as having been met. The solver reaching either of these two convergence metrics is sufficient for the optimization loop to be concluded. For all models, ϵ_g was set to 0.001, ϵ_{abs} was set to 0.0005, and ϵ_{rel} was set to 0.0005. Since a half model is employed, this means the relative change in objective function should be $C_D = 0.001 + 0.1\%$ of the previous iteration's value. This limit was placed as drag improvements of less than 0.001 C_D are typically considered to be within the error of the CFD simulation, for automotive purposes.

In some cases, poor numerical conditioning of the gradient at some control points of the FFD results in extremely large deformations of the mesh, violating the mesh quality criteria. This can be caused by from a number of reasons. Some of the control points can live within the geometry, therefore not being a part of the active flow region. Mapping the external flow solution to these points can then result in large gradients, due to large flow gradients that are present around the mesh points that influence the control point(s). Furthermore, scaling of the design variables only occurs at the first optimization iteration. For complex industrial problems, not scaling the design variables more frequently can yield large displacements, for which the local search direction is an inaccurate approximation over such a large step. Furthermore, poor conditioning of the Jacobian can be caused by the aforementioned numerical convergence issues of both the primal and adjoint solvers. In the previous chapter, inaccuracies in the sensitivity analysis were not quantitatively identified, since the adjoint information was used qualitatively, with engineering judgment employed in its interpretation. The optimization algorithms used do not have this same judgment and are susceptible to these convergence issues.

These mesh quality violations can cause the following optimization iteration to fail, since the simulation cannot compute a proper solution. In these cases, an *improved*, rather than *minimized*, solution can be isolated prior to the mesh quality violation occurring. While this is not strictly an optimized geometry, it can still be valuable from an industrial standpoint. The main purpose of applying shape optimization to automotive applications is to search for further improvements on the design. If an improvement was realized, this could still be useful for a car's development, despite it not being a rigorously optimized geometry.

5.6 Roof Spoiler Optimization

The results of the roof spoiler optimization will be presented by first investigation the evolution of the objective function (drag) as a function of optimization iteration, K , for the *base* and *improved* variants of the *Estateback*. These results of each variant will then be compared to each other, and the resultant geometry deformations will be compared.

5.6.1 Base Variant Results

The evolution of aerodynamic drag normalized to the initial function evaluation as a function of optimization iteration is presented in Figure 5.6 for the three optimization algorithms. It is noteworthy that the DBFGS did not make it past the first optimization iteration. This means that the first iteration for which the inverse Hessian matrix was calculated, resulting in a poorly-conditioned search direction, yielding non-physical mesh deformations, violating the mesh quality. This poor conditioning was likely due to the inversion of the Hessian matrix, which was discussed in Section 5.2.3. Since the DBFGS algorithm yielded mesh quality violations after the first optimization iteration, all three optimization algorithms would have had the same Jacobian vector calculated, because they all perform a steepest descent evaluation for the first optimization iteration. Since neither the CG, nor the BFGS algorithm yielded mesh quality issues, and *OpenFOAM* did not report the damping term being activated for the DBFGS solver, the problem in the search direction had to stem from the inversion of the Hessian matrix performed by the DBFGS algorithm. This indicates that the problem was poorly-conditioned in nature, and was non-convex/multimodal.




Figure 5.6: Optimization results for the three optimization algorithms. The objective function is normalized to the drag of the *base* variant at the initial optimization iteration

The BFGS trace is of interest, as it results in an increase in the objective function at the second optimization iteration, as well as at the fourth and fifth optimization iterations. This increase is feasible, due to the maximum number of iterations that are specified for the line search. It is also interesting to note that the solution converged, despite it being to a value greater than the initial objective function value. This suggests that there was some degree of multi-modality in the objective function's topology, with this being a local minimum. This multi-modality will be further discussed in the results of the *improved* variant's optimization results.

To understand if the problem with the BFGS solution was due to the line search, this case was repeated with a maximum of ten line search iterations, instead of five. By increasing the maximum number of line search iterations, the objective function was not forced to decrease at each iteration, but when it did increase, it was on the order of half a drag count. In this optimization however, the model had not converged, having failed on the fifth optimization iteration, due to mesh quality violations. The progression of this optimization run compared to the initial Conjugate Gradient and BFGS runs is shown in Figure 5.7. This figure shows that while additional line search iterations improved the BFGS algorithm's predicted improvement, it was still not as effective and required more function evaluations than CG. Furthermore, additional line search iterations did not improve the robustness of BFGS.




Figure 5.7: Optimization results for the 10-line search iteration BFGS case, compared to the initial CG and BFGS cases

Moreover, the fact that an improvement could not be achieved, suggests that there was poor conditioning of the Hessian matrix for this problem. This was likely due to multi-modality within the design space, despite areas of local convexity. Since the Conjugate Gradient method converged asymptotically towards a local minimum, this shows the local convexity is present,

which was able to be found because of small total displacements. The BFGS algorithm caused larger displacements, which led to it finding regions of local concavity, highlighted by the increase in drag in some objective functions. That being said, an improved design vector was predicted at iteration 3 for the BFGS algorithm, suggesting an improvement of 1.2%. This design could still be valuable in terms of directional understanding of how to improve the spoiler region of the car.

The Conjugate Gradient trace shows that there was constant improvement in the objective function, and a converged solution could be found. A maximum improvement in the objective function of 7.0% was predicted from the optimization algorithm, relative to the initial design vector. Since the Conjugate Gradient method largely relies on gradient information, with slight "steering" of the search direction from the β term, it is less sensitive to numerical error in the sensitivity derivatives, and will be a more robust method. The Conjugate Gradient method takes longer to yield an improvement since it will suggest smaller displacements at each optimization iteration, but this is favourable in the presence of a more complex objective topology, as is clearly exhibited.

5.6.2 Improved Variant Results

The evolution of aerodynamic drag normalized to the initial function evaluation, as a function of optimization iteration is presented in Figure 5.8, for the *improved Estateback* variant. As with the *base* variant, the DBFGS optimizer failed after the first optimization iteration, which would be due to a poorly conditioned search direction. It is evident that for industrial cases, the inversion step of the DBFGS algorithm yields numerical instabilities in the sensitivities.




Figure 5.8: Optimization results for each of the three optimization algorithms

The BFGS trace differed from the *base* variant, since its converged solution reduced the car’s drag. However, in the process the objective function was increased and subsequent iterations oscillated around this value, before converging to a solution with its drag reduced. This once again highlights that for industrial automotive applications, the BFGS algorithm is less robust, likely due to insufficient convergence of the primal equations, as well to the possibility of large displacements at each iteration, degrading the accuracy of the local Jacobian and inverse Hessian approximations. The fact that it converged for the *improved* variant and not the *base* variant is not indicative of the user-defined modifications resulting in a more favourable conditions objective function topology, but to rather to the change in the initial design vector that was sufficient to place the car in a more favourable region of the objective function’s topology, such that it could be guided towards a local minimum.

The Conjugate Gradient trace shows that a converged, improved objective function value could be reached. It is also valuable to note that this converged objective function value predicts greater improvement than the BFGS’ converged solution. As with the *base* variant, the Conjugate Gradient method was not only more robust, but it was also more effective at reducing the drag of the car. It is also valuable to note that the Conjugate Gradient optimization of the *improved* variant only resulted in a 4.7% improvement in drag, compared to 7.0% for the *base* variant. This is likely due to the user-driven modifications, that would have moved the initial design vector closer to the converged minimum. As such, a smaller total step was needed to reach that minimum, leading to a more accurate Jacobian prediction due to staying within the region of space in which the adjoint approximation is valid (convex).

The results of the spoiler optimization for the *improved* variant hint at more global multi-modality in the design space. Both the Conjugate Gradient and BFGS solutions yielded a converged solution, however the predicted drag reduction was different for both points. This suggests that the design space was not truly convex, which would likely extend to the *base Estateback* variant. These local minima were verified by performing an additional optimization run on the converged geometries. The repeat of the Conjugate Gradient model converged immediately due to insufficient change in the objective function. This showed that the Conjugate Gradient method found a local minimum, with less than 0.5% change in the drag when the optimizer was restarted using the new initial design vector. In the case of the BFGS repeat, it converged to a new solution after five major iterations, predicted an additional 2.5% reduction in drag. Since the BFGS mapping was capable of finding a new, stronger minimum than previously found, this further highlights the multi-modality of the design space, and suggests that the Conjugate Gradient algorithm is better suited to finding stronger minima in automotive applications.

This multi-modality is partly an artifact of the lack of wake resolution within the RANS solver, but is also likely due to the wake interactions between the roof flow and the underbody/rear fender flow. Minimizing with respect to C_D^2 could help improve the convexity of the design space to reduce some of these problems.

5.6.3 Comparing Optimization Results

This section details how the optimization of both variants compared to each other, normalized by the initial objective function value for the *base* variant, denoted by $J_{0,Base}$. The objective function evolution traces for each algorithm and variant are presented in Figure 5.9. As ex-

pected, the value of the objective function at optimization iteration one is the same for all algorithms for a given variant, since all three algorithms employed a steepest descent initial step, and were therefore limited to the same maximum initial displacement.




Figure 5.9: Optimization results for each of the three optimization algorithms, for both *Estateback* variants. Both variants’ drag is normalized by the drag of the *base* variant at the initial optimization iteration ($J_{0,Base}$)

Comparing the Conjugate Gradient traces for both *Estateback* variants highlights that there was a noticeably larger net improvement in the *base* variant’s drag, but the optimized *improved* variant is still predicted to have less drag than the optimized *base* variant. At the initial optimization iteration, the *improved* variant was predicted to have 4.8% less drag than the *base* variant, but the converged design vector was only predicted to have an improvement of 2.2%. This therefore shows the importance of using expert-driven design to reduce the net drag of a car. While the optimizer was capable of improving the *base* variant more than the *improved* variant, the optimized *improved* variant was still better than the optimized *base* variant. Since the MCF method that was used to drive the modifications of the *improved* variant, it is more favourable to use this method during the early phases of a design cycle to respect styling constraints, with optimization being reserved for later stages, when further improvements would be desired.

Table 5.3 highlights the performance of all of the optimizers for the two variants. The minimized objective function value relative to the initial objective function value for a variant (J^*/J_0), as well as the minimized objective function value relative to the *base* variant’s initial drag ($J^*/J_{0,Base}$) are presented to highlight the relative and net improvements found from each algorithm. J_K/J_0 is the normalized value of the aerodynamic drag for the final opti-

mization iteration, K . If J^*/J_0 and J_K/J_0 differ, this is not an indication of a converged solution not being found, but that the converged solution was not the minimal solution found during the optimization process. For the equivalent number of flow solutions, k , this value is in relation to the number of iterations the initial primal solution took. All subsequent primal evaluations cost 0.6 equivalent flow solutions, and all adjoint solutions cost 0.4 equivalent flow solutions.

Table 5.3: Summary of optimization progress for each algorithm and variant studied, for the spoiler FFD

Algorithm	Variant	K	k	J^*/J_0	$J^*/J_{0,Base}$	J_K/J_0	Converged
CG	<i>Base</i>	9	13.4	0.930	0.930	0.930	Yes
	<i>Improved</i>	9	10.4	0.953	0.908	0.953	Yes
BFGS	<i>Base</i>	6	16.2	0.988	0.988	1.038	Yes
	<i>Improved</i>	12	27.2	0.979	0.932	0.979	Yes
DBFGS	<i>Base</i>	1	2.4	0.993	0.993	0.993	No
	<i>Improved</i>	1	2.4	0.997	0.949	0.997	No

From Table 5.3, the most improved optimization iteration results for the *base* and *improved* variants were predicted to be from the Conjugate Gradient optimization. Furthermore, the Conjugate Gradient evaluations were able to perform reduce drag in fewer equivalent function evaluations than the BFGS algorithm did. This was because the BFGS optimization loops required many line search iterations before either the iteration limit, or an improved design was found. The Conjugate Gradient models tended to require only one function evaluation per line search execution, which further supports the Conjugate Gradient’s search direction being more robust, but also more efficient than the other tested methods.

To understand the modifications to the *base* and *improved* variants, the deformed geometry for the *base* variant at optimization Iteration 9 and the deformed geometry of the *improved* variant at optimization Iteration 9 are presented in Figure ??, for the Conjugate Gradient models. When examining the displacement, a positive displacement means the surface was deformed inwards, along the surface normal, and a negative displacement is outwards, against the surface normal. The BFGS and DBFGS models will not be presented as they did not produce a comparable reduction in drag.

For both the *base* and *improved* variant, the outboard end of the spoiler was flared downwards more than the inboard end, creating more downward curvature on the outboard end of the roofline, leading to more outboard flow acceleration. Additionally, having the trailing edge lower on the outboard end of the roof can help contract the overall size of the wake. This would require validation with higher fidelity methods, as RANS tends to underpredict the amount of separation on the rear fascia. Thus, it would likely perform relatively accurately up until separation occurred on the roofline.




Figure 5.10: Deformed surfaces, subject to the spoiler FFD

On the *base* variant, upstream of the spoiler, a bump was added to the roof, allowing for more flow acceleration upstream of the spoiler. This bump was not present on the *improved* variant, likely due to previous modifications that were guided by the MCF. Since the MCF yielded information that increasing flow acceleration at the trailing edge of the roof would help decrease drag, the roof was already given additional downwards curvature. As such, a bump was not needed to be generated to achieve this effect. It is also interesting to note that this bump was present on the *base* variant, since this was information that was directly distilled from the MCF guidance. However, the bump being developed was likely a result of the optimizer exploiting the geometry of the spoiler FFD, rather than being a physical solution. It can be seen that the bump present in the *base* variant was not tangent to the rest of the roofline. This suggested that the roofline should have been raised higher to accommodate for the additional curvature generated downstream of the bump. Having the bump present would create a local pressure buildup and potential thickening of the boundary layer, that would not be desired. Nonetheless, this highlights that the MCF and optimization should be used in tandem, whereby the MCF information guides early work in a car's design cycle to target key areas where drag can be minimized, with optimization extracting the final performance from the design.

5.7 Roof Spoiler Post-Optimality CFD

The designs on which post-optimality CFD was performed were: Conjugate Gradient optimization Iteration 9 and BFGS Iterations 4 and 6 for the *base* variant, as well as Conjugate Gradient Iteration 9 and BFGS Iteration 12 for the *improved* variant. The BFGS iterations were chosen to determine the accuracy of the RANS solver employed in the optimization process when large deformations were generated between optimization iterations, which is one of the hallmarks of the BFGS method. Additionally, BFGS Iterations 3 and 6 for the *base* variant were chosen to see if Iteration 3 actually resulted in reduced drag, and if Iteration 6 resulted in increased drag. The remaining options were chosen to verify if the predicted improvement was able to be replicated within a DDES model. The DBFGS iterations were not chosen for either variant, since only one optimization iteration was completed in these cases.

Figure 5.11 presents the percent difference between the post-optimality and optimizer's prediction of J/J_0 for the spoiler optimization. For the *base* variant, the change direction predicted by the optimizer was different than the direction calculated by the DDES model. In all three cases, if the optimizer predicted a change that increased drag, post-optimality CFD suggested it decreased drag, and vice versa. This trend was not present in the *improved* variant's optimization, with both analyzed designs predicting a reduction in drag with the post-optimality CFD also predicting a drag reduction. It is also valuable to see that in both cases of the *improved* variant's post-optimality analysis, they predicted roughly 1.4% more drag reduction in the post-optimality CFD, compared to the optimization value. It was not expected that the post-optimality CFD would provide the actual drag savings from an optimization design vector, but rather a more accurate sense of the improvement predicted. However, drag reduction predicted by post-optimality CFD was trusted, given the consistency across different designs shown when validating the DDES models. Experimental validation of the designs would be required to understand the true drag savings achieved.




Figure 5.11: Percent difference in J/J_0 from post-optimality CFD to optimizer predictions, for the spoiler design vectors chosen

It is possible that the optimizer was more in line with the post-optimality results, because the geometry had already been modified. The modifications to the *improved* variant were all made with the goal of reducing large separation points and making wakes more uniform. This would cause the SA turbulence model to perform more effectively in the optimizer, which explains why there was better agreement to the post-optimality DDES solution.

To understand the differences in predicted improvement between the post-optimality CFD and the optimizer, it is valuable to look at the surface pressures resulting from both methodologies. Figures 5.12 - 5.14 present comparisons of surface pressures for select post-optimality runs. All comparisons not presented here can be found in Appendix F.

Figure 5.12 presents the surface pressures for the *base* variant's Conjugate Gradient, Iteration 9 point. The key difference between the two flow fields is the significantly higher suction predicted by the post-optimality CFD around the bump that was created in the roof, and along the D-pillars. As noted, the bump generated by the optimizer for the Conjugate Gradient results was meant to exploit the negative MCF region that was previously discussed for the *base* variant. It is likely that the optimizer generated too much curvature in this region however, for which the flow acceleration could not be adequately refined, especially around the D-pillar. Based on the post-optimality results, the additional curvature along the D-pillar seems to have created a more noticeable separation region, which was not resolved in the optimization's RANS model. This had the effect of increasing suction on the upper half of the rear glass, which explains why a drag increase was predicted in the post-optimality CFD, but not in the optimization.




Figure 5.12: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 9 design point of the *base* variant

Figure 5.13 presents the surface pressures for the *base* variant's BFGS, Iteration 6 point. This point showed the opposite trend in post-optimality that the Conjugate Gradient point had. In this design, areas that the RANS model predicted to be high suction, resulted in less suction in the post-optimality CFD. The BFGS algorithm produced less D-pillar curvature than the Conjugate Gradient algorithm had, which means the post-optimality CFD predicting less suction in this region is reasonable. Furthermore, there was not as high a suction peak on the top of the roof, since the BFGS optimizer did not generate the distinct hump the Conjugate Gradient method did. This resulted in a higher base pressure, leading to a drag improvement being predicted, instead of the drag increase predicted by the optimizer.




Figure 5.13: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 6 design point of the *base* variant

Figure 5.14 presents the surface pressures for the *improved* variant's BFGS, Iteration 12 point. For this geometry, the post-optimality and optimization results were far more consistent than the *base* variant's results. In general, there was less deformation of the *improved* variant's surface, leading to smaller adverse pressure gradients on the roof's surface. Other than the different base pressure predicted by the DDES model, which was inherent due to its ability to predict separation more accurately, the two models showed strong agreement.




Figure 5.14: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 12 design point of the *improved* variant

In summary, the post-optimality CFD had better agreement with the optimizer for the *improved* variant simulations. This was due to the modifications that defined the *improved* variant, which sought to reduce the size of the rear wake and make it uniform. As a result, the optimizer’s underlying RANS solver was better able to deform the geometry in a manner that would yield a drag reduction. This further highlights the importance of expert-driven design modifications prior to the application of optimizers for automotive purposes.

5.8 Front Bumper Cheek Optimization

The results of the front bumper cheek optimization will be presented in the same manner as the spoiler optimization.

5.8.1 *Base* Variant Results

The evolution of aerodynamic drag normalized to the initial function evaluation as a function of optimization iteration is presented in Figure 5.15 for the front bumper cheek of the *base* variant. In contrast to the *base* variant’s spoiler optimization, none of the algorithms showed a clear path of improvement by deforming the front bumper cheek. There was a considerable amount of *zig-zagging* by the various algorithms, with a clear improvement being elusive. This was likely due to the more complex flow physics which occurred around the front bumper cheek, compared to the spoiler region. The flow that comes off the front bumper cheeks was largely dominated by the interactions it had with the front tires and their wake. Since the optimizer makes use of a SA RANS model, the ability to predict these wakes was susceptible to significant modelling errors due to the lack of resolved turbulent kinetic energy. As such, the sensitivities in this area may not only be incorrect, but the resultant primal evaluations may not have given a reasonable evaluation of the flow field.




Figure 5.15: Optimization results for each of the three optimization algorithms. The objective function is normalized to the drag of the *base* variant at the initial optimization iteration

In the case of the BFGS optimization, large deformations were generated for the front bumper cheek, which would likely not be acceptable from a styling perspective. However, the BFGS algorithm was capable of yielding more improvement than the Conjugate Gradient approach, as it converged early due to small surface displacements. This is in contrast to the spoiler optimization where the BFGS algorithm was not as effective as Conjugate Gradient.

5.8.2 Improved Variant Results

The evolution of aerodynamic drag normalized to the initial function evaluation as a function of optimization iteration is presented in Figure 5.16 for the *improved Estateback* variant.

It is noteworthy that the BFGS and DBFGS traces differed so significantly. This can be attributed to the inversion of Hessian matrix, occurring in the DBFGS algorithm causing variations in the search direction. These will not occur for the BFGS algorithm, as the *OpenFOAM* implementation directly solves for the inverse Hessian matrix. As with the *base* variant, the BFGS algorithm predicted greater improvement in drag than the Conjugate Gradient algorithm, largely driven by larger surface displacements. As with the *base* variant, little improvement was found by modifying the front bumper cheek region, due to the complex flow physics that are difficult to model using a RANS solver. This further highlights that optimization was likely best suited to areas in which the RANS solver could yield reasonable predictions: in areas that are not dominated by unsteady flow phenomena.




Figure 5.16: Optimization results for each of the three optimization algorithms

5.8.3 Comparing Optimization Results

This section details how the optimization of both variants compared to each other, normalized by the initial objective function value for the *base* variant, denoted by $J_{0,Base}$. The objective function evolution traces for each algorithm and variant are presented in Figure 5.17. As with the optimization of the spoiler, the *improved* variant displayed less relative drag reduction when deforming the front bumper cheeks. This should be expected, since the purpose of the *improved* variant was to push the *DrivAer Estateback* closer to a minimal drag design, meaning it should result in less reduction in drag from the optimization routine than the *base* variant. However, significantly less improvement was predicted by deforming the front bumper cheeks, since the maximum relative improvement was less than 0.5% compared to up to 7% relative improvement by deforming the spoiler. This is once again due to the more complex flow physics in this region, which can be difficult for a RANS-based optimizer to yield a reduction in drag.




Figure 5.17: Optimization results for each of the three optimization algorithms, for both *Estateback* variants. Both variants' drag is normalized by the drag of the *base* variant at the initial optimization iteration ($J_{0,Base}$)

Table 5.4 highlights the performance of all the optimizers for the two variants for the front bumper cheek optimization, similar to the spoiler optimization. Unlike the spoiler optimization, the BFGS algorithm resulted in the most improvement in drag for both the *base* and *improved* variants. However, it did not find a converged solution, meaning these designs can only be considered improvements, rather than optima. Part of the reason the Conjugate Gradient method did not yield a noticeable improvement, is because the first few iterations result in little total displacement of the surface with the maximum displacement being around three millimeters. Since this region was dominated by complex flow physics, these small displacements were difficult to resolve with a RANS solver. Even though the drag reduction was predicted from deformations of similar magnitude in Chapter 4, these designs were evaluated using DDES models which were better suited to resolving the front tire wakes. As such, the Conjugate Gradient evaluations converged early due to the convergence criteria used. Therefore, it could be more valuable to perform a BFGS optimization sweep first, on which the best iteration could then have a Conjugate Gradient optimization sweep performed.

Table 5.4: Summary of optimization progress for each algorithm and variant studied, for the front bumper cheek FFD box

Algorithm	Variant	K	k	J^*/J_0	$J^*/J_{0,Base}$	J_K/J_0	Converged
CG	<i>Base</i>	2	3.4	0.998	0.998	0.998	Yes
	<i>Improved</i>	2	5.8	1.000	0.952	1.000	Yes
BFGS	<i>Base</i>	5	11.8	0.996	0.996	0.998	No
	<i>Improved</i>	3	9.2	0.997	0.949	0.997	No
DBFGS	<i>Base</i>	2	5.8	0.998	0.998	1.005	No
	<i>Improved</i>	3	9.2	1.000	1.000	1.005	No

One aspect of the results that was consistent between the spoiler and front bumper cheek results was that the BFGS algorithm takes consistently more function evaluations when performing an optimization routine. As with the spoiler optimization, the BFGS algorithm resulted in a search direction where an improvement in the objective function required more than one line search iteration. This can be attributed to the additional curvature information that the BFGS algorithm provided the solver, compared to the Conjugate Gradient case. This allows for a larger traversal of the objective function topology, but it may then require additional line search iterations to meet the *Armijo* conditions.

From Table 5.4, the most improved optimization iteration results for the *base* and *improved* variants were predicted to be from the BFGS algorithm. The deformed geometry for the *base* variant at optimization Iteration 3 and the deformed geometry of the *improved* variant at optimization Iteration 3 are presented in Figure 5.18.


(a) *Base* Variant

Figure 5.18: Deformed surfaces, subject to the front bumper cheek FFD


(b) *Improved Variant*

Figure 5.18: Deformed surfaces, subject to the front bumper cheek FFD

The *base* improvement showed roughly three times as much deformation than the *improved* variant did in their most improved forms. The deformation realized for the *base* variant would not be practical however, from neither a manufacturing nor an aesthetics perspective. However, it is interesting to note that on the lower portion of the front bumper cheek, the optimizer sought to grow a tire air deflector, to push air outwards around the lower portion of the tire. The optimizer also tried to push the bumper cheek outwards, which would generate some local thrust due to the additional curvature, but would also help align the flow to be more parallel to the face of the front tires. This was present on the *improved* variant as well, but to a lesser degree.

The inwards deformation on the upper half of the front bumper cheek for the *base* variant is likely an effect of the RANS solver being unable to resolve the flow physics due to the front bumper cheek's interactions with the tires. While some inward curvature would help generate suction and local thrust, the bumper cheek was pushed in too much, causing some of the front tires to be exposed. There was some inwards deformation of the *improved* variant, however it was far more subtle, seeming to reduce the height of the flick. When this flick was initially generated, only the flick length along the front bumper cheek was changed, with all tested variants being 5 mm tall. Based on these deformations, it is evident that a RANS-driven gradient-based optimizer will perform more poorly in areas where the flow is not well behaved and features large areas of separation.

5.9 Front Bumper Cheek Post-Optimality CFD

The front bumper cheek designs on which post-optimality CFD was performed were: were Conjugate Gradient optimization Iteration 2 and BFGS Iterations 3 and 5 for the *base* variant, as well as Conjugate Gradient Iteration 2, BFGS Iteration 3, and DBFGS Iteration 3 for the *improved* variant. The DBFGS result for the *improved* variant was analyzed since the search direction for BFGS and DBFGS should have been similar, but there was a noticeable difference in their objective function values.

Figure 5.19 highlights how the chosen optimization iterations performed in the post-optimality DDES model for the front bumper cheek deformations. For both variants, there was no real trend in the post-optimality results, largely due to the complex flow physics around the front bumper cheek, which were less prevalent around the roof spoiler. Since the drag in this region was largely dominated by the front wheel drag, these are features that the RANS-based optimizer could not resolve with sufficient fidelity to yield effective change. However, as with the spoiler, the *improved* variant saw better drag correlation between the post-optimality CFD and the optimization results.




Figure 5.19: Percent difference in J/J_0 from post-optimality CFD to optimizer predictions, for the front bumper cheek design vectors chosen

Investigating the surface pressures for the *base* variant’s BFGS, Iteration 3, and the *improved* variant’s DBFGS, Iteration 3 helped identify where the disagreement occurred in the optimization and post-optimality results. Figures 5.20 and 5.21 present these comparisons, with additional post-processing images available in Appendix F.

For the *base* variant’s front bumper cheek optimization in Figure 5.20, the results show that there was agreement in the pressure predictions on the front faces of the wheelwell. However, the post-optimality CFD predicted more suction on the front bumper cheek, as a result of the additional curvature for the deformed geometry, that the RANS model does not. The largest discrepancy however was the pressure on the rear of the front tires and fender. For these areas, the post-optimality model predicted more suction, with the difference on the front tires being greater than $\Delta C_p = -0.3$. These two regions were key contributors to the drag of the vehicle, and were the source of one of the most dominant wake structures for the vehicle. Underpredicting the drag in these areas was a key contributor to why the negligible change in the optimization results yielded an increase in drag in the post-optimality CFD.




Figure 5.20: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 3 design point of the *base* variant

Figure 5.21 shows there was better agreement between the optimizer and post-optimality, since the *improved* model had smaller deformations. Without the large surface curvature that was present in the *base* variant's results, there were fewer regions where the RANS and DDES solutions were likely to differ. Furthermore, the deformation of the *improved* variant focused deformations on how the flow was jetted off the front bumper cheek flick, rather than growing a complex flick as the *base* variant had. As such, the pressure gradients in this area were less severe, and less air impinged on the front tires, leading to a smaller separation region. This yielded better pressure agreement between the DDES and RANS models, especially on the back of the front tires and fender.




Figure 5.21: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the DBFGS, Iteration 3 design point of the *improved* variant

In summary, since the front bumper cheek region was dominated by complex flow structures, the optimizer did not reduce drag measurably, and did not agree with the post-optimality CFD. This highlights the importance of the solver fidelity in the optimization routine. In areas where the flow physics are complex, it would be more effective to use a gradient-free optimizer, which could use DDES models for its drag evaluation, instead of a gradient-based RANS optimizer. This negates inaccuracies in a RANS model around adverse pressure gradients, and remove the need for post-optimality CFD.

6 Conclusions and Recommendations

The aim of this research project was to investigate the applicability and efficacy of gradient-based methods to reduce the aerodynamic drag of a canonical automotive geometry. An expert-driven Momentum Contribution Field (MCF) approach, followed by gradient-based optimization was used to achieve the drag savings. This found that the combined method was more effective than the using each method on their own. This chapter provides the conclusions reached along with suggestions and recommendations to improve the research performed and to inspire future research on this topic.

6.1 Conclusions

This research studies a new, unified design approach using continuous adjoint-based methods for automotive aerodynamics. Through adjoint-based sensitivity analyses, expert-guided and computer-guided approaches were investigated to determine their individual and combined efficacy when applied to the canonical *DrivAer Estateback* geometry. Both RANS and DDES simulations were applied to the *Estateback* model, with validation of the vehicle’s drag and surface pressures performed using the DDES model. Drag calculations for four variations of the *DrivAer Estateback* showed the DDES models to be accurate within 4% of wind tunnel values, with a consistent drag underprediction.

Using the RANS-based continuous adjoint methodology, the MCF with respect to drag could be calculated for the *DrivAer Estateback*. Modifications to the car guided by this field’s information predicted an 8.8% reduction in drag compared to the baseline model. This was achieved through modifications to the front splitter, front bumper cheeks, rear fenders, and roof spoiler. It was consistently found that using the MCF information resulted in more improvement than ignoring it.

Areas in which the RANS-based information yielded designs that consistently increased drag were also highlighted: around the hood’s trailing edge and the rear diffuser. Drag increases around the hood were noted due to inaccurate sensitivities from the RANS MCF. A methodology was described where these MCF predictions could be improved using time-averaged velocity, pressure, and face flux fields from a DDES simulation. These time-averaged fields could be used for the primal computation of the adjoint solver, providing higher accuracy in these primal quantities, than from a pure RANS solution.

When the RANS-based optimization methodology was used, the *base* and *improved* variants of the *Estateback* were optimized, using freeform deformation boxes to constrain surface displacement. For this optimization, the roof spoiler and front bumper cheek regions of both

variants were modified, using the Conjugate Gradient, BFGS, and DBFGS algorithms. For the spoiler region optimization, a relative drag reduction of up to 7% was predicted for the *base* variant, and up to 4.7% for the *improved* variant. The baseline variant predicted greater drag reduction, since the MCF-guided modifications of the *improved* variant placed it closer to a local minimum. However, the *improved* variant predicted a greater net improvement in drag when combining the MCF and optimization methods than just MCF or optimization modifications made to the *base* variant. For both variants of the *Estateback*, the Conjugate Gradient algorithm was found to be the fastest and most robust for reducing drag.

Optimization of the front bumper cheek highlighted the inaccuracies of the RANS solver and subsequent sensitivities in areas dominated by separation. Negligible relative improvement was predicted for either *Estateback* variant, due to the complex flow physics around and downstream of the front tires. For both regions, it was noted that the BFGS algorithm tended to yield more displacement earlier in the optimization routine, which tended to cause oscillatory changes in drag because of the inexact line search method used. Additionally, the BFGS and DBFGS algorithms tended to yield poorly conditioned sensitivities, due to the curvature information contained in their Hessian approximations, which caused the optimization routine to fail due to mesh quality violations. Poorer convergence of the primal solver in the front bumper cheek region directly impacts the adjoint computation, which informed the conditioning of the gradients used to update the search direction. This highlights that RANS-based gradient optimizers are likely to be insufficient for realistic automotive geometries in regions dominated by adverse pressure gradients and separation. This restricts the regions on the car where the optimizer will be effective.

Post-optimality CFD of key designs for the roof spoiler and front bumper cheek optimization routines highlighted areas in which the RANS-based optimizer suffered in accuracy. For the roof spoiler optimization, the *base* variant's predicted change did not show a consistent trend with the post-optimality CFD, influenced by adverse pressure gradients that were not effectively resolved by the RANS solver. The *improved* variant's post-optimality CFD agreed much better with the RANS model, with the same direction of change consistently predicted and the post-optimality solution tending to predict a consistent 1.4% additional drag reduction. Post-optimality CFD for the front bumper cheek optimization further emphasized the inability of effective drag minimization from a RANS-based solver. Greater disagreement was found in both the drag change direction and magnitude between the post-optimality and optimization CFD.

This research highlights the importance of using both expert-guided and computer-guided methodologies to reduce drag and to improve the efficacy of automotive aerodynamic design. Optimization results yielded greater improvement, and were more stable when applied to the *improved* expert-guided design variant of the *Estateback*. This shows that the two methods complement each other, with the expert-guided method being better suited to early phases of a car's design cycle. For example, it can be used to identify key areas where drag can be reduced as well as guide large initial changes to the design. Gradient-based optimization can then be applied to this developed variant of the car to yield additional drag reduction. It was concluded that caution is warranted when optimizing regions with high separation. The gradient-based optimizers had poorer performance in these regions due to sensitivity inaccuracies and poorer local primal solution convergence.

6.2 Recommendations for Future Work

In this work, a number of areas were identified that could yield additional validation of the results, as well as methods to improve the efficacy of the outlined design approach, along with additional areas in which this design framework can be applied.

Despite the DDES models used to evaluate design changes being validated against existing experimental data, the various stages of the Estateback's design modifications were not verified. Wind tunnel evaluation of the designs conceptualized in this work would be valuable to provide additional CFD and MCF validation.

A number of improvements to the tested optimizer were noted: the first would be the ability to place restrictions on the total displacement of the freeform deformation box in each direction. Large displacements generated by the BFGS algorithm caused the gradient approximation to suffer accuracy. Limiting the total displacement could still yield an optimization iteration where the objective function increases due to the inexact line search, but it could damp control point displacement such that the optimizer could guide the solution back towards a local minimum. Furthermore, being able to rescale the design variables at each iteration would be valuable. This could work with total displacement limitations to damp certain control points that lead to excessive deformations. Both of these improvements would help improve optimizer robustness.

Having more objective functions with their corresponding adjoint derivations could improve the robustness of the solution. It was noted that minimizing the square of drag could be advantageous, since the squared term helps modify the design space to become more convex. It is not guaranteed to remove multi-modality, but it could improve the numerical conditioning of the problem.

While the gradient-based optimization yielded design improvements for the spoiler region, it was limited by the underlying RANS solver. It would be worthwhile to explore gradient-free methods such as response surface methods, which can make use of DDES models for design evaluations. Exploiting methods that use higher-fidelity CFD simulations should also allow for more effective design exploration in regions dominated by separation, such as the front bumper cheek. The penalty of using gradient-free methods however, is they are limited to a fewer number of design variables to be computationally feasible. Therefore, a high-fidelity freeform deformation box would be too expensive to optimize an area of a car. Therefore, more accurate results could be achieved, but the deformation would have larger changes than a comparable gradient-based optimizer. Furthermore, making use of gradient-free methods using DDES solvers negates the need for post-optimality CFD.

Finally, an extension of this work should apply the MCF method to a high-downforce vehicle, with the MCF computed with respect to lift. This can be used to evaluate the method's efficacy over a broader range of objective functions, for vastly different design concepts. Particular interest should be paid as to how the MCF can be used to guide interactions between different downforce-generating devices, as well as to help guide the placement of additional devices. Furthermore, since downforce-generating cars do not need to produce a set amount of downforce, it would be worthwhile to investigate if there are the same limitations on the MCF's guidance in positive MCF regions.

Bibliography

- [1] V. Masson-Delmotte, P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, A. Pirani P.R. Shukla, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J. B. R. Matthews, Y. Chen, X. Zhou, M. I. Gomis, E. Lonnoy, T. Maycock, M. Tignor, and T. Waterfield, editors. *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*. IPCC, 2018.
- [2] International Energy Agency. Energy Technology Perspectives 2020. Technical report, International Energy Agency, 2020.
- [3] Setting CO₂ emission performance standards for new passenger cars and for new light commercial vehicles, and repealing regulations (EC) No 443/2009 and (EU) No 510/2011, 2019.
- [4] Nick Kurczewski and Brian Normile. Electric Cars With the Longest Range, May 2022.
- [5] Michael Hinterbuchinger. Aerodynamics electrified, Porsche Taycan Turbo. In *12th FKFS Conference Progress in Vehicle Aerodynamics and Thermal Management, Stuttgart Germany*, September 2019.
- [6] Alexander Nastov and David Caples. Aerodynamic Development of the 2019 Chevrolet Corvette C7 ZR1. In *SAE Technical Paper Series*. SAE International, April 2019.
- [7] Ziyu Guo, Yingchao Zhang, and Wei Ding. Optimization of the aerodynamic drag reduction of a passenger hatchback car. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 233(8):2819–2836, 2019.
- [8] F. Angrand. Optimum design for potential flows. *International Journal for Numerical Methods in Fluids*, 3(3):265–282, 1983.
- [9] W. Kyle Anderson and Daryl L. Bonhaus. Aerodynamic Design on Unstructured Grids for Turbulent Flows. Technical report, 1997.
- [10] W.Kyle Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *Computers & Fluids*, 28(4-5):443–480, may 1999.
- [11] E. M. Papoutsis-Kiachagias and K. C. Giannakoglou. Continuous Adjoint Methods for Turbulent Flows, Applied to Shape and Topology Optimization: Industrial Applications. *Archives of Computational Methods in Engineering*, 23(2):255–299, December 2014.
- [12] E.M. Papoutsis-Kiachagias, N. Magoulas, J. Mueller, C. Othmer, and K.C. Giannakoglou. Noise reduction in car aerodynamics using a surrogate objective function and the continuous adjoint method with wall functions. *Computers & Fluids*, 122:223–232, November 2015.

-
- [13] Shamira Bagnall Hare. Exploring Adjoint Optimisation to Improve the Aerodynamics of a Formula 1 Car. Masters thesis, University College Dublin, April 2023.
- [14] Carsten Othmer. Adjoint methods for car aerodynamics. *Journal of Mathematics in Industry*, 4(1):6, 2014.
- [15] Thomas Blacha, Misha Marie Gregersen, Moni Islam, and Henry Bensler. Application of the adjoint method for vehicle aerodynamic optimization. In *SAE Technical Paper Series*. SAE International, April 2016.
- [16] Gianluca Francesconi Ing, Luca Miretti, Laura Lorefice, Francesco Pitillo Ing, and Nicola Paola Ing. Application of Adjoint Methods on Drag Reduction of Current Production Cars. In *CO2 Reduction for Transportation Systems Conference*. SAE International, may 2018.
- [17] E. M. Papoutsis-Kiachagias, V. G. Asouti, K. C. Giannakoglou, K. Gkagkas, S. Shimokawa, and E. Itakura. Multi-point aerodynamic shape optimization of cars based on continuous adjoint. *Structural and Multidisciplinary Optimization*, 59(2):675–694, October 2018.
- [18] Angelina Heft, Thomas Indinger, and Nikolaus Adams. Introduction of a New Realistic Generic Car Model for Aerodynamic Investigations. In *SAE 2012 World Congress & Exhibition*. SAE International, 04 2012.
- [19] Ping He, Charles A. Mader, Joaquim R.R.A. Martins, and Kevin J. Maki. An aerodynamic design optimization framework using a discrete adjoint approach with OpenFOAM. *Computers & Fluids*, 168:285–303, may 2018.
- [20] Emmanuel Guilmineau. Numerical Simulations of Flow around a Realistic Generic Car Model. *SAE International Journal of Passenger Cars: Mechanical Systems*, 7(2):646 – 653, 2014.
- [21] S.R. Ahmed, G. Ramm, and G. Faltin. Some Salient Features Of The Time-Averaged Ground Vehicle Wake. In *SAE International Congress and Exposition 1984*. SAE International, February 1984.
- [22] Daniel Wood, Martin A. Passmore, and Anna-Kristina Perry. Experimental Data for the Validation of Numerical Methods - SAE Reference Notchback Model. *SAE International Journal of Passenger Cars - Mechanical Systems*, 7(1):145–154, April 2014.
- [23] Georgios Karpouzas, E. Papoutsis-Kiachagias, Thomas Schumacher, Eugene De Villiers, K.C. Giannakoglou, and Carsten Othmer. Adjoint optimization for vehicle external aerodynamics. *International Journal of Automotive Engineering*, 7:1–7, 01 2016.
- [24] Juliane Nies, Chenyi Zhang, Max Tanneberger, Timo Kuthada, Felix Wittmeier, and Jochen Wiedemann. Introduction of the AeroSUV-A New Generic SUV Model for Aerodynamic Research. In *WCX SAE World Congress Experience*. SAE International, April 2019.
- [25] Tim A. Albring, Max Sagebaum, and Nicolas R. Gauger. Efficient Aerodynamic Design using the Discrete Adjoint Method in SU2. In *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. American Institute of Aeronautics and Astronautics, June 2016.
- [26] Siva Nadarajah and Antony Jameson. A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization. In *38th Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, January 2000.

-
- [27] Simone Landi and Claudio Santarelli. Simcenter TV Broadcasts: Simcenter STAR-CCM+ - Faster Than Ever! FSAE Aerodynamics with Simcenter STAR-CCM+, July 2021.
- [28] F. R. Menter. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8):1598–1605, August 1994.
- [29] I.S. Kavvadias, E.M. Papoutsis-Kiachagias, G. Dimitrakopoulos, and K.C. Giannakoglou. The continuous adjoint approach to the $k-\omega$ turbulence model with applications in shape optimization. *Engineering Optimization*, 47(11):1523–1542, December 2014.
- [30] A.S. Zymaris, D.I. Papadimitriou, K.C. Giannakoglou, and C. Othmer. Continuous adjoint approach to the Spalart–Allmaras turbulence model for incompressible flows. *Computers & Fluids*, 38(8):1528–1538, September 2009.
- [31] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Computers in Physics*, 12(6):620, 1998.
- [32] OpenCFD Ltd. OpenFOAM: User Guide v2112 snappyHexMesh, December 2021. Online; Accessed August 22, 2023.
- [33] OpenCFD Ltd. OpenFOAM: User Guide - 4.4 Mesh Generation with the snappyHexMesh Utility, 2023. Online; Accessed August 22, 2023.
- [34] OpenCFD Ltd. OpenFOAM: User Guide v2112 Reynolds Averaged Simulation (RAS), December 2021. Online; Accessed August 23, 2023.
- [35] P. SPALART and S. ALLMARAS. A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, January 1992.
- [36] F. R. Menter, M. Kuntz, and R. Langtry. *Ten Years of Industrial Experience with the SST Turbulence Model*, pages 625–632. Begell House, Inc., 2003.
- [37] OpenCFD Ltd. OpenFOAM: User Guide v2112 Spalart-Allmaras, December 2021. Online; Accessed August 23, 2023.
- [38] OpenCFD Ltd. OpenFOAM: User Guide v2112 k-omega Shear Stress Transport (SST), December 2021. Online; Accessed August 23, 2023.
- [39] M. Strelets. Detached eddy simulation of massively separated flows. In *39th Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, January 2001.
- [40] OpenCFD Ltd. OpenFOAM: User Guide v2112 k-omega-SST Detached Eddy Simulation (DES), December 2021. Online; Accessed August 23, 2023.
- [41] P. R. Spalart, S. Deck, M. L. Shur, K. D. Squires, M. Kh. Strelets, and A. Travin. A new version of detached-eddy simulation, resistant to ambiguous grid densities. *Theoretical and Computational Fluid Dynamics*, 20(3):181–195, may 2006.
- [42] F. R. Menter and M. Kuntz. Adaptation of eddy-viscosity turbulence models to unsteady separated flow behind vehicles. In *The Aerodynamics of Heavy Vehicles: Trucks, Buses, and Trains*, pages 339–352. Springer Berlin Heidelberg, 2004.
- [43] Philippe R. Spalart. Hybrid RANS-LES methods. In *Advanced Approaches in Turbulence*, pages 133–159. Elsevier, 2021.

-
- [44] Angelina Heft, Thomas Indinger, and Nikolaus Adams. Investigation of Unsteady Flow Structures in the Wake of a Realistic Generic Car Model. In *29th AIAA Applied Aerodynamics Conference, Honolulu, Hawaii*, number 2011-3669. American Institute of Aeronautics and Astronautics, June 2011.
- [45] R.I Issa, A.D Gosman, and A.P Watkins. The computation of compressible and incompressible recirculating flows by a non-iterative implicit scheme. *Journal of Computational Physics*, 62(1):66–82, January 1986.
- [46] S.V Patankar and D.B Spalding. A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows. *International Journal of Heat and Mass Transfer*, 15(10):1787–1806, October 1972.
- [47] OpenCFD Ltd. OpenFOAM: User Guide v2206 pimpleFoam, 2022. Online; Accessed April 20, 2023.
- [48] Evangelos Papoutsis-Kiachagias, Konstantinos Gkaragkounis, Andreas-Stefanos Margetis, Themis Skamagkis, Varvara Asouti, and Kyriakos Giannakoglou. adjointOptimisationFoam: An OpenFOAM-Based Framework for Adjoint-Assisted Optimisation. In *14th International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control, Athens, Greece*, number 7960-18382. Institute of Structural Analysis and Antiseismic Research National Technical University of Athens, 2021.
- [49] K.C. Giannakoglou I.S. Kavvadias, E.M. Papoutsis-Kiachagias. On the proper treatment of grid sensitivities in continuous adjoint methods for shape optimization. *Journal of Computational Physics*, 301(9,), 2015.
- [50] Parallel CFD and Optimization Unit, National Technical University of Athens. *adjointOptimisationFoam, an OpenFOAM-based optimisation tool*. Parallel CFD and Optimization Unit, June 2020.
- [51] R. Fletcher. Function minimization by conjugate gradients. *The Computer Journal*, 7(2):149–154, February 1964.
- [52] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, 1999.
- [53] M. J. D. Powell. Algorithms for nonlinear constraints that use lagrangian functions. *Mathematical Programming*, 14(1):224–248, December 1978.
- [54] Ernest C Perry, Steven E Benzley, Mark Landon, and Richard Johnson. Shape optimization of fluid flow systems. In *Proceedings of ASME FEDSM'00. 2000 ASME Fluids Engineering Summer Conference*, 2000.

Appendices

A Supplemental Mesh Dependency Study Images

Additional plots showing the divisions for a sample background mesh and the resultant effect the number of background mesh divisions had on the surface and volume mesh are presented in this section.

A.1 Sample *blockMesh* Output

The output of the *blockMesh* utility for the very coarse mesh sizing used in the mesh dependency study is presented in Figure A.1. This highlights the number of divisions in the length, width, and height directions (x , y , and z).





Figure A.1: Background *blockMesh* used for the 84 x 28 x 14 *Very Coarse* mesh

A.2 *DrivAer Estateback* Surface Meshes

This section presents images showing the overall surface mesh around the *DrivAer Estateback* for each of the five meshes. When defining the refinement of the surface in *snappyHexMeshDict*, the car's surfaces were meshed between cell level 8 - 11. Table A.1 summarizes what the corresponding length in mm of each of the cell levels corresponds to for each of the surface meshes. Figure A.2 provides additional close-ups of the mesh around the mirror are shown to highlight the effect of the mesh refinement on the different details.

Table A.1: Summary of surface cell level lengths in mm for each of the mesh dependency meshes

Mesh	Level 8	Level 9	Level 10	Level 11
<i>Very Coarse</i>	7.14	3.57	1.79	0.89
<i>Coarse</i>	6.67	3.33	1.67	0.83
<i>Medium</i>	6.25	3.13	1.56	0.78
<i>Fine</i>	5.88	2.94	1.47	0.74
<i>Very Fine</i>	5.56	2.78	1.39	0.69



(a) Region of interest




(b) *Very Coarse*


Figure A.2: Surface mesh samples for the different background meshes



(c) *Coarse*



(d) *Medium*



(e) *Fine*

Figure A.2: Surface mesh samples for the different background meshes



(f) *Very Fine*

Figure A.2: Surface mesh samples for the different background meshes

A.3 *DrivAer Estateback* Volume Meshes

This section presents images investigating the effect of the *blockMesh* sizing on the volume mesh, being investigated on the centreline plane of the car, and around the tires. Since the layer sizes are defined as relative aspect ratios of the local surface sizing, as the *level 0* size decreases, the size of the boundary layer also changes, adding refinement to the geometry. The sample volume meshes are presented in Figures A.3 - A.7.

Figure A.3: Volume mesh images resulting from the 84 x 28 x 14 *Very Coarse* mesh




Figure A.4: Volume mesh images resulting from the $90 \times 30 \times 15$ *Coarse* mesh




Figure A.5: Volume mesh images resulting from the $96 \times 32 \times 16$ *Medium* mesh




Figure A.6: Volume mesh images resulting from the $102 \times 34 \times 17$ *Fine* mesh




Figure A.7: Volume mesh images resulting from the 108 x 34 x 18 *Very Fine* mesh

B Additional Post-Processing of *DrivAer Estateback* Variants

In this section, additional post-processing highlighting the differences in the flow field around the four tested *Estateback* configurations are discussed. This serves as supplemental investigation regarding the effects the floor and mirror configurations have on the global flow field. These differences are examined through the DDES simulations performed during model validation, with wake structures and the surface pressure fields being investigated.

B.1 Wake Structures

The wake structures of the different configurations explained in Table 3.1 were examined through slices along the car's length, contoured to total pressure coefficient. The slice was clipped to a maximum total pressure coefficient of 0.9, such that the downstream structures would be visible. The notable change in the wake structures came from the deletion of the side mirrors, which remove their low energy wake content from the slices. The omission of the mirrors resulted in a strengthening of the A-pillar vortex, which resulted in greater loss along the top of the car. These images are presented in Figures B.1 to B.4.




Figure B.1: Wake structures of the E_D_wM_wW configuration. Surface contoured by pressure coefficient




Figure B.2: Wake structures of the E_D_woM_wW configuration. Surface contoured by pressure coefficient




Figure B.3: Wake structures of the E_S_wM_wW configuration. Surface contoured by pressure coefficient




Figure B.4: Wake structures of the E_S_woM_wW configuration. Surface contoured by pressure coefficient

B.2 Surface Pressure Coefficient

The surface pressures of the different *Estateback* configurations are examined in this section. Figures B.5 to B.8 present the pressure distribution on the front and rear of each model, with Figures B.9 to B.12 presenting the underbody distributions. In general, the removal of the mirrors leads to lower pressure on the rear fascia. This stems from the higher energy wake at the back of the car, leading to a stronger shear layer at the separation point.




Figure B.5: Front and rear views of the surface pressure coefficient of the E_D_wM_wW configuration




Figure B.6: Front and rear views of the surface pressure coefficient of the E_D_woM_wW configuration




Figure B.7: Front and rear views of the surface pressure coefficient of the E_S_wM_wW configuration




Figure B.8: Underbody surface pressure coefficient of the E_S_woM_wW configuration




Figure B.9: Underbody surface pressure coefficient of the E_D_wM_wW configuration




Figure B.10: Underbody surface pressure coefficient of the E.D.woM.wW configuration




Figure B.11: Underbody surface pressure coefficient of the E.S.wM.wW configuration




Figure B.12: Underbody surface pressure coefficient of the E.S.woM.wW configuration

C Additional Post-Processing of *DrivAer Estateback* Modifications

This section provides supplemental images for the geometries not discussed in Section 4.6. The geometries already presented in Section 4.6 will not be presented again.

C.1 Additional Splitter Post-Processing

This section presents post-processing for Splitter Options 1 and 2, showing the underbody pressure coefficient for both designs in Figures C.1 and C.2.




Figure C.1: Comparison between the baseline splitter and Splitter Option 1 of pressure coefficient on the floor






Figure C.2: Comparison between the baseline splitter and Splitter Option 2 of pressure coefficient on the floor

C.2 Additional Front Bumper Cheek Post-Processing

This section presents post-processing for Front Bumper Cheek Options 1 and 2, showing the pressure drag coefficient around the tires for the two designs in Figures C.3 and C.4.




(a) Front Three Quarter View




(b) Rear Three Quarter View

Figure C.3: Comparison of pressure drag coefficient around the front tires for the baseline and Option 1 front bumper cheeks



(a) Front Three Quarter View



(b) Rear Three Quarter View

Figure C.4: Comparison of pressure drag coefficient around the front tires for the baseline and Option 2 front bumper cheeks

C.3 Additional Rear Fender Post-Processing

This section presents post-processing for the Rear Fender Infill Options 1, 3 and 4. Both pressure drag coefficient on the rear fascia, as well as a slice showing the total pressure coefficient in the wake will be presented for these designs in Figures C.5 - C.10.




Figure C.5: Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 1




Figure C.6: Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 3




Figure C.7: Pressure drag coefficient on the rear fascia: baseline fender versus Infill Option 4




Figure C.8: Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 1




Figure C.9: Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 3




Figure C.10: Total pressure coefficient on the $z = 310$ mm plane: baseline fender versus Infill Option 4

C.4 Additional Roof Spoiler Post-Processing

This section presents post-processing for the Roof Spoiler Options 1, 3 and 4. Both pressure drag coefficient on the rear fascia, as well as a slice showing the normalized velocity along car centreline will be presented for these designs in Figures C.11 - C.16.




Figure C.11: Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 1




Figure C.12: Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 3




Figure C.13: Comparison of pressure drag coefficient on the rear fascia between the baseline spoiler and Spoiler Option 4




Figure C.14: Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 1




Figure C.15: Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 3




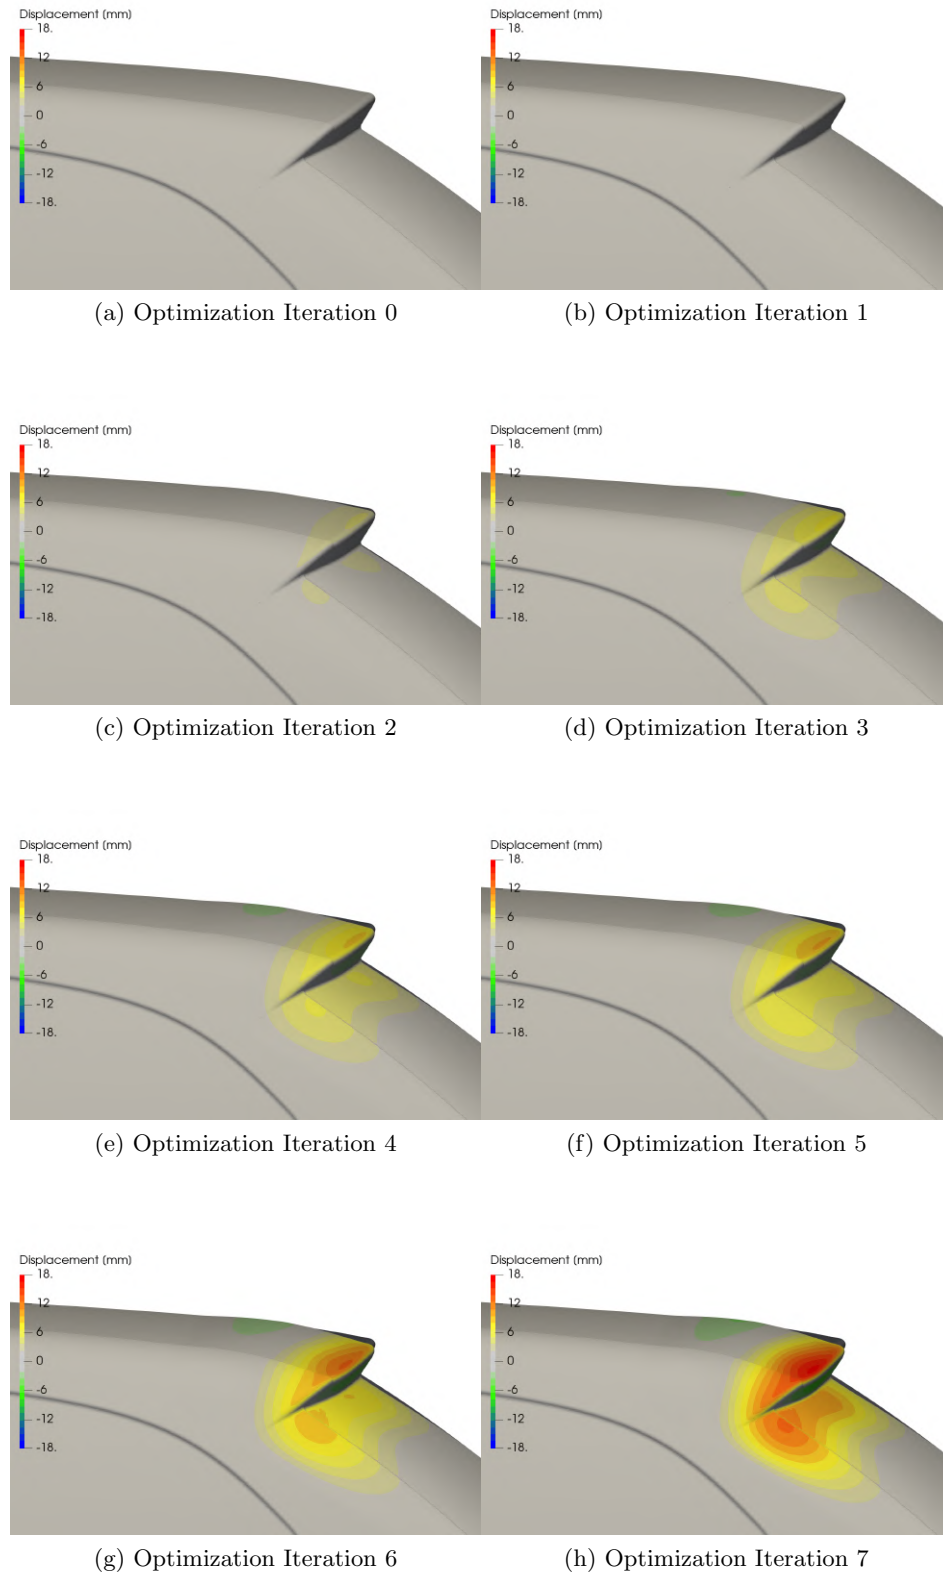
Figure C.16: Comparison of normalized velocity on the $y = 0$ mm plane between the baseline spoiler and Spoiler Option 4

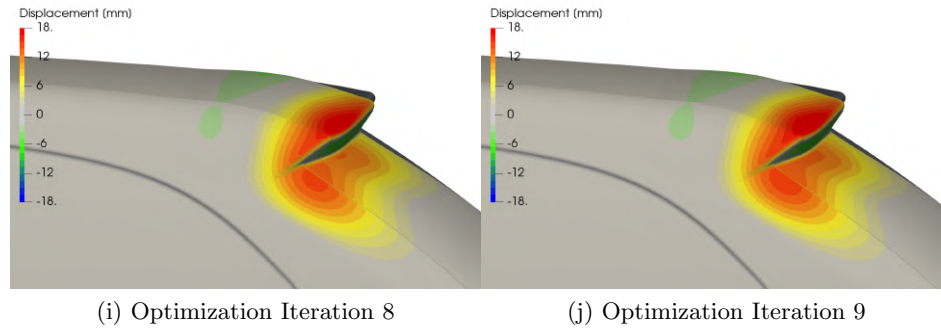
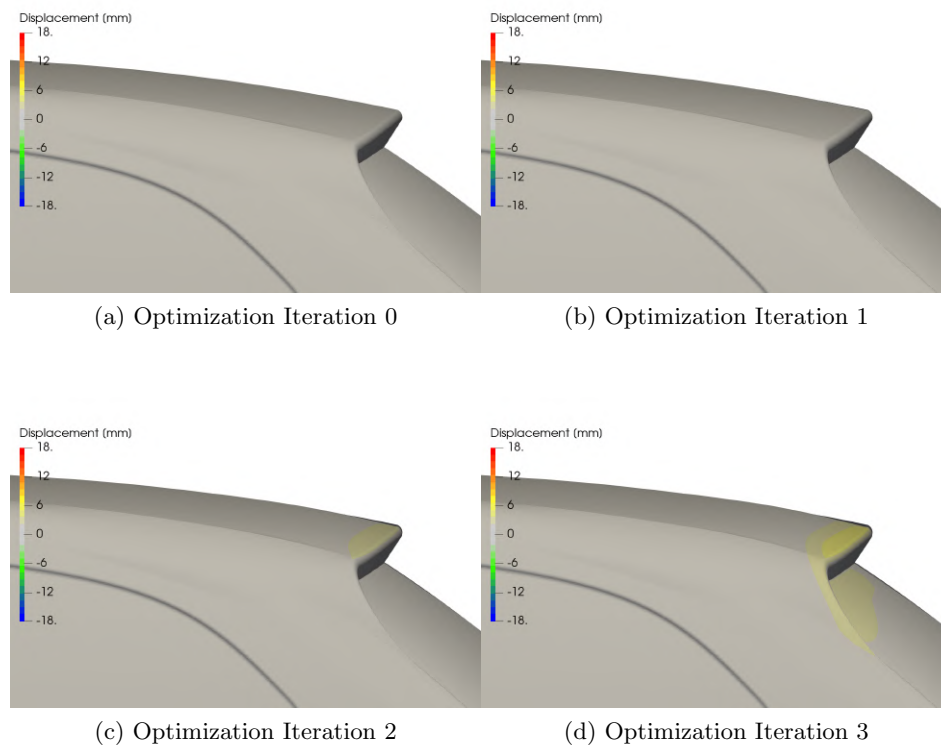
D Mesh Deformation Images for the Spoiler Optimization

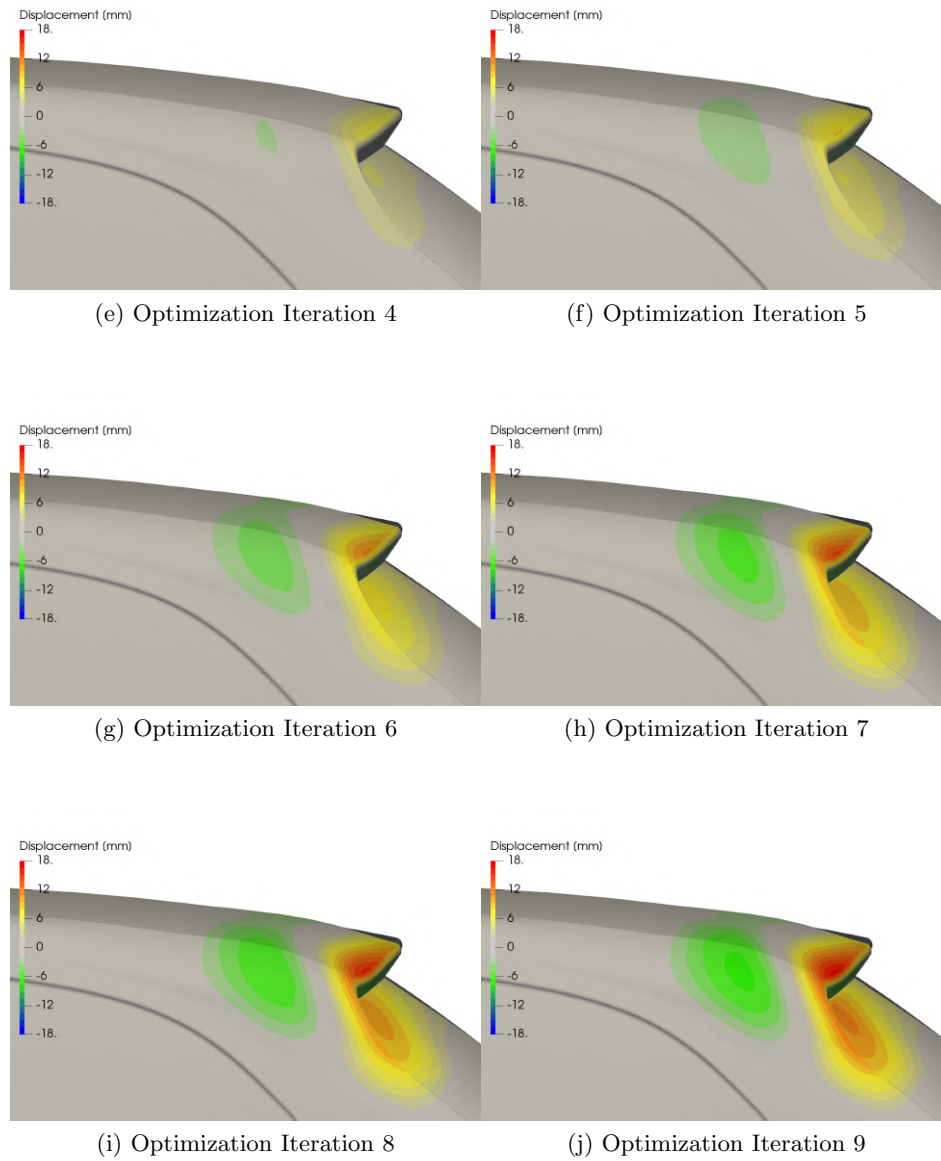
In this section, images showing the progression of the mesh deformation for both the base and improved variants roof spoiler will be presented. This serves as supplemental data regarding the direction each algorithm took when deforming the spoiler region of the *Estateback*. As mentioned, negative displacement indicates that the mesh was deformed outwards, against the surface normal, and a positive displacement is inwards, along the surface normal.

D.1 Conjugate Gradient Optimization

When applying the Conjugate Gradient method to the two variants, similar deformations were predicted. For both the *base* and *improved* variant, the outboard end of the spoiler was flared downwards more than the inboard end, creating more downward curvature on the outboard end of the roofline, leading to more outboard flow acceleration. On the *base* variant, upstream of the spoiler, a bump was added to the roof, allowing for more flow acceleration. As mentioned, this bump feature was likely not driven by actual flow physics, but the optimizer exploiting the geometry of the FFD. This feature was not as prominent on the *improved* variant, which showed that maintaining roof tangency upstream of the FFD was desired, as discussed in Section 5.6.3. Figures D.1 and D.2 show how the surface evolved during the optimization iterations.

Figure D.1: *Base* variant Conjugate Gradient spoiler deformation progression

Figure D.1: *Base* variant Conjugate Gradient spoiler deformation progressionFigure D.2: *Improved* variant Conjugate Gradient spoiler deformation progression

Figure D.2: *Improved* variant Conjugate Gradient spoiler deformation progression

D.2 BFGS Optimization

The BFGS algorithm lead to similar mesh deformation patterns as the Conjugate Gradient, namely the outboard end of the roof being lower than the inboard end. One key difference for the *base* variant, is it was not given downwards curvature along the roofline, which would result in less flow acceleration. This helps explain why the *base* BFGS optimization sweep tended to result in designs with a higher predicted drag. For the *improved* variant, since there was already downwards curvature along the roofline, this was exploited. It is valuable to note that the BFGS algorithm tended to result in larger displacements between each iteration compared to Conjugate Gradient. Figures D.3 and D.4 show the surface evolution for the two variants.

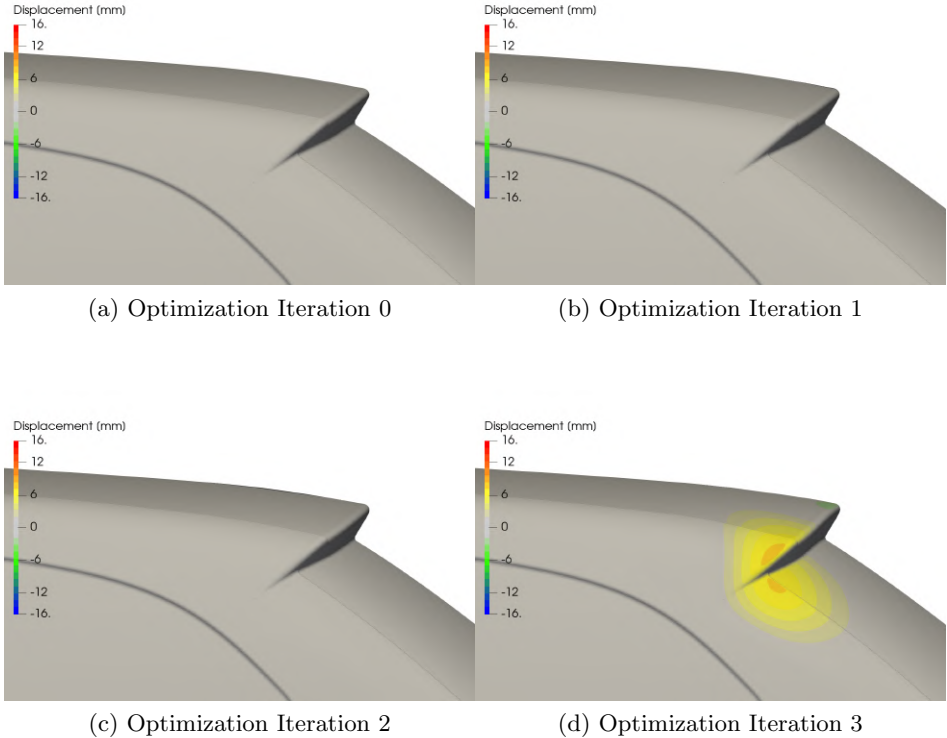


Figure D.3: *Base* variant BFGS spoiler deformation progression

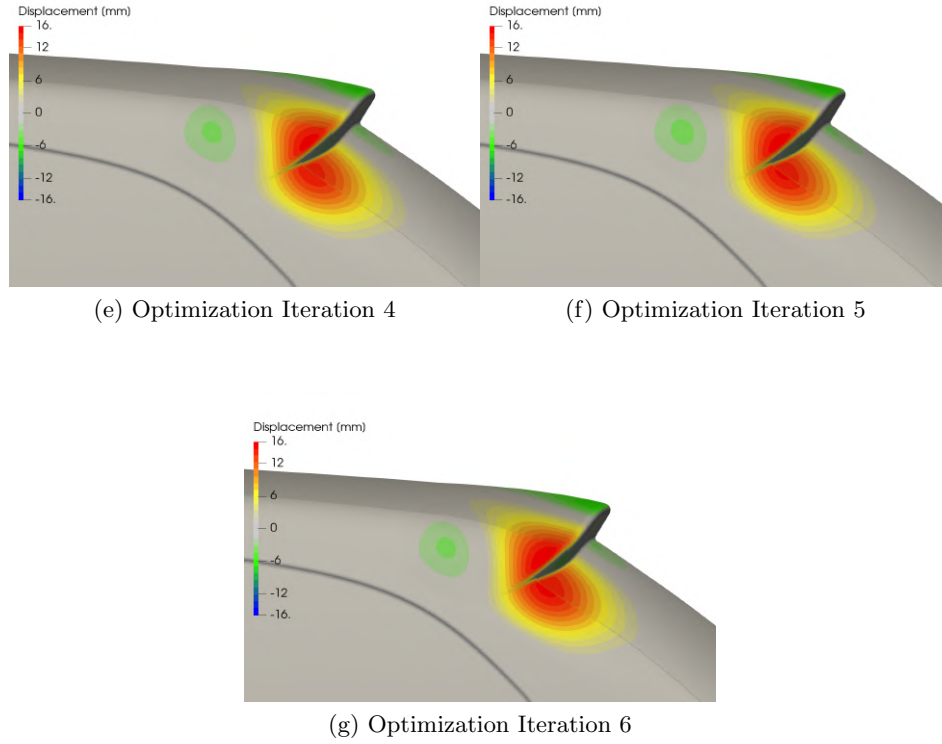


Figure D.3: *Base* variant BFGS spoiler deformation progression

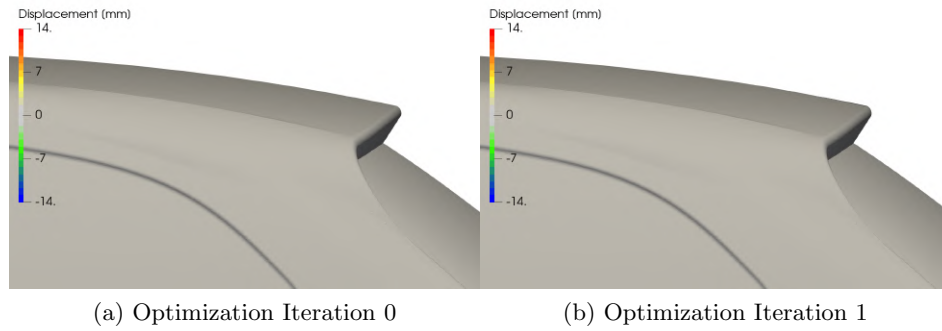




Figure D.4: *Improved* variant BFGS spoiler deformation progression




(c) Optimization Iteration 2

(d) Optimization Iteration 3




(e) Optimization Iteration 4

(f) Optimization Iteration 5



(g) Optimization Iteration 6

(h) Optimization Iteration 7



(i) Optimization Iteration 8

(j) Optimization Iteration 9

Figure D.4: *Improved* variant BFGS spoiler deformation progression

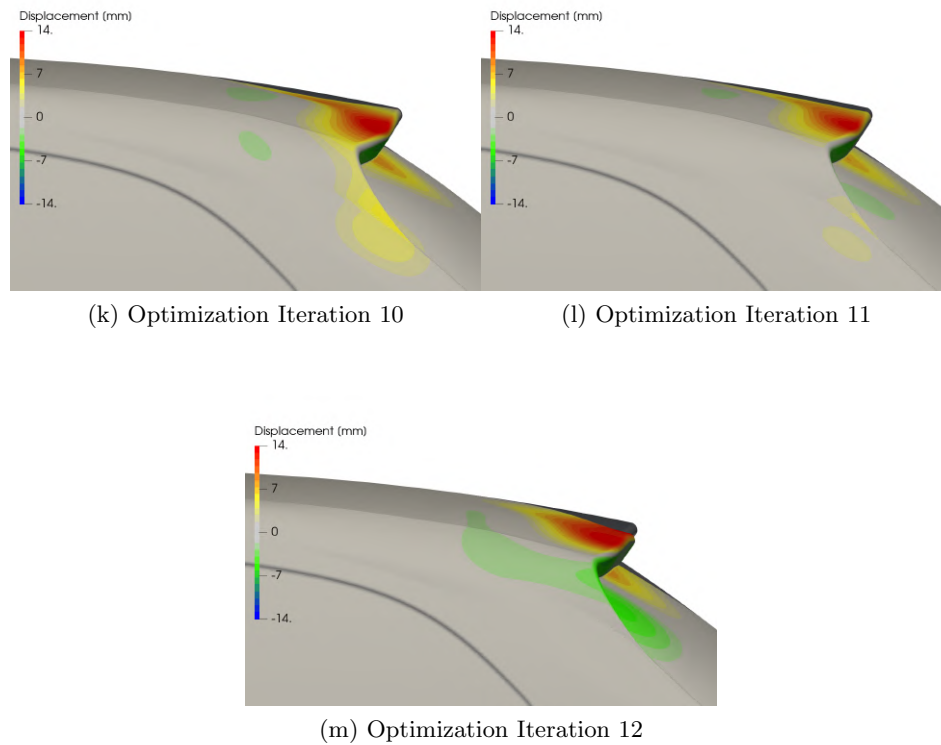


Figure D.4: *Improved* variant BFGS spoiler deformation progression

D.3 DBFGS Optimization

Since the DBFGS algorithm tended to yield mesh deformations that violated mesh quality criteria, it tended to not deform the surface much, being limited to 1 mm of maximum displacement. These images are therefore equivalent to optimization iterations zero and one for the Conjugate Gradient and BFGS methods. The ranges are however scaled down to show what displacement did occur. Furthermore, the *base* variant shows the surface deformation that led to the mesh quality issues, which entailed large, nonsensical surface deformations. Figures D.5 and D.6 present the mesh deformation for the DBFGS algorithm.

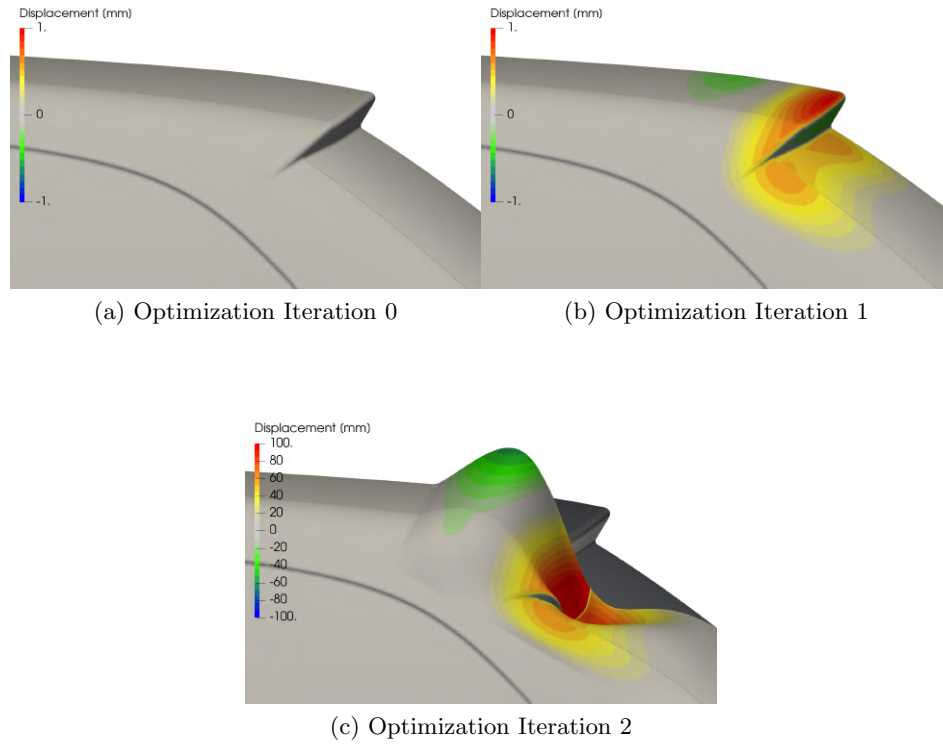


Figure D.5: *Base* variant DBFGS spoiler deformation progression

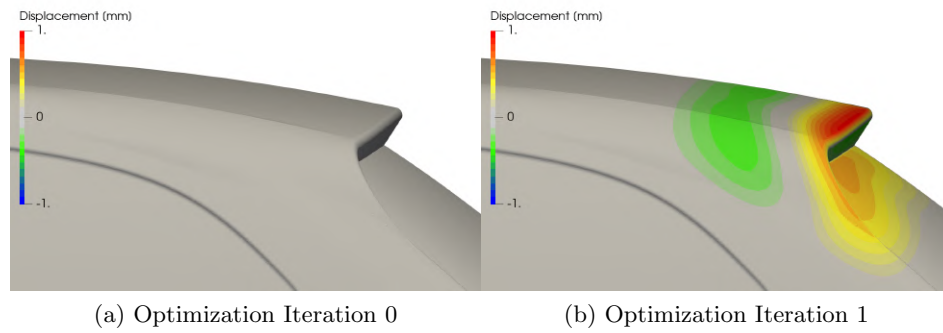


Figure D.6: *Improved* variant DBFGS spoiler deformation progression

E Mesh Deformation Images for the Front Bumper Cheek Optimization

In this section, images showing the progression of the mesh deformation for both the *base* and *improved* variants' front bumper cheek will be presented, serving as supplemental data for the optimization routines.

E.1 Conjugate Gradient Optimization

When applying the Conjugate Gradient method to the two variants, minimal total deformation was predicted. Small initial displacements yielded insufficient change in the objective function to prevent early convergence. Figures E.1 and E.2 show how the surface evolved during the optimization iterations.

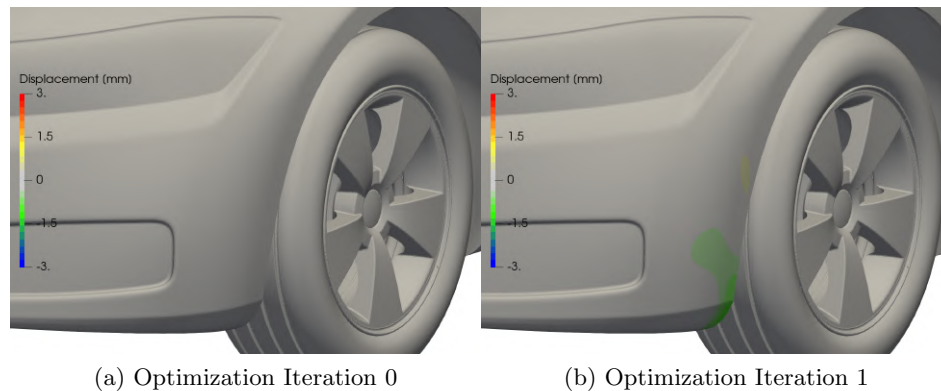




Figure E.1: Base variant Conjugate Gradient front bumper cheek deformation progression




(c) Optimization Iteration 2

Figure E.1: Base variant Conjugate Gradient front bumper cheek deformation progression



(a) Optimization Iteration 0

(b) Optimization Iteration 1



(c) Optimization Iteration 2

Figure E.2: Improved variant Conjugate Gradient front bumper cheek deformation progression

E.2 BFGS Optimization

The BFGS algorithm lead to very different deformations for the *base* and *improved* variants. The *base* variant saw significantly more deformation of the region, with more waviness in the result. The *improved* variant showed significantly less displacement. Furthermore, the *base* variant saw more inward deformation of the upper half of the bumper cheek than the *improved* variant. Both variants saw the cheek flared outwards however, as discussed in Section 5.8.3. Figures E.3 and E.4 show the surface evolution for the two variants.

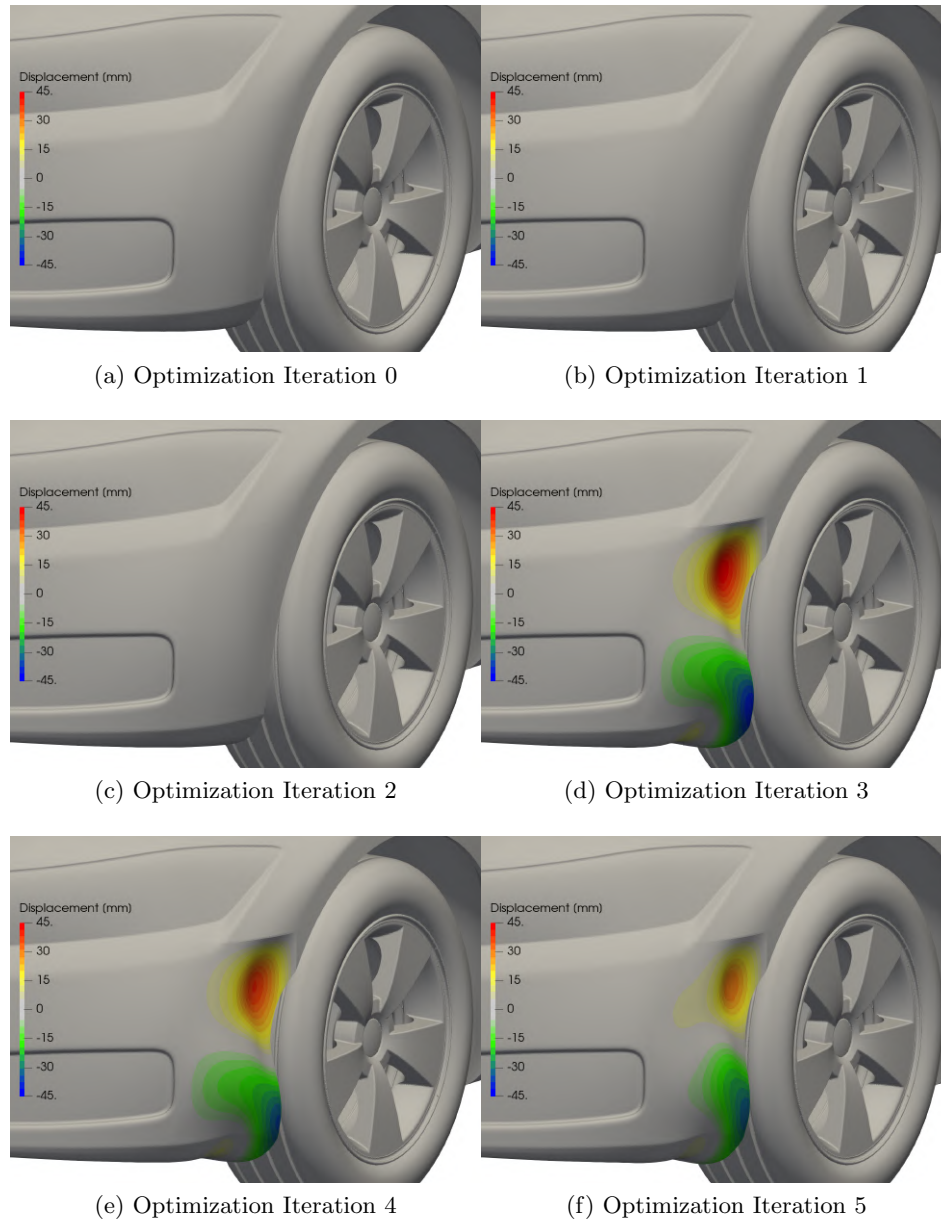

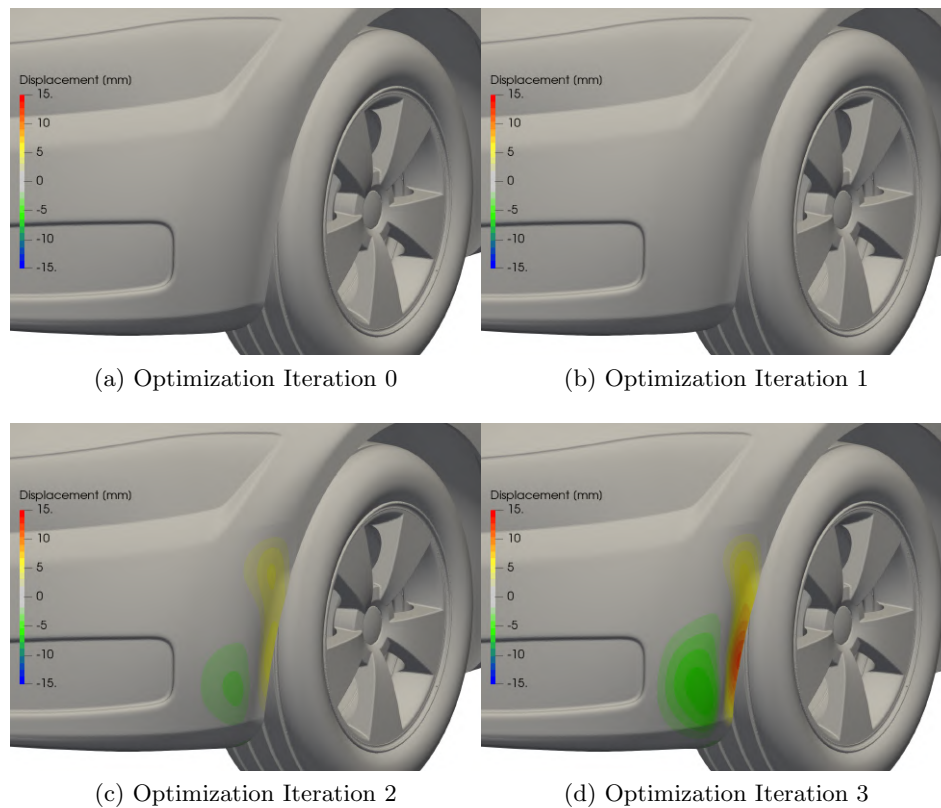


Figure E.3: *Base* variant BFGS front bumper cheek deformation progression



(g) Optimization Iteration 6

Figure E.3: *Base* variant BFGS front bumper cheek deformation progressionFigure E.4: *Improved* variant BFGS front bumper cheek deformation progression

E.3 DBFGS Optimization

Since the DBFGS algorithm tended to yield mesh deformations that violated mesh quality criteria, it tended to not deform the surface much for the *base* variant, as with the spoiler region. However, the *improved* variant showed larger total deformations and progressed through more optimization iterations before violating mesh criteria. It then resulted with a deformation field similar to the BFGS optimization. Figures E.5 and E.6 present the mesh deformation for the DBFGS algorithm.

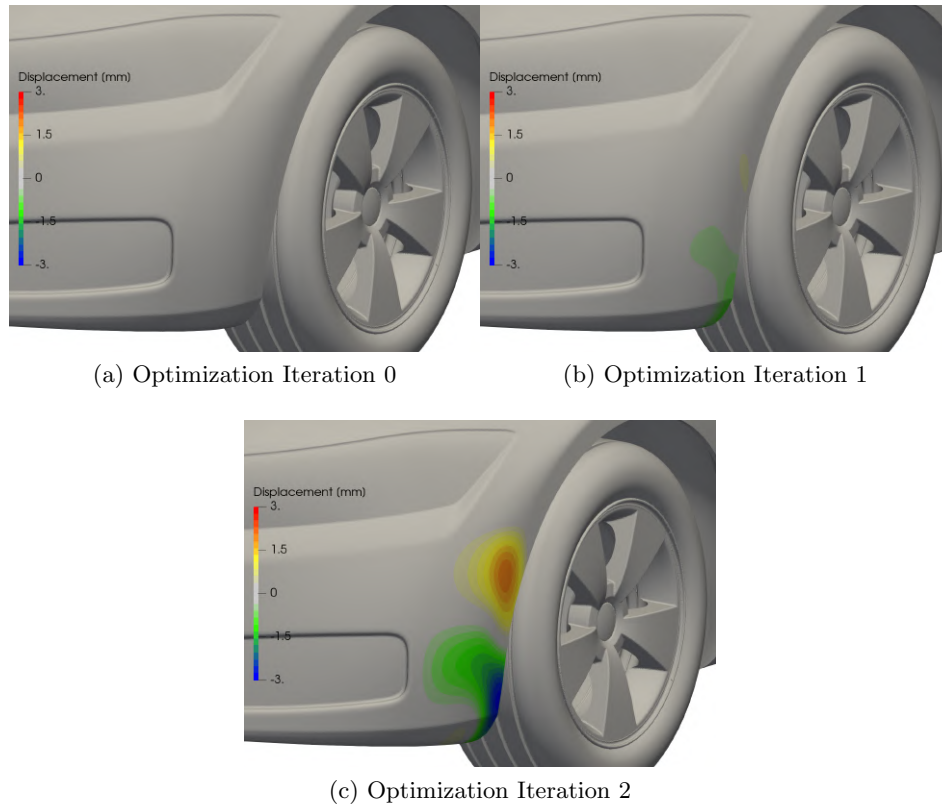


Figure E.5: *Base* variant DBFGS front bumper cheek deformation progression




Figure E.6: *Improved* variant DBFGS front bumper cheek deformation progression

F Additional Post-Processing of Post-Optimality CFD

In this section, further images comparing the surface pressures for the roof spoiler and front bumper cheek's post-optimality CFD campaign will be presented, serving as supplemental comparisons between the DDES and RANS models.

F.1 Additional Roof Spoiler Images

Figures F.1 and F.2 present the surface pressures for the *base* variant's BFGS Iteration 4 and the *improved* variant's Conjugate Gradient Iteration 9 designs.




Figure F.1: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 4 design point of the *base* variant




Figure F.2: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 9 design point of the *improved* variant

F.2 Additional Front Bumper Cheek Images

Figures F.3 - F.5 present the surface pressures for the *base* variant's Conjugate Gradient Iteration 3 and BFGS Iteration 5, as well as the *improved* variant's Conjugate Gradient 2 and BFGS Iteration 3 design points.




Figure F.3: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 3 design point of the *base* variant




Figure F.4: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the BFGS, Iteration 5 design point of the *base* variant




Figure F.5: Comparison of surface pressures for between the post-optimality CFD and optimization CFD for the Conjugate Gradient, Iteration 3 design point of the *improved* variant