

ABSTRACT

Title of Dissertation: MULTI-FIDELITY PARAMETRIC SENSITIVITY
FOR LARGE EDDY SIMULATION

Nikhil Oberoi
Doctor of Philosophy, 2023

Dissertation Directed by: Professor Johan Larsson
Department of Mechanical Engineering

Designing engineering systems involving fluid flow under uncertainty or for optimality often requires performing many computational fluid dynamics (CFD) calculations. For low-fidelity turbulence modeling simulations such as Reynolds-averaged Navier-Stokes (RANS), such a framework has been established and is in use. However, for high-fidelity turbulence resolving simulations such as large eddy simulations (LES), the relatively high computational cost of even a single calculation hinders the development of such a framework. The overarching goal of this work is to aid LES in becoming a usable engineering design tool.

In this thesis, a computationally affordable approach to estimate parametric sensitivities of engineering relevant quantities of interests in an LES is explored. The method is based on defining a RANS problem that is constrained to reproduce the LES mean flow field. The proposed method is described and assessed for a shock/boundary layer interaction problem, where the shock angle and wall temperature are considered variable or uncertain. In the current work, a proof-of-concept of the proposed method is demonstrated. The

method offers qualitative improvements to the sensitivity prediction of certain flow features as compared to standalone RANS simulations, while using a fraction of the LES cost. Different cost functions to infer auxiliary RANS variables are also examined and their influence on the sensitivity estimation is assessed. Overall, the results serve as an important proof-of-concept of the method and suggests the most promising path for future developments.

MULTI-FIDELITY PARAMETRIC SENSITIVITY ESTIMATION FOR LARGE
EDDY SIMULATION

by

Nikhil Oberoi

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2023

Advisory Committee:

Professor Johan Larsson, Chair/Advisor
Professor Christoph Brehm, Dean's representative
Professor Kenneth Kiger
Professor Arnaud Trouvé
Professor Jacob Wenegrat

© Copyright by
Nikhil Oberoi
2023

Ravi asks Chenab, how is Sutlej doing?

ਰਾਵੀ ਤੇ ਛੇਨਾਬ ਪੁਛਦਾ, ਕਿ ਹਾਲ ਸਤਲੁਜ ਦਾ ?

راوی تو چناب پُچھدا کی حال ستلُج دا ؟

रावी तों चनाब पुछदा, कि हाल सतलुज दा ?

Acknowledgements

I am indebted to everyone who helped this work come to fruition.

First and foremost, I would like to thank Prof. Johan Larsson for giving me the opportunity to explore scientific frontiers. His intellect, exceptional observational skills, erudite research vision, and empathetic understanding of human situations have shaped me as a researcher. The suggestions and feedback during the entire Ph.D. journey have helped me improve my analytical, writing, and coding skills. I owe my professional career trajectory to him. A shared curiosity for almost anything with him is one of the most treasured memories of my time as a Ph.D. student.

Next, I would like to thank the committee members (in alphabetical order), Prof. Brehm, Prof. Kiger, Prof. Trouvé, and Prof. Wenegrat for taking time out to go through this thesis and sharing your valuable insights and suggestions. This has undoubtedly enhanced the quality of the thesis.

The work presented here would not have been possible without the generous compute time provided by the University of Maryland supercomputing resources (<http://hpcc.umd.edu>) and the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program at the Argonne Leadership Computing Facility. This work was funded by the Air Force Office of Scientific Research (AFOSR; grant FA9550-19-1-0210) and by the NNSA Predictive Science Academic Alliance Program (PSAAP; grant DE-NA0003993). Their support is gratefully acknowledged.

The lab space has been a fun working environment. I am thankful to all my labmates, from the past, and the present - Vedant, Ali, Abishek, Walter, Siavash, John, Raihan, and Kamrul. They have through the interactions made the time spend at the lab worthwhile. I have learned from their experiences and feedback when presenting the work. Especially during the lockdown days, the camaraderie shared with Ali and Vedant helped keep the spirits high. The food rides in Ali's Toyota Camry were a special highlight. Special mention goes to Adam at Massachusetts Institute of Technology, our discussions on topics that were both relevant and unrelated to our PSAAP collaborations were quite enriching. The thesis would not have reached its completion without Vedant's help while editing the draft.

I would not have been here without the support of my parents and extended family. Their unwavering and unconditional encouragement and understanding have been instrumental in my coming so far. Friends are the family you choose. Right from high school to graduate school I have been lucky to have made amazing friends who have always encouraged me, listened to my rants, balanced my perspective, heard my bad jokes, were an audience to my cooking experiments, supported me emotionally and financially, and shared many happy memories with me. Thank you, Sachit, Karishma, Aman, Chirrag, Kriti (Manav Sthali School), Lalit, Vivek, Parteek, Ayush, Rahul, Govind (IIT (BHU) Varanasi), Bhavesh, Kamal, Tanmay, Abhishek (University of Illinois at Urbana-Champaign), Spandan, Keshav, Vedant, Akarsh, Anshuman, Debankur and Sanket (University of Maryland at College Park). Nothing would have been possible without you all. Special thanks to the housemates Spandan, Keshav, Akarsh, and Sanket for keeping a familial environment, especially during the pandemic lockdown. Our random discussions were always fun to partake in.

I would like to express my appreciation to the Behavioral Health Services at the University

Health Center, and my therapist Justin Sybenga for their help in navigating life while I was here. Additionally, I extend my thanks to the graduate staff at the department for making the paperwork seamless.

This work would have been incomplete without the extensive anonymous help available online in form of course notes, YouTube videos, tutorials, providing answers to problems on forums, and open-source libraries. Special shoutout to the podcast and YouTube video creators, your content kept me entertained and taught me new things.

Lastly, I struggle to put it into words, but I am extremely grateful for my origins. It has made me who I am. The Indian subcontinent, through its history, food and cuisines, languages, and diverse cultures provided me with rabbit holes to explore when I wanted to digress from my thesis work. It always has and will hold a special place in my heart.

Table of Contents

Acknowledgements	iii
Table of Contents	vi
List of Tables	viii
List of Figures	ix
List of Abbreviations	xiv
Chapter 1: Introduction	1
1.1 Overview	1
1.2 UQ methods in CFD	5
1.3 Parametric sensitivity estimation	7
1.4 Objective of the thesis	9
1.5 Publications	10
Chapter 2: Modeling and numerical background	12
2.1 High-fidelity model: LES	13
2.2 Low-fidelity model: RANS	15
2.2.1 RANS closure models and Boussinesq approximation	17
2.3 Proposed multi-fidelity sensitivity algorithm	19
2.4 Choice of RANS model	23
2.4.1 Inference of auxiliary variables	24
2.5 CFD solvers	26
2.5.1 hybrid	26
2.5.2 tortuga	27
2.5.3 Finite-volume numerics	27
2.6 Solver verification	31
Chapter 3: Application to a shock/boundary layer interaction	41
3.1 Problem set-up	43
3.1.1 LES setup	45
3.1.2 RANS setup: standalone, frozen and LES-constrained	46
3.2 Results: nominal condition	48
3.2.1 Influence of the forcing term	58
3.2.2 Distribution of the forcing term	60

3.2.3	Grid assessment for MFSA	63
3.3	Sensitivity estimation	65
3.3.1	LES perturbation setup	65
3.3.2	RANS and MFSA perturbation setup	66
3.3.3	Results: wall quantities	67
3.3.4	Results: wall normal profiles	71
3.3.5	Results: output pdfs	76
3.4	Concluding remarks	76
Chapter 4:	Inference of turbulent diffusivities from LES	79
4.1	Inference cost functions	79
4.1.1	Rescaling of terms	81
4.1.2	Inference equations	81
4.2	Augmenting the inference process	83
4.2.1	Regularization term	84
4.2.2	Smoothness constraint	85
4.3	Effective inference equations	86
4.4	Numerical solution of the inference equations	88
4.4.1	Post processing	90
4.5	Comparison of inferred fields from methods M1 and M2	91
4.6	Choice of hyperparameters and grid on the inferred profiles	95
4.7	Simulations using M1- and M2-inferred data	99
4.7.1	Results: frozen RANS simulations at base state	99
4.7.2	Results: MFSA sensitivity predictions	105
4.8	Concluding remarks	112
Chapter 5:	Summary	114
5.1	Future directions	116
Appendix A:	Calculation of the Euler-Lagrange equations	118
A.1	M1 inference equations	119
A.2	M2 inference equations	121
A.3	Regularization term	123
A.4	Smoothness constraint	124
Bibliography		126

List of Tables

2.1	Sequence of the 5 grids used for TMR flat plate validation case. The Δy^+ at the wall, at the station of interest $x/L = 0.97$ are computed a posteriori, and indicates the quality of grid for wall-bounded flows.	32
2.2	Convergence history of c_f and δ at $x/L = 0.97$ for the sequence of 5 grids used.	40
4.1	Sequence of the 3 grids used for computation of M2 viscosity, describing the near-wall refinement as compared to the LES grid	99

List of Figures

1.1	Schematic illustrating the accuracy and cost trade-off between RANS, LES and DNS, particularly for wall-bounded flows.	2
1.2	Illustration of an UQ exercise using CFD. The first row exemplifies the desired outputs from a CFD toolbox. The second row shows the functional dependence of a QoI J on the input parameters a_1, a_2 and the uncertainty in input parameters via their joint pdf $f(a_1, a_2)$. The dashed black lines indicate the nominal values of the parameters. The parametric gradient of J at that condition is another useful output. The last row shows the uncertainty in J because of uncertain inputs.	3
2.1	This illustrates the objective of the work. The solid curves show the parametric profile of a QoI J w.r.t parameter a . The dashed lines tangent to these curves mark their slope at some nominal condition a_b . We expect to replicate the black dashed line using the blue dashed line from the proposed MFSA. . . .	20
2.2	A canonical 2-D finite-volume grid illustrating the cell centers where the data is stored, the cell faces and the cell sizes.	28
2.3	Schematic diagram of the NASA TMR flat plate test case illustrating the boundary conditions and station of interest at which the code tortuga is compared.	31
2.4	Δx and Δy distributions on the finest grid G5. Each coarser level of grid doubles the grid spacing in x and y at each point.	33
2.5	The first row shows the contours of streamwise velocity u along with the boundary layer height δ (solid black line). The second row shows the contours of μ_t , the turbulent dynamic viscosity, obtained from the SA model. The values are taken from the finest grid G5.	34
2.6	Grid convergence history of wall-normal profiles of u at the station of interest $x/L = 0.97$	36
2.7	Grid convergence history of wall-normal profiles of μ_t at the station of interest $x/L = 0.97$ in outer and inner scaling.	37
2.8	Grid convergence history of c_f v/s streamwise x/L	37
2.9	The rate of convergence for c_f and δ in comparison to the theoretical prediction of second order accuracy. The mesh size is $h = 1/\sqrt{n_x n_y}$ and the error in each grid is computed w.r.t the results on the finest grid.	39

3.1	Schematic illustration of an inlet of a scramjet engine. The isolator region experiences SBLI which is the application case for MFSA (inspired by Khan & Hasan, 2021)	42
3.2	Schematic representation of an SBLI flow configuration. The boundary layer is coming from left to right, experiences an oblique shock. This causes a separation bubble to form, where there is recirculating flow. The flow eventually reattaches. The wall again turns the flow causing a reflected shock.	42
3.3	Contours of ρ obtained from an LES at nominal condition. Yellow to orange shows an increase in ρ . The black lines indicate streamlines. The streamline starting from the wall and ending near the impingement location is the separation bubble. The top most streamlines indicate the turning of the freestream flow due to shocks. The streamline below it starts in the incoming boundary layer and indicates that the boundary layer height amplifies post-shock interaction.	44
3.4	Schematic plot illustrating the domain used and choice of boundary conditions for LES and RANS simulations. The LES domain is three-dimensional while RANS domain (highlighted in gray) is a subset of LES domain and excludes the spanwise direction.	45
3.5	c_f and c_h curves highlighting the discrepancy among different fidelities at baseline condition. The standalone RANS (red line) shows systematic errors when compared to LES data (black line). Frozen RANS with LES inferred closures by itself also shows discrepancies when compared to LES results. Frozen RANS is examined for two cases - fixed $Pr_t = 0.9$ (green curve) and variable LES-inferred Pr_t (orange curve).	50
3.6	Temperature contours for different fidelities contrasting the different fidelities - LES, standalone RANS and frozen RANS with LES inferred closures.	52
3.7	The streamlines in the freestream ($y \gtrsim 2.75\delta_r$) and outer part of boundary layer ($y \lesssim 2\delta_r$) are contrasted between different fidelities - LES (black curve), standalone RANS (red curve) and frozen RANS (fixed $Pr_t = 0.9$ (green curve) and variable LES inferred Pr_t (dashed orange curve))	53
3.8	The streamlines near the separation region and the bubble profiles are compared for different fidelities - LES (black curve), standalone RANS (red curve) and frozen RANS with LES inferred closure (fixed $Pr_t = 0.9$ (green curve) and variable LES inferred Pr_t (orange curve)). The brown lines without any arrow indicates the bubble profile.	55
3.9	The wall-normal profiles of the streamwise velocity u_1 highlights the discrepancies for different fidelities at different stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer (second row), near separation (third row), post-shock reattachment (last row). The black line shows LES data, red is standalone RANS, green is frozen RANS with fixed Pr_t and dashed-orange lines are frozen RANS with variable Pr_t	56

3.10	The wall-normal profiles of temperature T highlights the discrepancies for different fidelities at different stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer (second row), near separation (third row), post-shock reattachment (last row). The black line shows LES data, red is standalone RANS, green is frozen RANS with fixed Pr_t and dashed-orange lines are frozen RANS with variable Pr_t	57
3.11	c_f and c_h curves show the extent to which the LES-constrained RANS (blue line) replicates LES data (black dots) at the base condition, whereas the standalone RANS (red line) shows systematic errors when compared to LES.	59
3.12	The bubble profile, marked by the local bubble height (h_b) shows the extent to which LES-constrained RANS (blue line) replicates LES data (black dots) at the base condition, whereas the standalone RANS (red line) highlighting the systematic errors between LES and RANS.	60
3.13	Distribution of the forcing function used in the proposed multi-fidelity sensitivity algorithm framework. The forcing function is the residual of the LES data on the standalone RANS equation. The absolute value for each component is contrasted in log scale.	62
3.14	Comparison of standalone RANS (top) and LES-constrained RANS (bottom) with LES data overlaid as black dots on two grids - the LES grid and the LES grid coarsened by a factor of two everywhere, the “half” grid.	64
3.15	Comparison of base LES sensitivities of c_f w.r.t. T_w and β against standalone RANS and MFSA. The dashed vertical lines mark separation and reattachment locations. The gray region marks the 95% confidence interval of the LES predictions.	68
3.16	Comparison of base LES sensitivities of c_h w.r.t. T_w and β against standalone RANS and MFSA. The dashed vertical lines mark separation and reattachment locations. The gray region marks the 95% confidence interval of the LES predictions.	69
3.17	Comparison of sensitivities of u_1 profiles normal to wall in the incoming boundary layer, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.	72
3.18	Comparison of sensitivities of u_1 profiles normal to wall in the post-shock reattachment region, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.	73
3.19	Comparison of sensitivities of T profiles normal to the wall in the incoming boundary layer, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.	74
3.20	Comparison of sensitivities of T profiles normal to wall in the post-shock reattachment region, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.	75
3.21	The pdf of scalar QoIs are compared between LES, standalone RANS and MFSA.	77

4.1	Contours of μ_t near the separation bubble are compared for M1 and M2 inference methods. The black lines mark a streamline near the edge of the incoming boundary layer (top) and separation bubble (bottom), respectively.	91
4.2	The first row shows the $\bar{\rho}$ contours from base case LES. Each succeeding row illustrates the μ_t and Pr_t obtained from the M1 and M2 inference methods at different locations marked by vertical dashed lines in the $\bar{\rho}$ contours. These stations represent the attached boundary layer region, near-separation region and the post-shock reattachment region.	92
4.3	Influence of the hyperparameter ϵ on the inferred μ_t profiles at a station in the incoming boundary layer (top row) and the post-shock reattachment (bottom row) regions.	95
4.4	Influence of the hyperparameter ω_s on the inferred μ_t profiles at a station in the incoming boundary layer (top row) and the post-shock reattachment (bottom row) regions.	96
4.5	Influence of the wall normal grid distribution on the inferred μ_t fields at a station in the incoming boundary layer (top row), and the post-shock reattached boundary layer.	98
4.6	The streamlines in the freestream ($y \gtrsim 2.75\delta_r$) and the outer part of boundary layer ($y \lesssim 2\delta_r$) are compared for inference choices with LES (black curves). The M1-inferred frozen RANS streamlines are the red curves and M2-inferred frozen RANS streamlines are the blue curves. Here, both frozen RANS cases are with fixed $\text{Pr}_t = 0.9$. The brown line without any arrow indicates the bubble profile.	101
4.7	The streamlines near the separation region and the bubble profiles are compared for inference choices with LES (black curves). The M1-inferred frozen RANS streamlines are the red curves and M2-inferred frozen RANS streamlines are the blue curves. The brown line in each plot is the corresponding separation bubble. Here, both frozen RANS cases are with fixed $\text{Pr}_t = 0.9$.	102
4.8	Comparison of c_f and c_h profiles for frozen RANS with inferred fields from M1 and M2 methods with the baseline LES data.	103
4.9	The wall-normal profiles of the streamwise velocity u_1 comparing the differences for each inference methods at separate stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer, near separation (second row: left to right), inside the bubble and post-shock reattachment (third row: left to right). The black line shows LES data. The red and magenta lines are M1 frozen RANS, with fixed Pr_t and variable Pr_t respectively. The blue and green lines are M2 frozen RANS, with fixed Pr_t and variable Pr_t respectively.	104

4.10	The wall-normal profiles of the temperature T comparing the differences for each inference methods at separate stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer, near separation (second row: left to right), inside the bubble and post-shock reattachment (third row: left to right). The black line shows LES data. The red and magenta lines are M1 frozen RANS, with fixed Pr_t and variable Pr_t respectively. The blue and green lines are M2 frozen RANS, with fixed Pr_t and variable Pr_t respectively.	106
4.11	Comparison of base LES sensitivities of c_f w.r.t. T_w and β against MFSA prediction using M1- and M2-inferred variables. The dashed vertical lines mark separation and re-attachment locations. The gray region marks the 95% confidence interval of the LES predictions.	107
4.12	Comparison of base LES sensitivities of c_h w.r.t. T_w and β against MFSA prediction using M1- and M2-inferred variables. The dashed vertical lines mark separation and re-attachment locations. The gray region marks the 95% confidence interval of the LES predictions.	108
4.13	Comparison of sensitivities of u_1 wall-normal profiles w.r.t T_w and β , just before separation, between M1- and M2-inferred fields. The gray region marks the 95% confidence interval of the LES predictions.	110
4.14	Comparison of sensitivities of T wall-normal profiles w.r.t T_w and β , just before separation, between M1- and M2-inferred fields. The gray region marks the 95% confidence interval of the LES predictions.	111

List of Abbreviations

CFD	Computational Fluid Dynamics
LES	Large Eddy Simulation
RANS	Reynolds-Averaged Navier-Stokes
UQ	Uncertainty Quantification
QoI	Quantity of Interest
pdf	probability distribution function
WENO	Weighted Essentially Non-Oscillatory
MFSA	Multi-Fidelity Sensitivity Algorithm
LE	Lyapunov Exponent
LSS	Least Square Shadowing
NILSS	Non-Intrusive Least Square Shadowing
MC	Monte Carlo
PCE	Polynomial Chaos Expansion

Chapter 1: Introduction

1.1 Overview

Turbulence is the state of flow in which the fluid particles exhibit unsteady and chaotic motion, which is characterized by an efficient mixing of particles and properties like density, concentration, temperature, etc. It is a multi-scale phenomenon where the dynamical interactions happen at a continuous range of length and time scales, and are governed by the Navier Stokes equation. Computational fluid dynamics (CFD) has helped gain a better understanding about this state of flow. It has also been used by engineers to design equipment used under such conditions. From a numerical standpoint, turbulence can be studied via three broad approaches, namely (i) direct numerical simulation (DNS) where every scale in space and time is resolved, (ii) large eddy simulation (LES) where the large scales are resolved and the effects of the small scale dynamics are modeled, (iii) Reynolds-averaged Navier Stokes (RANS) where no dynamics of turbulent state are resolved and all the turbulence action is modeled, and lastly, (iv) engineering correlations and simplifications. For this thesis, DNS and LES are the high-fidelity models for turbulent flow simulations while RANS is a low-fidelity model. The cost of simulation is in commensuration with its fidelity, as illustrated in fig. 1.1. (cf. [Peherstorfer *et al.*, 2018](#)).




Figure 1.1: Schematic illustrating the accuracy and cost trade-off between RANS, LES and DNS, particularly for wall-bounded flows.

For the discussion ahead, from the high-fidelity models, only LES is considered. In terms of accuracy and cost, it is similar to DNS. LES has been around for half a century, yet its use has been largely limited to academic cases. In LES, most of the important spatio-temporal scales are resolved, and the solution is chaotic. However, from an engineering standpoint, we are often interested in time-averaged quantities-of-interest (QoIs), which in turn usually depend only on the mean flow state. For instance, the skin friction coefficient c_f depends only on the mean velocity and viscosity fields near the wall. The more affordable approach to predict the mean solution is to use the RANS simulations, in which one solves directly for the mean state. The drawback is that modeling is required for the unclosed turbulence terms, which in general introduces systematic errors. A particularly important class of problems for which RANS models are known to have systematic errors are that of separated flows.

In spite of LES being more accurate, RANS equations are ubiquitous in any engineering




Figure 1.2: Illustration of an UQ exercise using CFD. The first row exemplifies the desired outputs from a CFD toolbox. The second row shows the functional dependence of a QoI J on the input parameters a_1, a_2 and the uncertainty in input parameters via their joint pdf $f(a_1, a_2)$. The dashed black lines indicate the nominal values of the parameters. The parametric gradient of J at that condition is another useful output. The last row shows the uncertainty in J because of uncertain inputs.

design related calculation. The quick turnaround time of RANS allows many simulations to be run at reasonable computational cost. This helps when undertaking an Uncertainty Quantification (UQ) exercise or design optimization calculation. The schematic flow shown in the first row of fig. 1.2, exemplifies outputs in addition to the QoIs themselves. The accuracy of LES comes at a high computational price, which limits its use for auxiliary design calculations such as UQ. This motivates the current work, where we explore a computationally affordable approach which aids in making LES calculations more useful for an engineering design problem. NASA CFD Vision 2030 (cf. Slotnick *et al.*, 2014) also states the need for LES techniques at high-lift conditions as an important need. A canonical LES can be $O(10^2 - 10^4)$ times as expensive as the RANS calculation for the same flow field.

Fig. 1.2 shows a typical UQ exercise where there is uncertainty in the input parameters a_1, a_2 , expressed through its joint probability distribution function (pdf). The input parameters a_1, a_2 could be flow Reynolds number Re , angle of attack to the flow α , or any other boundary condition. The contours of QoI J shows its functional dependence on the input parameters, which is not explicitly known. The QoI J can depend on the input parameters a_1, a_2 and the flow state $\mathbf{v}$, which in turn depends on input parameters and is computed through CFD. Here, J could be the skin friction c_f , wall heat transfer q_w , lift force on the airfoil etc. Typically, one is interested in some nominal baseline condition. The objective of UQ is to then find the pdf $p(J)$ given the pdf of input parameters. The pdf of a variable conveys the likelihood of what values it can take. Through this pdf one may evaluate $E(J)$, the expected value, which may differ from simply evaluating the QoI at base conditions and σ_J , the standard variability in J . Additionally, the parametric gradient $\nabla_a J$ at the baseline condition is also of use, particularly if one needs to perform design optimization for optimal

performance.

1.2 UQ methods in CFD

Uncertainties can be classified as being either aleatory or epistemic. Aleatory uncertainties are due to natural randomness, whereas the epistemic uncertainties arise due to imperfect modeling. A particular model may miss out on some detail of the physics, thus bringing out discrepancies with the “actual” truth. Aleatory uncertainty is usually modeled under the traditional frequentist interpretation of probability and sampling based techniques are used, whereas for epistemic uncertainty, probability is seen as a measure of belief and Bayesian inference type methods are employed to validate the model (cf. [Sivia & Skilling, 2006](#)). For this thesis, the focus is more on the uncertainty of the aleatoric kind.

A typical UQ exercise with traditional Monte Carlo methods involves running many thousand simulations to obtain the probability density function (pdf) of the output QoIs. The statistical error is proportional to the square-root of the number of samples, hence numerous samples are required. Multi-fidelity Monte Carlo methods have been developed where models of different fidelities are used to get estimators for expected values and variances of a QoI (cf. [Peherstorfer *et al.*, 2018](#); [Qian *et al.*, 2018](#)). The different fidelities are used as black boxes, with the outputs from each fidelity blended together to generate desired estimators of the expectation and variance. The bulk of the research effort has been aimed at finding optimal methods to blend these predictions. This is typically done using the control variate framework (cf. [Peherstorfer *et al.*, 2018](#)) and estimators are obtained by optimizing some requirements. The objective of blending the outputs is to remove the

systematic errors present in the lower fidelity model without actually altering the model. A low-fidelity model could also be solving the exact equations on a coarser grid and not just solving an approximate model (cf. [Geraci et al., 2017](#)). For even multi-fidelity Monte Carlo, many LESs would be needed, making it cost prohibitive.

Alternatively, the stochastic behavior can be expressed as a series of multivariate polynomials whose domain is an array of random variables. Common methods include the non-intrusive polynomial chaos expansion (NIPCE) and stochastic collocation method (SCM). For instance,

$$J(\mathbf{v}; \mathbf{a}(\xi)) = \sum_{i=1}^{\infty} \alpha_i \Psi_i(\xi),$$

where $\xi = \{\xi_j\}_{j=1}^d$ are the random variables modeling the input uncertainty. NIPCE uses set of orthogonal polynomials like Hermite, Legendre etc. for Ψ_i and tries to find coefficients α_i in the expansion using Galerkin projection, after sampling at certain key values. SCM assumes knowledge of α_i and uses Ψ_i as interpolation polynomials like the Lagrange polynomials to approximate the response surface, by knowing the response at certain sampled locations. These methods still require some form of sampling to be done, although much less than that in MC based methods (cf. [Hosder et al., 2007](#)). The cost is decided by order of polynomials chosen, choice of quadrature rule and number of input parameters d . In spite of all the developments to alleviate the cost, one would still have to run multiple LESs, which is cost prohibitive.

To avoid sampling completely, the polynomial chaos expansion (PCE) can be used intrusively. One then solves for both the baseline values and the uncertainty simultaneously. These methods other than being intrusive also suffer from the “curse of dimensionality” and

the overall cost has exponential dependence on d . If d is more than $O(1)$, it quickly becomes cost prohibitive, especially for LES (cf. [Najm, 2009](#)).

1.3 Parametric sensitivity estimation

An alternative to performing a full UQ exercise aimed at finding the pdf of the QoIs is to use the error propagation formula to estimate the variability in the QoI. In the simplest case of a single parameter, this is $\sigma_J = |dJ/da|\sigma_a$ where J is a QoI, a is a parameter, and σ denotes their respective standard deviation. The derivative dJ/da is the parametric sensitivity of J with respect to a . For multiple inputs, it becomes $\sigma_J^2 = (\nabla_a J)^T \Sigma_a \nabla_a J$, where Σ_a is the covariance matrix of all the input parameters in $\mathbf{a}$. The parametric sensitivity is not only useful as a surrogate for UQ, but also for design optimization.

The simplest approach to compute the parametric sensitivity is to use the finite difference method in parameter space,

$$\frac{dJ}{da_i} = \frac{J(\mathbf{v} + \Delta\mathbf{v}; \mathbf{a} + \Delta a_i \mathbf{e}_i) - J(\mathbf{v}; \mathbf{a})}{\Delta a_i} + O(\Delta a_i), \quad (1.1)$$

where dJ/da_i is the parametric gradient w.r.t. $a_i \in \mathbf{a}$. This requires running full LES cases at perturbed conditions which is infeasible for a problem with many parameters. Adjoint-based methods can then be useful, as for each QoI J , they calculate the parametric gradient $\nabla_a J$ from a single simulation (cf. [Luchini & Bottaro, 2014](#)). However, for LES these diverge exponentially due to the “butterfly effect” of chaos (cf. [Lea et al., 2000](#)). The presence of positive Lyapunov exponents makes the linearized dynamics unstable, rendering all techniques using linearized equations unsuitable for LES. For RANS, this problem does not

arise since it solves directly for the mean state, thus suppressing the chaotic dynamics. There are robust frameworks for sensitivity analysis using output-based adjoints (cf. [Fidkowski & Darmofal, 2011](#); [Jameson, 1988](#)) in the RANS paradigm.

Computing the parametric sensitivity of an output from a chaotic system is an active area of research. One of the earliest ideas is the ensemble adjoint method (cf. [Lea *et al.*, 2000](#)). The problem of ill-posedness appears if the QoI is a long time-averaged quantity (cf. [Lea *et al.*, 2000](#)), which is mostly the case in turbulent flows, except, maybe in aeroacoustics where noise events occur in burst (cf. [Kim *et al.*, 2014](#)). To circumvent this problem, methods such as the least square shadowing (LSS) (cf. [Wang, 2013, 2014](#)) have been developed. This involves finding a trajectory that “shadows” the base state and then uses it to compute meaningful derivatives of long time-averaged QoI. Finding this trajectory involves solving an optimization problem, which leads to a linear system of dimension $N \times \tau$ where N is the size of the dynamical system and τ is the temporal sampling locations. For a typical LES, N is $O(10^6)$ and τ is $O(10^3)$, which means the LSS method would be very expensive (cf. [Chandramoorthy *et al.*, 2017](#); [Ni & Wang, 2017](#)). The multiple shooting shadowing (MSS) (cf. [Blonigan & Wang, 2018](#)) method tries to reduce the cost of LSS by reformulating the optimization problem to consider only a few checkpoints in the time series, say K . This gives a linear system of size $K \times N$ where $K \ll \tau$; while cheaper, this is still infeasibly expensive for LES. Alternatively, the non-intrusive LSS (NILSS) modifies the LSS by projecting the optimization problem onto the unstable subspace of the attractor, containing all the positive Lyapunov exponents (cf. [Ni & Wang, 2017](#)). The cost of NILSS is proportional to number of positive Lyapunov exponents, which is often $\ll N$. While there is a cost reduction as compared to LSS, the number of positive LEs are not known a priori and can vary from case

to case. Contemporary work in the field is based on Ruelle’s linear response theory of non-equilibrium systems (cf. [Chandramoorthy *et al.*, 2017](#)) and the method is called space-split sensitivity (S3) (cf. [Chandramoorthy *et al.*, 2018](#); [Śliwiak & Wang, 2022, 2023](#)), which is under development.

An alternative to all these approaches that aim to find the exact chaotic sensitivity is to instead find an approximate sensitivity through the mean equations. This has been done in the context of output-based grid-adaptation for the Kuramoto-Sivashinsky equation (cf. [Larsson, 2018](#)). The extension of this underlying idea to the case of LES forms the crux of this work.

1.4 Objective of the thesis

The main objective of this thesis is to explore a multi-fidelity sensitivity algorithm (MFSA) for LES, as alluded in the previous section. The method involves running a single LES at the nominal condition and using the results of that simulation to define a constrained or corrected RANS problem, which can then be sampled (for UQ) or used to derive an adjoint (for sensitivity estimation). The method can be viewed as utilizing the single LES prediction to remove bias and using either multiple RANS evaluations or a single RANS-adjoint to estimate the variability. This idea of blending modeling fidelities is conceptually similar to multi-fidelity Monte Carlo, but the present approach is inherently intrusive into the lower-fidelity model.

The accuracy of the method is assessed for an LES of a shock/boundary layer interaction (SBLI) flow configuration. The flow serves as a good test case because it has a combination

of wall bounded and separated flow. While RANS is arguably very good to capture the former, it is known to be inaccurate for the latter kind of flow. In related work in the same research group (cf. [Arias-Ramírez, 2022](#)), a similar method based on the algebraic Prandtl mixing-length turbulence model was applied to LES of flow over NACA 0012 airfoil and the influence of assumptions like frozen turbulence and other numerical errors were assessed. This related work showed reasonable accuracy of the predicted sensitivity but with clear room for improvement. The research hypothesis of the present work was that we should gain accuracy by moving to a better RANS turbulence model. The present work thus uses the Spalart-Allmaras turbulence model which solves an additional transport equation, requiring additional inference from the LES data.

1.5 Publications

Peer reviewed journals:

- (Under review) Oberoi N., Arias-Ramírez W., and Larsson J., 2023. Multi-fidelity parametric sensitivity estimation for large eddy simulation with the Spalart-Allmaras model, *Journal of Turbulence*.
- (In preparation) Oberoi N. and Larsson J., 2023. Effective inference of turbulent diffusivity from large eddy simulation, *Journal TBD*.
- (Under review) Arias-Ramírez W., Oberoi N., and Larsson J., 2023. A multi-fidelity approach to sensitivity estimation in large eddy simulation, *AIAA Journal*.
- Larsson, J., Kumar, V., Oberoi, N., Renzo, M. D., and Pirozzoli, S., 2022. Large-

eddy simulations of idealized shock/boundary-layer interactions with crossflow. *AIAA Journal*, 60(5), 2767-2779.

- Di Renzo, M., Oberoi, N., Larsson, J., and Pirozzoli, S., 2022. Crossflow effects on shock wave/turbulent boundary layer interactions. *Theoretical and Computational Fluid Dynamics*, 36(2), 327-344.

Conference presentations:

- Oberoi, N., Arias-Ramírez, W., Larsson, J., 2021. A computationally affordable multi-fidelity approach to parametric studies. In *APS Division of Fluid Dynamics Meeting* (pp. T02-009)
- Oberoi, N., Arias Ramírez, W., Larsson, J., 2019. Inferring an effective eddy viscosity from High Fidelity Data. In *APS Division of Fluid Dynamics Meeting* (pp. NP05-022).

Chapter 2: Modeling and numerical background

This work deals with compressible fluid flows over a wall with impinging shocks. The flow is governed by the compressible Navier-Stokes equations, encompassing the conservation laws. The fluid is assumed to be Newtonian with Fourier heat conduction and ideal gas relation is followed. The state variable $\mathbf{v} = [\rho, \mathbf{u}, T, p]$, where ρ denotes the density field, $\mathbf{u}$ denotes the velocity with three components $\{u_i\}_{i=1}^3$, p is the pressure field and T is the temperature field. For this thesis, an equivalence between the coordinate description (x_1, x_2, x_3) and (x, y, z) , and velocity components (u_1, u_2, u_3) and (u, v, w) is followed. For the discussion of equations in the thesis, Einstein indicial notation is used, where twice-repeating indices imply summation.

2.1 High-fidelity model: LES

The compressible Navier-Stokes equations in a Cartesian coordinate system are,

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \frac{\partial (\rho u_j)}{\partial x_j} &= 0, \\ \frac{\partial (\rho u_i)}{\partial t} + \frac{\partial (\rho u_i u_j)}{\partial x_j} + \frac{\partial p}{\partial x_i} - \frac{\partial \sigma_{ij}}{\partial x_j} &= 0, \text{ and} \\ \frac{\partial (\rho e_0)}{\partial t} + \frac{\partial (\rho e_0 u_j)}{\partial x_j} + \frac{\partial (p u_j)}{\partial x_j} + \frac{\partial q_j}{\partial x_j} - \frac{\partial (u_i \sigma_{ij})}{\partial x_j} &= 0. \end{aligned} \tag{2.1}$$

Here $\rho, \mathbf{u}, p, T$ are the instantaneous density, velocity, pressure and temperature fields respectively over a domain $\Omega \subseteq \mathbb{R}^3$, σ_{ij} is the viscous stress tensor and q_j is Fourier heat flux (conduction) with

$$\sigma_{ij} = 2\mu \left(S_{ij} - \frac{1}{3} S_{kk} \delta_{ij} \right), \quad q_j = \kappa \frac{\partial T}{\partial x_j}, \quad S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right).$$

S_{ij} is the symmetric part of the velocity deformation tensor, μ is the dynamic viscosity, which is obtained from the power law $\mu/\mu_{\text{ref}} = (T/T_{\text{ref}})^{3/4}$ with subscript ref denoting a reference state, $\kappa = \mu c_p / \text{Pr}$, c_p is the specific heat at constant pressure and Pr is the Prandtl number, ratio of momentum and thermal diffusivities. Total specific energy $e_0 = e + u_k u_k / 2$, $e = c_v T$ is the specific internal energy and c_v is specific heat at constant volume. If the fluid is further assumed to be calorifically perfect, then c_p, c_v are constant with specific heat ratio $\gamma = c_p / c_v$

and gas constant $R = c_p - c_v$. The equations are closed by using the equation of state

$$p - \rho RT = 0. \quad (2.2)$$

For compactness, the system of equations is denoted as $N(\mathbf{v}; \mathbf{a}) = 0$. Here $\mathbf{a}$ is the set of input parameters such as Reynolds number Re , Mach number Ma , wall temperature T_w , angle of attack α , oblique shock angle β , etc. These equations describe the full dynamics of the system. Given suitable boundary conditions, they form a completely deterministic system.

Turbulence is a complex and chaotic state of a fluid flow, characterized by dynamics that span a wide range of length and time scales. In the turbulent regime, solving these equations, resolving all the scales without any further modeling is called a direct numerical simulation (DNS) of the flow configuration. The largest scale in a turbulent flow is often decided by the geometry and the flow configuration. These turbulent structures, also called large eddies, break into smaller structures or eddies via the energy cascade (cf. [Pope, 2000](#)). At the smallest scales, the viscous effects balance the inertial effects and dissipate the energy cascaded down from the large eddies. The ratio of the large and small scales appearing in a turbulent flow is proportional to the flow Reynolds number Re , and hence the cost of a DNS increases rapidly with the Reynolds number.

In large eddy simulations (LES), only the largest scales are resolved to alleviate the cost of the simulation. The small scale eddies are assumed to be universal, not affected by the flow geometry and hence modeled (cf. [Pope, 2000](#); [Sagaut, 2006](#); [Toosi, 2019](#)). This is typically done by adding μ_{sgs} and κ_{sgs} to the corresponding fluid properties, where sgs

denotes the sub-grid scale i.e. the eddies that were not resolved. In this thesis, LES is viewed as a coarse-grid representation of the DNS, without any explicit sub-grid scale model i.e. $\mu_{sgs} = 0$ and $\kappa_{sgs} = 0$. This is also called implicit-LES in literature (cf. [Toosi & Larsson, 2020](#)). This is especially important for simulating wall-bounded flows (cf. [Larsson *et al.*, 2022](#)), which is the application in focus for this work. LES though cheaper than a DNS, is still a computationally intensive exercise in itself. From an LES of a turbulent flow, the time series of the state variables are determined. These are then averaged (in statistically homogeneous directions) to obtain the mean state and the desired quantities of interest (QoIs) are computed.

2.2 Low-fidelity model: RANS

Most of the QoIs that are useful from an engineering standpoint often only require the averaged fluid state. That being the case, eqn. (2.1) is statistically averaged to obtain the equations governing the averaged fluid state $\bar{\mathbf{v}}$ in eqn. (2.3). Each flow field can be decomposed into its mean and fluctuating components, with two kinds of means used for simplification of the equations, $\overline{(\cdot)}$ and $\widetilde{(\cdot)}$ are the ensemble (Reynolds) and density-weighted (Favré) averages respectively. Here, $\widetilde{(\cdot)} = \overline{\rho(\cdot)}/\bar{\rho}$ denotes the Favré average, while $(\cdot)' = (\cdot) - \overline{(\cdot)}$ and $(\cdot)'' = (\cdot) - \widetilde{(\cdot)}$ are the corresponding fluctuations. The averaged equations

with such a decomposition are expressed as

$$\begin{aligned}
\frac{\partial \bar{\rho}}{\partial t} + \frac{\partial (\bar{\rho} \tilde{u}_k)}{\partial x_k} &= 0, \\
\frac{\partial (\bar{\rho} \tilde{u}_i)}{\partial t} + \frac{\partial (\bar{\rho} \tilde{u}_i \tilde{u}_j)}{\partial x_j} + \frac{\partial \bar{p}}{\partial x_i} - \frac{\partial}{\partial x_j} (\bar{\sigma}_{ij} - \bar{\rho} R_{ij}) &= 0, \\
\frac{\partial (\bar{\rho} \tilde{e}_0)}{\partial t} + \frac{\partial (\bar{\rho} \tilde{u}_j \tilde{e}_0)}{\partial x_j} + \frac{\partial (\tilde{u}_j \bar{p})}{\partial x_j} - \frac{\partial}{\partial x_j} (\bar{\sigma}_{jk} \tilde{u}_k + \overline{\sigma_{jk} u_k''}) \\
&+ \frac{\partial}{\partial x_j} \left(\bar{q}_j + h_{Tj} + \bar{\rho} R_{jk} \tilde{u}_k + \frac{1}{2} \overline{\rho u_j'' u_k'' u_k''} \right) = 0, \text{ and} \\
\bar{p} - \bar{\rho} R \tilde{T} &= 0.
\end{aligned} \tag{2.3}$$

These equations are referred to as the Reynolds-averaged Navier-Stokes (RANS) equations (cf. [Wilcox et al., 1998](#)), and are represented in shorthand as, $M(\bar{\mathbf{v}}, \mathbf{V}) = 0$, with $\bar{\mathbf{v}} = [\bar{\rho}, \tilde{\mathbf{u}}, \tilde{T}, \bar{p}]$ and $\mathbf{V} = [R_{ij}, h_{Tj}, \overline{\sigma_{jk} u_k''}, \overline{u_j'' u_k'' u_k''}]$. Here, the Reynolds stress tensor $R_{ij} = \overline{u_i'' u_j''}$ and turbulent heat flux $h_{Tj} = \overline{u_j'' h''} = c_p \overline{\rho u_j'' T''}$. The average of the viscous stress $\bar{\sigma}_{ij}$ and heat conduction $\bar{q}_j$ are approximated as,

$$\bar{\sigma}_{ij} \approx 2\mu(\tilde{T}) \left(\tilde{S}_{ij} - \frac{1}{3} \tilde{S}_{kk} \delta_{ij} \right), \quad \bar{q}_j \approx \kappa(\tilde{T}) \frac{\partial \tilde{T}}{\partial x_j},$$

where μ, κ are determined using the power law as functions of $\tilde{T}$. The choice of decomposition allows the RANS equations in eqn. (2.3) to have a similar structure as that of the Navier-Stokes equations in eqn. (2.1). RANS has been the backbone of engineering design calculations under turbulent flow regimes (cf. [Bush et al., 2019](#)). In addition, by averaging, RANS eliminates the chaotic dynamics which makes it amenable to adjoint-based methods, further

making the design process more efficient and effective (cf. [Fidkowski & Darmofal, 2011](#)).

However, the RANS equations still do not form a closed system, due to the additional correlation terms that appear in the equation. This is a consequence of non-linear terms in the Navier-Stokes equations. The additional terms are often also referred to as turbulent fluxes and require some form of modeling, so that the RANS equations can be used in a predictive sense. This is the so-called closure problem of RANS.

2.2.1 RANS closure models and Boussinesq approximation

The turbulent fluxes for a quantity ϕ typically appear in the RANS equations as $\widetilde{u_j''\phi''}$. On an average sense, these terms account for the convective transport of ϕ from different eddies present in the turbulence spectrum. In spirit, this is similar to how at molecular level the molecules would diffuse the field ϕ through random vibrations, albeit turbulence does it at a more macro level. This motivates the gradient transport hypothesis for turbulent fluxes (cf. [Pope, 2000](#)), where,

$$\widetilde{u_j''\phi''} = -\Gamma_t \frac{\partial \phi}{\partial x_j}, \quad (2.4)$$

with Γ_t being the corresponding turbulent diffusivity. Under this assumption, R_{ij} can be modeled against the mean shear rate linearly using ν_t as,

$$R_{ij} = -2\nu_t S_{ij}^d + \frac{R_{kk}}{3} \delta_{ij}, \quad S_{ij} = \frac{1}{2} \left(\frac{\partial \widetilde{u}_i}{\partial x_j} + \frac{\partial \widetilde{u}_j}{\partial x_i} \right). \quad (2.5)$$

This is also called the Boussinesq approximation. The turbulent thermal flux $q_{T_k} = \widetilde{u_k'' T''}$, appearing in the energy equation is modeled against meant thermal gradient $g_k = \partial \tilde{T} / \partial x_k$ as

$$q_{T_k} = -\lambda_t g_k. \quad (2.6)$$

Γ_t denotes the set of turbulent diffusivities incorporated e.g., ν_t to model the Reynolds stress and λ_t to model turbulent thermal flux. RANS models can have an algebraic model to compute ν_t or use transport equations of auxiliary variables which are eventually used to compute ν_t . For instance, the auxiliary variables for the $k - \epsilon$ RANS model (cf. Wilcox, 1998) are $k = R_{jj}/2$, the turbulent kinetic energy, and ϵ , the dissipation rate. Alternatively, the Spalart-Allmaras RANS model (cf. Spalart & Allmaras, 1992; Wilcox, 1998) solves for $\tilde{\nu}$ which is used to compute ν_t . For λ_t , turbulent Prandtl number $Pr_t = \nu_t / \lambda_t$ is used. Typically, it is set to be 0.9 throughout the flow field.

The remaining terms in eqn. (2.3) are modeled as

$$\widetilde{u_i'' \sigma_{ij}} + \frac{1}{2} \widetilde{u_j'' u_k'' u_k''} = - \left(\nu + \frac{\nu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j}, \quad (2.7)$$

in the presence of a transport equation for k , otherwise are dropped from the equation. Here, σ_k is a model coefficient that appears in the transport equation for k , if used in the RANS model.

For brevity, let Π be the set of auxiliary variables used to obtain $\Gamma_t = [\nu_t, \lambda_t]$ elements. This closed RANS system is denoted by $M(\bar{\mathbf{v}}, \Pi; \mathbf{a}) = 0$, which serves as a low-fidelity

predictive model for turbulent flow with input parameters $\mathbf{a}$.

In the RANS equations, the turbulent fluxes R_{ij} and q_{Tj} appear in the equations after being scaled by $\bar{\rho}$ and $\bar{\rho}c_p$, respectively. In accordance, the turbulent diffusivities are rescaled as $\mu_t = \bar{\rho}\nu_t$ and $\kappa_t = \bar{\rho}c_p\lambda_t$. The sub-grid scale model variables μ_{sgs} and κ_{sgs} are developed in a similar spirit. In RANS, all the turbulence is modeled, in LES turbulence associated with small scales is modeled. Arguably, both LES and RANS are similar in spirit, with the extent of turbulence modeling deciding the fidelity.

Even with the extra equations that the RANS system solves, the overall computational cost to compute the averaged flow state is orders of magnitude cheaper than compared to that of LES. Often, the flow state is statistically stationary i.e., the mean state does not vary in time, as a result in RANS only the steady state equations are needed to be solved and the full time series is not desired. Also, the flow state can be statistically homogeneous in spatial dimensions as well, reducing the overall dimension of the system to be solved. These two together bring down the computational cost of RANS significantly. However, as is the case with most low-fidelity models, they work well for flow regimes where they have been calibrated, but can struggle to replicate the accuracy in other flow regimes. RANS is known to have reduced accuracy when simulating wall-bounded flows where the flow has a tendency to separate from the wall in presence of adverse pressure gradients.

2.3 Proposed multi-fidelity sensitivity algorithm

In this section, the proposed multi-fidelity sensitivity algorithm (MFSA) is devised. Running one LES at nominal base state $\mathbf{a}_b$ provides us $\mathbf{v}^{\text{LES}}$ and subsequently the QoIs




Figure 2.1: This illustrates the objective of the work. The solid curves show the parametric profile of a QoI J w.r.t parameter a . The dashed lines tangent to these curves mark their slope at some nominal condition a_b . We expect to replicate the black dashed line using the blue dashed line from the proposed MFSA.

$J(\bar{\mathbf{v}}^{\text{LES}}; \mathbf{a}_b)$. Next, by running LES at different $\mathbf{a}$, one can obtain the parametric response (black curve in fig. 2.1) as well as the parametric sensitivity at $\mathbf{a}_b$ (dashed black tangent line in fig. 2.1). However, doing this for an LES is computationally prohibitive, as discussed previously. One can repeat the same process for RANS and obtain its parametric response (the red curve and correspondingly the red dashed line tangent to the parametric curve at $\mathbf{a}_b$ in fig. 2.1). Given the non-chaotic behavior and a single RANS simulation being very cheap, the process of finding the parametric sensitivity is amenable to plenty of methods. However, there can be systematic differences between parametric responses of LES and RANS (as visualized in fig. 2.1). They may predict different flow regimes for the same parameter value. For instance, consider a canonical boundary layer under a non-zero pressure gradient, for which LES and RANS may predict attached and separated flow, respectively, or vice versa. As a consequence, we may end up with different parametric sensitivities as well.

In the proposed framework, we want to address this possible mismatch at nominal conditions. The key idea in MFSA is to find a RANS system around the LES data at nominal parameters $\mathbf{a}_b$. The RANS system should be able to replicate the LES data at that condition, and using suitable sensitivity estimation methods, predict a parametric sensitivity (the blue dashed tangent line in fig. 2.1) which closely approximates the true LES sensitivity.

The first step in replicating the base LES solution using RANS is to infer the auxiliary variables Π from the LES solution, resulting in Π^{LES} . This process is discussed in sec. 2.4.1. However, since the closure modeling is imperfect, $M(\bar{\mathbf{v}}^{\text{LES}}, \Pi^{\text{LES}}; \mathbf{a}_b) \neq 0$, i.e., the LES data with inferred turbulence model variables does not exactly satisfy the RANS system. This may happen due to imperfect modeling of the underlying assumption. The inferred fields will be optimal but not perfect.

To ensure that there is a RANS-like equation that satisfies LES data at nominal condition, the equations need to be augmented. We propose to modify the RANS system as

$$M(\bar{\mathbf{v}}, \Pi; \mathbf{a}) - M(\bar{\mathbf{v}}^{\text{LES}}, \Pi^{\text{LES}}; \mathbf{a}_b) = 0. \quad (2.8)$$

We refer to eqn. (2.8) as LES-constrained RANS in this work. This LES-constrained RANS system at nominal conditions replicates $\bar{\mathbf{v}}^{\text{LES}}$ and holds true for any choice of inferred Π^{LES} , by construction. This subtraction of the RANS residual when operating on the LES solution is similar to the method of manufactured solutions often used in CFD code verification (Oberkampf & Trucano, 2002).

The mean mass conservation equation in the RANS system requires no closure, so it should have no residual of LES data. For the momentum conservation equation, the extra term accounts for the discrepancy between the divergence of the R_{ij} from LES data and the one modeled using the inferred ν_t^{LES} and S_{ij}^d as per eqn. (2.5). This approach is similar to many machine learning methods using DNS/LES data to improve RANS models (cf. Brunton *et al.*, 2020; Wang *et al.*, 2017; Wu *et al.*, 2019). The proposed MFSA involves running one LES, obtaining $\bar{\mathbf{v}}^{\text{LES}}$ and inferring Π^{LES} , and then solving eqn. (2.8) at different perturbed states to evaluate the sensitivity dJ/da_i by using a curve fit. Alternatively, one may linearize eqn. (2.8) about the LES data at nominal conditions and solve this tangent equation, or use it to derive the adjoint equation and solve that.

2.4 Choice of RANS model

In this work, the Spalart-Allmaras (SA) RANS model ([Allmaras & Johnson, 2012](#); [Spalart & Allmaras, 1992](#)) is used with $\Pi = [\tilde{\nu}, \text{Pr}_t]$, and $\tilde{\nu}$ satisfying

$$\begin{aligned} \frac{\partial \tilde{\nu}}{\partial t} + u_j \frac{\partial \tilde{\nu}}{\partial x_j} - c_{b1} (1 - f_{t2}) \tilde{S} \tilde{\nu} + \left[c_{w1} f_w - \frac{c_{b1}}{\kappa^2} f_{t2} \right] \left(\frac{\tilde{\nu}}{d} \right)^2 \\ - \frac{1}{\sigma} \left[\frac{\partial}{\partial x_j} \left((\nu + \tilde{\nu}) \frac{\partial \tilde{\nu}}{\partial x_j} \right) + c_{b2} \frac{\partial \tilde{\nu}}{\partial x_i} \frac{\partial \tilde{\nu}}{\partial x_i} \right] = 0. \end{aligned} \quad (2.9)$$

Here, $\tilde{\nu}$ relates to the momentum eddy viscosity ν_t as

$$\nu_t = \tilde{\nu} f_{\nu 1}, \quad (2.10)$$

with $f_{\nu 1} = \chi^3 / (\chi^3 + c_{\nu 1}^3)$ and $\chi = \tilde{\nu} / \nu$. The default values of the coefficients and the definitions of terms are used ([Spalart & Allmaras, 1992](#)). To account for turbulent the heat transfer, the standard SA model uses $\text{Pr}_t = 0.9$. Alternatively, the base LES data can be used to infer Pr_t as a field variable. Comparison between these two choices is discussed when comparing results from the proposed MFSA.

In the sister work to this thesis, [Arias-Ramírez \(2022\)](#) applied the MFSA with the Prandtl mixing-length turbulence model to an LES of NACA 0012 airfoil. The Prandtl mixing-length model (cf. [Wilcox, 1998](#)) is an algebraic RANS model. The results from [Arias-Ramírez \(2022\)](#) suggest that the MFSA can work but needs an improvement in the RANS closure model. Accordingly, the SA-RANS model is chosen which has a transport equation for one of its auxiliary variables, and is known to be useful for aerospace applications. In general,

transport equation based RANS models provide a better accuracy as compared to algebraic RANS models (cf. [Wilcox, 1998](#)).

2.4.1 Inference of auxiliary variables

The auxiliary variables for the chosen SA-RANS model are $\Pi = [\tilde{\nu}, \text{Pr}_t]$, and are inferred from ν_t and λ_t (see eqn. (2.10) and definition of Pr_t). These eddy diffusivities are in turn obtained from LES data by minimizing the error in the gradient-diffusion hypotheses of eqns. (2.5) and (2.6).

We find ν_t^{LES} by minimizing the Frobenius norm of the error tensor $E_{ij}(\nu_t) = R_{ij}^d + 2\nu_t S_{ij}^d$, where S_{ij}^d and R_{ij}^d are obtained from LES data. Additionally, we require the inferred eddy viscosity fields to be smooth in spite of statistical noise in R_{ij} and u_i fields emanating from imperfect averaging. Further, when $|S| = \sqrt{S_{ij}^d S_{ij}^d} \rightarrow 0$, the minimization problem becomes ill-posed, typically observed in the freestream region of the flow. We regularize this problem by using an indicator function

$$\chi_s = \begin{cases} 1 & |S|^2 < \epsilon_s \\ 0 & |S|^2 \geq \epsilon_s. \end{cases} \quad (2.11)$$

In regions where $\chi_s = 1$, $\nu_t^{\text{LES}} = \nu_t^\infty = 1.5\nu_{\text{ref}}$, with ν_{ref} being the reference fluid kinematic viscosity. This is similar to the freestream boundary condition for $\tilde{\nu}$ in a typical Spalart-Allmaras RANS simulation. These additional constraints are combined into the cost functional

$$\mathcal{J}_{\nu_t} = \int_{\Omega} \left[\frac{\omega_s}{U_{\text{ref}}^2} \frac{\partial \nu_t}{\partial x_k} \frac{\partial \nu_t}{\partial x_k} + \frac{1}{U_{\text{ref}}^4} E_{ij} E_{ij} + \frac{\chi_s}{U_{\text{ref}}^2 \delta_{\text{ref}}^2} (\nu_t - \nu_t^\infty)^2 \right] d\Omega, \quad (2.12)$$

where ω_s is the relative weight of the smoothing constraint. The minimum of $\mathcal{J}_{\nu_t}$ is obtained by setting $\delta\mathcal{J}_{\nu_t} = 0$. This gives us a partial differential equation

$$-U_{\text{ref}}^2 \omega_s \frac{\partial^2 \nu_t}{\partial x_k \partial x_k} + 2S_{pq}^d S_{pq}^d \nu_t + R_{ij}^d S_{ij}^d + \frac{U_{\text{ref}}^2 \chi_s}{\delta_{\text{ref}}^2} (\nu_t - \nu_t^\infty) = 0, \quad (2.13)$$

whose solution is ν_t^{LES} . The equation is solved with homogeneous Dirichlet boundary conditions

$\nu_t^{\text{LES}} = 0$ at walls and homogeneous Neumann boundary condition everywhere else. We

similarly infer λ_t^{LES} by minimizing the Euclidean norm of the modeling error vector $e_k(\lambda_t) = q_{T_k} + \lambda_t g_k$. Here, the fields q_{T_k} and g_k are sampled from the LES data at baseline conditions.

The cost functional for λ_t with similar smoothness and freestream constraints is

$$\mathcal{J}_{\lambda_t} = \int_{\Omega} \left[\frac{\omega_s}{U_{\text{ref}}^2} \frac{\partial \lambda_t}{\partial x_k} \frac{\partial \lambda_t}{\partial x_k} + \frac{1}{U_{\text{ref}}^2 T_{\text{ref}}^2} e_k^{\text{LES}} e_k^{\text{LES}} + \frac{\chi_g}{U_{\text{ref}}^2 \delta_{\text{ref}}^2} (\lambda_t - \lambda_t^\infty)^2 \right] d\Omega, \quad (2.14)$$

with the indicator function

$$\chi_g = \begin{cases} 1 & |g|^2 < \epsilon_g, \\ 0 & |g|^2 \geq \epsilon_g. \end{cases} \quad (2.15)$$

This gives us a partial differential equation

$$-T_{\text{ref}}^2 \omega_s \frac{\partial^2 \lambda_t}{\partial x_k \partial x_k} + g_k g_k \lambda_t + g_k q_{T_k} \frac{\chi_g T_{\text{ref}}^2}{\delta_{\text{ref}}^2} (\lambda_t - \lambda_t^\infty) = 0. \quad (2.16)$$

The solution to this equation is λ_t^{LES} and is solved with similar boundary conditions as of

eqn. (2.13). Here, T_{ref} , U_{ref} , and δ_{ref} are reference values for temperature, velocity and

length, respectively. For eqns. (2.13) and (2.16), $\omega_s = 0.001$ satisfactorily removes the statistical noise. For this work, the thresholds were calibrated to $\epsilon_s = 0.001U_{\text{ref}}^2/\delta_{\text{ref}}^2$ and $\epsilon_g = 0.001T_{\text{ref}}^2/\delta_{\text{ref}}^2$. Lastly, we ensure $\nu_t^{\text{LES}}, \lambda_t^{\text{LES}} \geq 0$ by setting $\nu_t^{\text{LES}} = \max(\nu_t^{\text{LES}}, 0)$ and $\lambda_t^{\text{LES}} = \max(\lambda_t^{\text{LES}}, 0)$.

Near the wall, $\nu_t^{\text{LES}}, \lambda_t^{\text{LES}} \sim h^3$, where h is the wall distance height. By definition, this implies that $\tilde{\nu}^{\text{LES}} \sim h$ and $\text{Pr}_t^{\text{LES}} \sim 1$ near the wall. To ensure the correct near-wall behavior even on coarse LES grids, the correct near-wall behaviors are enforced for $\tilde{\nu}^{\text{LES}}$ and Pr_t^{LES} after their values are computed from the LES data. The inference process is discussed in detail in chapter 4.

2.5 CFD solvers

The solver needs to be able to perform an LES and simulate the RANS equations. The in-house codes `hybrid` and `tortuga` are used for the same. Different solvers are used for each fidelity.

2.5.1 hybrid

For LES, we use an in-house finite-difference code `hybrid`, based on a hybrid numerical scheme. Near shocks, a fifth-order-accurate weighted essentially non-oscillatory (WENO) scheme with Roe flux-splitting is used to approximate the inviscid fluxes, whereas a sixth-order-accurate central difference scheme in the split form is used in the remainder of the domain. The viscous terms are treated by a conservative scheme that has the resolution characteristics of a sixth-order scheme, and the system is integrated in time using an explicit

fourth-order-accurate Runge-Kutta method. The code uses a gentle sixth-order artificial dissipation term to reduce numerical noise. This code has been extensively validated and used for production for LES data for SBLI flow configurations (cf. [Johnsen *et al.*, 2010](#); [Larsson *et al.*, 2022](#); [Volpiani *et al.*, 2018](#)).

2.5.2 *tortuga*

For the RANS (standalone and LES-constrained) equations we use an in-house finite-volume code *tortuga*. A hexahedral structured grid is used with HLLC Riemann solver with third-order WENO scheme for inviscid fluxes and second-order viscous flux reconstruction on the cell faces. The steady state equations are solved in a pseudo-time-marching manner with local time stepping using an explicit third-order-accurate Runge Kutta scheme. The code was validated on the 2D zero pressure gradient flat plate validation case that is a part of the NASA Turbulence Modeling Resource (cf. [Rumsey, 2013](#)), which is discussed in sec. [2.6](#).

2.5.3 Finite-volume numerics

In a finite-volume setting, the grid can either be cell-centered, vertex-centered or cell-vertex. A cell-centered approach is used in the *tortuga* framework. For a field Ψ , which can be a scalar, vector, or a general tensor derived from the state variable $\mathbf{v}$, the discrete data is stored as cell averages on the cell centers as,

$$\int_{\Delta_{i,j}} \Psi d\Delta_{i,j} = \Psi_{i,j} \Delta_i \Delta_j + O\left(\Delta_i^2, \Delta_j^2\right) \quad (2.17)$$




Figure 2.2: A canonical 2-D finite-volume grid illustrating the cell centers where the data is stored, the cell faces and the cell sizes.

The gray shaded region in fig. 2.2 marks the cell area (or volume in 3-D) $\Delta_{i,j}$, the black dot marks the cell center where the average values are stored. Since the grid is structured there is a local i, j ordered pair associated with each cell center, for the example illustrated in fig. 2.2. The average value is taken to be the same as the cell center value which has a numerical error proportional to the grid spacing square (mid-point rule for integration).

Generally, a flux $\mathcal{F}$ which can be a vector or a tensor appears in the conservation laws as a divergence (cf. eqn. (2.1), eqn. (2.3)). Using Gauss' divergence theorem, this can be converted to surface fluxes as

$$\int_{\Delta_{i,j}} \text{div} \mathcal{F} d\Delta_{i,j} = \int_{\partial\Delta_{i,j}} \mathcal{F} \cdot \mathbf{n} ds_{i,j}. \quad (2.18)$$

The face normal $\mathbf{n}$ is always taken to be pointing outwards, as illustrated by the red arrows on the faces of cell $\Delta_{i,j}$ in fig. 2.2. The surface flux is discretized by approximating it to be face center value throughout (using mid-point rule of integration) as

$$\begin{aligned} \int_{\partial\Delta_{i,j}} \mathcal{F} \cdot \mathbf{n} ds_{i,j} &= (\mathcal{F} \cdot \mathbf{n})_{i+1/2,j} \Delta_j + (\mathcal{F} \cdot \mathbf{n})_{i-1/2,j} \Delta_j \\ &\quad + (\mathcal{F} \cdot \mathbf{n})_{i,j+1/2} \Delta_i + (\mathcal{F} \cdot \mathbf{n})_{i,j-1/2} \Delta_i \\ &\quad + O\left(\Delta_i^2, \Delta_j^2\right). \end{aligned} \quad (2.19)$$

The face centers are marked by the red outward pointing triangles in fig. 2.2. Face between cell $\Delta_{i,j}$ and $\Delta_{i+1,j}$ is denoted by subscript $i + 1/2, j$ and so on for other faces. The projected flux $\mathcal{F} \cdot \mathbf{n}$, which will be a function of any field Ψ , on these face centers, is reconstructed from the cell-center data.

The discussion here is summarized from [LeVeque \(2002\)](#) which focuses on use of finite-volume methods for hyperbolic problems. These usually contain the inviscid fluxes which are the convective terms in eqn. (2.1) and eqn. (2.3). These fluxes are constructed using a combination of WENO interpolation and HLLC approximate Riemann solver (cf. [Johnsen et al., 2010](#); [Shu, 1999](#)) in tortuga.

Viscous fluxes appear as

$$\frac{\partial}{\partial x_j} \left(a \frac{\partial \Psi}{\partial x_j} \right),$$

where a is the non-uniform isotropic diffusion coefficient. Upon finite-volume discretization, $a \partial \Psi / \partial n$ is needed to be evaluated at each face. Here, $\partial \Psi / \partial n = \partial \Psi / \partial x_k n_k$ where n_k are the components of the outward-face-normal $\mathbf{n}$. At each face, this can be computed using only the adjacent cell center value the outward-face-normal points to. For instance, at face $i + 1/2, j$

$$\left. \frac{\partial \Psi}{\partial n} \right|_{i+1/2, j} = 2 \frac{\Psi_{i+1, j} - \Psi_{i, j}}{\Delta_i + \Delta_{i+1}} + O \left(\Delta_i^2, \Delta_{i+1}^2, 1 - \frac{\Delta_{i+1}}{\Delta_i} \right). \quad (2.20)$$

However, if there is anisotropic diffusion where we have a diffusion matrix A_{pq} , then the diffusive terms are

$$\frac{\partial}{\partial x_p} \left(A_{pq} \frac{\partial \Psi}{\partial x_q} \right).$$

This upon discretization needs $n_p A_{pq} \partial \Psi / \partial x_q$ to be evaluated at each face. This requires information of the full gradient of Ψ and hence is slightly more complex to compute. In




Figure 2.3: Schematic diagram of the NASA TMR flat plate test case illustrating the boundary conditions and station of interest at which the code *tortuga* is compared.

chapter 4, sec. 4.4 we need to construct a matrix of these linear operations - isotropic and anisotropic diffusion terms acting on the array of the discrete field variables. The added complexity of computing the anisotropic diffusion term makes it difficult to compute an effective matrix for the same. It serves as a prelude to the discussion ahead in sec. 4.4.

2.6 Solver verification

The verification and validation of the SA-RANS solver of the in house code *tortuga* was done on the NASA TMR flat plate case (cf. Rumsey, 2013). Fig. 2.3 shows the schematic sketch of the test case. At the inflow we have uniform inflow of $Ma = 0.2$ and $Re =$

$\rho_\infty U_\infty L / \mu = 5 \times 10^6$. In general, μ varies with temperature T , however in this case, T is constant everywhere, and so is μ . Ahead of the flat plate leading edge, symmetry boundary condition is imposed for all variables. The wall is adiabatic and solid i.e., enforces no-slip boundary condition on the velocity components. At the outflow, there are subsonic outflow conditions, so the inflow pressure is enforced, since it is a zero pressure gradient flow, and all other variables are interpolated from the interior of the domain. At the farfield boundary away from the wall, values are extrapolated from the neighboring cells. The wall starts at $x/L = 0$.

Grid name	n_x	n_y	Δy_w^+
G1	32	24	1.527
G2	64	48	0.744
G3	128	96	0.370
G4	256	192	0.185
G5	512	384	0.092

Table 2.1: Sequence of the 5 grids used for TMR flat plate validation case. The Δy^+ at the wall, at the station of interest $x/L = 0.97$ are computed a posteriori, and indicates the quality of grid for wall-bounded flows.

There is a sequence of 5 grids as tabulated in table 2.1, with stretching in both the x and y directions. The grid is coarsened at each level by a factor of 2 everywhere, as we move from G5 to G1. In x , the grid is fine near the start of flat plate (cf. fig. 2.4) and we split the domain into two blocks about $x = 0$. The block boundary is illustrated in fig. 2.4. In y , the grid spacing varies linearly in the wall-normal direction. The height of the first cell near the wall is documented in table 2.1 in plus units at the station of interest. In wall bounded flows near-wall, wall-normal grid spacing is mentioned in “plus” or “wall” units




Figure 2.4: Δx and Δy distributions on the finest grid G5. Each coarser level of grid doubles the grid spacing in x and y at each point.




Figure 2.5: The first row shows the contours of streamwise velocity u along with the boundary layer height δ (solid black line). The second row shows the contours of μ_t , the turbulent dynamic viscosity, obtained from the SA model. The values are taken from the finest grid G5.

(cf. Pope, 2000). The wall velocity scale $u_\tau = \sqrt{\tau_w/\rho_w}$, where $\tau_w = \mu_w \partial u / \partial y|_w$ is the wall shear stress, and ρ_w, μ_w are the density and dynamic viscosity at the wall, respectively. The plus unit length scale is $\delta_\nu = \mu_w / (\rho_w u_\tau)$.

Typically, for RANS solvers, the first wall-normal cell should have a height $y^+ \lesssim 1$. The near-wall requirements vary from model to model, with the $k-\omega$ RANS model needing a super fine grid near the wall and SA allowing for relatively relaxed requirements (cf. Wilcox, 1998). This depends on the auxiliary variables the model includes and their near-wall behavior.

Fig. 2.5 shows the streamwise velocity u and turbulent dynamic viscosity μ_t , scaled by the freestream values. The boundary layer height, illustrated by δ (black curve in fig. 2.5), starts growing from $x/L = 0$, the start of the plate. The curve of δ is the locus of all the points where streamwise velocity u is 99% of its freestream value U_∞ .

Next, we look at grid convergence of the results, by examining the c_f v/s x/L curves, and wall-normal profiles of u (outer and wall scaling), μ_t at the station of interest (cf. the blue line in fig. 2.3), for each grid. Also, scalar quantities such as c_f and δ at that location will be assessed for rate of grid convergence. Here, c_f is the local skin friction, defined by $c_f = 2\tau_w / (\rho_\infty U_\infty^2)$.

Fig. 2.6 illustrates the wall normal profiles of u at $x/L = 0.97$, the station of interest, on the sequence of grids. We note that the results converge to typical turbulent boundary layer. Refinement of the grid beyond G3, has very little change in the eventual profile observed. In the outer units, where wall distance is scaled by local δ , u reaches its freestream value by $y/\delta = 1.0$. In the inner or wall scaling, all the profiles have linear profile near the wall $y^+ \lesssim 5$, and follow the log-law beyond $y^+ \approx 80$. The log law dictates that the velocity




Figure 2.6: Grid convergence history of wall-normal profiles of u at the station of interest $x/L = 0.97$




Figure 2.7: Grid convergence history of wall-normal profiles of μ_t at the station of interest $x/L = 0.97$ in outer and inner scaling.




Figure 2.8: Grid convergence history of c_f v/s streamwise x/L

profile follows,

$$u^+ = \frac{1}{\kappa} \ln(y^+) + B. \quad (2.21)$$

The values of the constants typically, are $\kappa = 0.41$ and $B = 5.2$ (cf. [Pope, 2000](#)). Eventually, the profiles saturate to the freestream values. In plus units, c_f for incompressible flow can be rewritten as $c_f = 2/U_\infty^2$ at each streamwise station. The values differ for each grid, but start converging G3 onwards. The convergence history of μ_t follows a similar trend where G1 and G2 values differ the most from the rest of the grids and we start observing grid convergence from G3 onwards. Beyond G3, perceptible differences are not seen in the plots and results appear grid converged.

In [fig. 2.8](#) the convergence history of c_f against the streamwise direction is illustrated. The large variations amongst different grids in their c_f prediction near the leading edge of the plate reflect the very fine grid requirements in both x and y at that location. The values of c_f start converging G3 onwards, congruent to previous observations. The difference in the solution on the sequence of grids G1-G5, possibly comprises of errors accumulated at the local station and the errors which flow downstream from the leading edge region where the grid may have been slightly coarser.

In [table 2.2](#), the changes in c_f and δ saturate as we move to finer grids, indicating we are converging to grid independent values. The rate of convergence follows an expected second order rate w.r.t. a grid size, as compiled in [fig. 2.9](#). The grid size is $h = 1/\sqrt{n_x n_y}$, and the error for each grid is computed w.r.t. to the finest grid data. For c_f , the value obtained by the NASA codes FUN3D and CFL3D with quadrilateral meshes (recall this is a




Figure 2.9: The rate of convergence for c_f and δ in comparison to the theoretical prediction of second order accuracy. The mesh size is $h = 1/\sqrt{n_x n_y}$ and the error in each grid is computed w.r.t the results on the finest grid.

Grid name	$c_f \times 10^3$	δ/L
G1	2.6953	0.015514
G2	2.7175	0.014184
G3	2.7231	0.014222
G4	2.7245	0.014110
G5	2.7250	0.014078

Table 2.2: Convergence history of c_f and δ at $x/L = 0.97$ for the sequence of 5 grids used.

2-D simulation) on the finest grid G5 is $\approx 2.705 \times 10^{-3}$. The value *tortuga* converges to is off by around 0.5%. Given the code shows grid convergence, matches theoretical prediction for wall-bounded flows (cf. fig. 2.6), we consider the code to be verified and validated to perform SA-RANS simulations for wall-bounded flows.

Chapter 3: Application to a shock/boundary layer interaction

We apply the MFSA to a shock/boundary layer interaction (SBLI) problem. Oblique shocks (shocks that are not orthogonal to the flow) appear in supersonic flow when it has to turn or suddenly change directions. SBLI occurs when such a shock wave interacts with the boundary layer on a nearby solid surface. This phenomenon can cause significant changes in the behavior of the boundary layer, including separation, reattachment, and turbulence amplification (cf. [Clemens & Narayanaswamy, 2014](#); [Gaitonde, 2015](#); [Gaitonde & Adler, 2023](#)). Examples of SBLI in engineering application include scramjet (supersonic air-breathing propulsion) inlets, reentry vehicles, supersonic and hypersonic flight.

Figure [3.1](#) illustrates a canonical inlet for a supersonic propulsion air-breathing engine. The isolator of this inlet typically experiences SBLI and this affects the performance of the engine. Figure [3.2](#) sketches a typical SBLI flow configuration. The shocks in the freestream cause large pressure jumps which are seen by the boundary layer as an adverse pressure gradient. This leads to flow separating from the wall. Nearby turbulence provides the separated flow momentum which helps it eventually reattach to the wall. Due to this, a recirculating region is formed, which is marked as the separation bubble in fig. [3.2](#). The




Figure 3.1: Schematic illustration of an inlet of a scramjet engine. The isolator region experiences SBLI which is the application case for MFSA (inspired by [Khan & Hasan, 2021](#))




Figure 3.2: Schematic representation of an SBLI flow configuration. The boundary layer is coming from left to right, experiences an oblique shock. This causes a separation bubble to form, where there is recirculating flow. The flow eventually reattaches. The wall again turns the flow causing a reflected shock.

incident shock is caused by some external geometry turning the flow in the freestream. The wall eventually turns the flow back to being parallel to the wall, causing the reflected shock. This is well illustrated from the freestream streamline in fig. 3.3. Upon reattachment, the boundary layer height and turbulence within it experience amplification (cf. Di Renzo *et al.*, 2022). In an SBLI configuration, the coordinates are scaled by the boundary layer height at some reference location in the incoming boundary layer (δ_r) and the inviscid impingement location on the wall (x_{imp}) is treated as the origin.

3.1 Problem set-up

The flow setup comprises of multiple flow regimes - wall bounded flow, shock-induced separation, and flow reattachment approaching wall equilibrium. Traditionally all RANS models have been calibrated on wall bounded flow, free shear flow and homogeneous isotropic flows. Presence of separated flow regime is known to detrimentally affect RANS models accuracy, so it would be a good test case to gauge how well the method is able to constrain the RANS model to some base LES data. Also, we can assess how well the method performs in multiple flow regimes just by this one test case.

The nominal baseline condition is a shock with an impingement angle $\beta = 35.00^\circ$ to an incoming boundary layer of $Ma_e = 2.28$ and $Re_{\delta_2} = \rho_e U_e \theta / \mu_w \approx 1000$, coming from left to right in the x -direction. The length scale for the Reynolds number is the momentum thickness θ (cf. Schlichting & Gersten, 2016), which is defined as

$$\theta = \int_0^\infty \frac{\rho u_1}{\rho_e U_e} \left(1 - \frac{u_1}{U_e}\right) dy, \quad (3.1)$$




Figure 3.3: Contours of ρ obtained from an LES at nominal condition. Yellow to orange shows an increase in ρ . The black lines indicate streamlines. The streamline starting from the wall and ending near the impingement location is the separation bubble. The top most streamlines indicate the turning of the freestream flow due to shocks. The streamline below it starts in the incoming boundary layer and indicates that the boundary layer height amplifies post-shock interaction.

with U_e being the local freestream streamwise velocity value at any station. The wall enforces no slip and isothermal temperature $T_w = 0.5T_r$. Here, the subscript e denotes the freestream/boundary layer edge condition and w denotes the wall condition. The recovery temperature $T_r = T_e(1 + 0.5\text{Pr}^{1/3}(\gamma - 1)Ma_e^2)$ denotes the isothermal temperature at which the wall behaves as an adiabatic wall. In high-speed flows, due to viscous heating, the boundary layer heats up in comparison to the freestream and conducts heat from the wall if the wall is not at recovery temperature. At the recovery temperature the wall-normal gradients vanish at the wall, this can be thought as a local maxima location along the wall-normal direction for canonical cases. If a wall is at $T_w < T_r$, it is called a cold wall and it takes away heat from the boundary layer, at recovery temperature wall becomes adiabatic, and for $T_w > T_r$, it is called a hot wall, which conducts heat to the fluid. This effect is dependent on the Mach number Ma_e , so for low-speed flows the recovery temperature is close to the freestream value, but increases as the speed goes up. For this work, we use ideal gas laws for the equation of state with $\gamma = 1.4$, Newtonian fluid with Fourier heat flux, power law




Figure 3.4: Schematic plot illustrating the domain used and choice of boundary conditions for LES and RANS simulations. The LES domain is three-dimensional while RANS domain (highlighted in gray) is a subset of LES domain and excludes the spanwise direction.

for viscosity with exponent being 0.75 and $Pr = 0.7$.

3.1.1 LES setup

The hybrid code is used to perform the LES calculations. Fig. 3.4 displays the domain with $L_1 \times L_2 \times L_3 = 100\delta_0 \times 40\delta_0 \times 8\delta_0$. Here, δ_0 is the height of incoming boundary layer, and subscripts 1, 2 and 3 denote streamwise, wall-normal and the spanwise directions, respectively. The grid is uniform in the streamwise and spanwise directions with the grid resolution $dx_1^+ \approx 18$ and $dx_3^+ \approx 9$ (in wall units). The first point above the wall is at $x_{2,w}^+ \approx 0.8$ (in wall units). Here, the wall units are derived from the incoming boundary layer. The grid stretches in x_2 as we move away from the wall. These grids were found to be sufficient for grid convergence (cf. Larsson *et al.*, 2022).

A boundary layer velocity profile is specified at the inflow and digital filtering is used to generate turbulence. The Walz's relation (cf. Walz, 1969) is used for specifying the

temperature profile. For outflow, a sponge layer is deployed. At the wall, Dirichlet conditions are specified for velocity (no slip) and temperature (isothermal). At the top and upper parts of the inflow profile, pre- and post-shock state values are specified for velocity, temperature and pressure. The shock strength depends on the incoming Mach number and β , while the profile is determined by β and shock impingement location.

The LES data is averaged in statistically homogeneous directions (x_3 and t) from samples obtained in a simulation. Ensemble average by definition requires infinite samples which is not computationally feasible, so we will always have a statistical error, proportional to $\sqrt{N_s}$, where N_s is the number of samples we have. This error ε_J for each QoI J is estimated using the batch method ([Arias-Ramírez & Larsson, 2022](#)).

3.1.2 RANS setup: standalone, frozen and LES-constrained

For the RANS there are three cases - standalone, frozen and LES-constrained. In all the cases, to solve the equations, the finite-volume code `tortuga` is used. Standalone RANS refers to vanilla SA-RANS equations (eqns. (2.3) and (2.9)), without any input from the LES simulation. The LES-constrained solves the SA-RANS equations with the additional forcing term computed using LES inferred data (cf. eqn. (2.8)). To illustrate the need for the forcing function in the LES-constrained equations, an additional set of frozen RANS simulations are run at nominal conditions. The frozen RANS simulation corresponds to the case where the inferred Π^{LES} is specified to the RANS solver and the conservation laws are solved without

solving for any auxiliary equation, i.e., solving,

$$M\left(\bar{\mathbf{v}}, \Pi^{\text{LES}}; \mathbf{a}_b\right) = 0. \quad (3.2)$$

All the RANS simulations (standalone, frozen and LES-constrained) are performed on a subdomain of the LES geometry (excluding the spanwise direction). The sub-domain is highlighted in gray in fig. 3.4. This is primarily done to avoid regions of LES where the flow-field maybe unphysical due to the boundary conditions. At the inflow, some part of the streamwise length is used to get a fully-developed turbulent boundary layer (cf. [Dhamankar et al., 2018](#); [Wu, 2017](#)), while at the outflow the flow-field is affected by the presence of sponge layers to reduce the acoustic waves (cf. [Bodony, 2006](#)). For the RANS inflow, the LES solution at the nominal condition and the respective inferred $\tilde{\mathbf{v}}$ is sampled. The inflow location is determined such that the incoming turbulent boundary layer is fully-developed. The reference values for boundary layer height (δ_r) and free stream values are drawn from the inflow location. At the RANS outflow, the primitive variables satisfy the characteristic boundary conditions. The pre- and post-shock states for the primitive variables p, u_i, T on the top boundary are specified using the incoming Mach number, β and shock impingement location. The transition from pre- to post-shock state is done using a hyperbolic tangent function. The wall boundary conditions are the same as that of the LES. An important feature to note here is that we apply the same boundary conditions for standalone, frozen and LES-constrained RANS, with additional forcing terms for the last case (cf. eqn. (2.8)).

In terms of grid resolution, the standalone RANS showed grid-converged results even with a grid that was twice as coarse as compared to the LES grid. This is a direct consequence

of the SA model relaxing near-wall requirements compared to other RANS models. However, for LES-constrained RANS, it was necessary to use the same LES grid to reconstruct the baseline flow-field. This is discussed in sec. 3.2.3. For frozen RANS, the same LES grid is used.

Lastly, the modifications to the inference of the auxiliary variables particular to the SBLI LES base case are discussed. Around the intersection of the incident and reflected shocks, just above the separation bubble, the inferred $\nu_t^{\text{LES}}, \lambda_t^{\text{LES}} < 0$. We enforce continuity of these inferred quantities, by raising ν_t^{LES} and λ_t^{LES} to their freestream values. This also helps in avoiding dividing by 0 when computing Pr_t^{LES} . Further, ν_t^{LES} and λ_t^{LES} decay at different rates towards the freestream, resulting in orders of magnitude of variation in the corresponding Pr_t^{LES} . We subsequently constrain Pr_t^{LES} between 0.1 and 2.5, since it is expected to be an $\mathcal{O}(1)$ quantity.

3.2 Results: nominal condition

In this section the outputs of different fidelities (LES and RANS) are analyzed for the same baseline condition. In addition, to motivate the need for a forcing term in the MFSA framework, LES is also compared to frozen RANS (both with variable Pr_t inferred from LES and uniform $\text{Pr}_t = 0.9$).

From an engineering standpoint, arguably the most important quantities are skin

friction coefficient (c_f) and the Stanton number (c_h), which are defined as,

$$c_f = \frac{2\tau_w}{\rho_{e,\text{ref}}U_{e,\text{ref}}^2}, \quad \tau_w = \mu_w \left. \frac{\partial u_1}{\partial x_2} \right|_w, \quad (3.3a)$$

$$c_h = \frac{q_w}{\rho_{e,\text{ref}}c_p(T_w - T_r)}, \quad q_w = \frac{c_p\mu_w}{\text{Pr}} \left. \frac{\partial T}{\partial x_2} \right|_w. \quad (3.3b)$$

Here the subscript ref indicates the reference values, which are sampled from LES data at the RANS inflow location and subscript e denotes the freestream values. In fig. 3.11, both the standalone RANS and frozen RANS agree with the LES predictions of c_f and c_h in the incoming boundary layer region. However, all the RANS-based simulations diverge from the LES predictions near the separation bubble (marked by $c_f \leq 0$). Standalone RANS shows a similar qualitative behavior to that of LES, but differs in the location of separation and reattachment locations ($c_f = 0$), and the values of peak negative friction and peak heat transfer inside the bubble. Post reattachment, these errors propagate downstream where eventually standalone RANS and LES converge. Frozen RANS on the other hand experiences large jumps in both c_f and c_h just before separation, which is not seen in either LES or standalone RANS. In addition, it also predicts a very shallow bubble compared to LES, but agrees well with the reattachment location. The c_f curve of the frozen RANS simulations agrees well with the LES data post reattachment. On the other hand, the c_h curve shows large discrepancy with LES data even in the post reattachment region. The choice of Pr_t seems to have very little influence over the frozen RANS predictions. It does alter the values of c_f and c_h , especially in the bubble and post reattachment, but does not bring any qualitative change in profile of these wall quantities.




Figure 3.5: c_f and c_h curves highlighting the discrepancy among different fidelities at baseline condition. The standalone RANS (red line) shows systematic errors when compared to LES data (black line). Frozen RANS with LES inferred closures by itself also shows discrepancies when compared to LES results. Frozen RANS is examined for two cases - fixed $\text{Pr}_t = 0.9$ (green curve) and variable LES-inferred Pr_t (orange curve).

Next, in fig. 3.6 the contours of T in the domain are compared among different fidelities. The freestream region with shocks is captured equally well by each method. The incoming shock is imposed as a boundary condition and is similarly resolved in all the fidelities. The reflected shock is a byproduct of a solution and is strongly influenced by the separation location. Since the separation location is different for each case, the reflected shock traces different contour lines. In contrast to LES data, all the RANS simulations predict a larger T value in the interaction region (where the incident and reflected shock intersect). This carries over in the boundary layer post reattachment. Largest discrepancy is noted for the frozen RANS cases. The T values over the bubble and into the reattached boundary layer are 20% larger than the LES results. In addition, it overpredicts the peak T values in the incoming boundary layer as well. The variable Pr_t case shows no difference when contrasted with the uniform $Pr_t = 0.9$ results.

Subsequently, the streamlines of the flow-fields are analyzed. In fig. 3.7, first the streamlines in freestream and the outer part of the incoming boundary layer are considered. The freestream streamline indicates the flow turning from the incident and reflected shocks. The first downward turn is similar for all the fidelities, given it is an imposed boundary condition and influenced by inviscid behavior. The next turning, where the flow again becomes wall parallel, occurs at different locations. This is influenced by the separation location, the farther downstream the separation, the later the reflected shock turns the streamlines. This is in congruence with observations from fig. 3.5

The streamlines in the outer layer of the boundary layer ($y \lesssim 1.5\delta_r$) are affected by the separation bubble, which are portrayed in fig. 3.8. All the fidelities show agreement in the incoming boundary layer, but diverge as they pass over the separation bubble. Fig. 3.8




Figure 3.6: Temperature contours for different fidelities contrasting the different fidelities - LES, standalone RANS and frozen RANS with LES inferred closures.




Figure 3.7: The streamlines in the freestream ($y \gtrsim 2.75\delta_r$) and outer part of boundary layer ($y \lesssim 2\delta_r$) are contrasted between different fidelities - LES (black curve), standalone RANS (red curve) and frozen RANS (fixed $\text{Pr}_t = 0.9$ (green curve) and variable LES inferred Pr_t (dashed orange curve))

highlights the differences in the separation bubble height and size for each fidelity. The LES separation has a long laminar-region mouth, which none of the other fidelities capture. The standalone RANS predicts a larger bubble as compared to LES, whereas the frozen RANS results show a shallower bubble. The frozen RANS bubble size is shallower compared to the LES one. The streamlines for the frozen RANS case first turn towards the wall and then move over the bubble. This is in contrast to both LES and standalone RANS. Arguably this explains the peak in c_f just before separation.

Lastly, the wall normal profiles are compared at different fidelities. Since LES is the benchmark, streamwise locations are chosen using LES data. Locations in the incoming boundary layer, just before separation and in the post-shock attachment are chosen. The vertical lines in the ρ -contours in fig. 3.9 and fig. 3.10 indicate the stations chosen.

The streamwise velocity wall-normal profile in fig. 3.9 show that in the incoming boundary layer region, all the fidelities agree well, with minor differences in the buffer layer. At the station just before separation, the LES profile differs from all the fidelities, probably the LES boundary layer starts showing the tendency to separate, whereas the others still remain attached. In the post-shock state, similar observations are echoed. The standalone RANS has discrepancies with the LES profiles, however in this case, the frozen RANS agrees very well with the LES profile.

In fig. 3.10, a different trend is observed where all the fidelities disagree with LES. Standalone RANS has the minimum discrepancy with LES. Frozen RANS on the other hand overpredicts the peak values in the boundary layer at each station. The choice of Pr_t seems to have very little impact on the overall behavior of the frozen RANS simulations.

In conclusion, it is noted that the low-fidelity model predictions differ from the high-




Figure 3.8: The streamlines near the separation region and the bubble profiles are compared for different fidelities - LES (black curve), standalone RANS (red curve) and frozen RANS with LES inferred closure (fixed $\text{Pr}_t = 0.9$ (green curve) and variable LES inferred Pr_t (orange curve)). The brown lines without any arrow indicates the bubble profile.




Figure 3.9: The wall-normal profiles of the streamwise velocity u_1 highlights the discrepancies for different fidelities at different stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer (second row), near separation (third row), post-shock reattachment (last row). The black line shows LES data, red is standalone RANS, green is frozen RANS with fixed Pr_t and dashed-orange lines are frozen RANS with variable Pr_t




Figure 3.10: The wall-normal profiles of temperature T highlights the discrepancies for different fidelities at different stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer (second row), near separation (third row), post-shock reattachment (last row). The black line shows LES data, red is standalone RANS, green is frozen RANS with fixed Pr_t and dashed-orange lines are frozen RANS with variable Pr_t

fidelity LES results at the baseline condition. This supports the initial hypotheses of SBLI being a good test candidate to apply and assess MFSA on. The first step towards interaction between the two fidelities is to directly use auxiliary variables inferred from LES data and use that to solve the RANS conservation laws. Arguably, this choice performs worse than standalone RANS at places, especially for, wall quantities c_f, c_h, T values across the domain and bubble size. The overshoot in the frozen RANS c_f curves suggests that the flow experiences acceleration, just before separation. The downward acceleration just before separation seen in the streamlines of frozen RANS simulations also, which is not a typical LES behavior. However, it does a much better job in the reattachment region for the streamwise velocity. Probably this is the manifestation of imperfect closure modeling, which augurs the need for forcing term to constrain the RANS solver to mimic LES data at nominal conditions.

3.2.1 Influence of the forcing term

The effect of the forcing function in eqn. (2.8) is illustrated in fig. 3.11. and fig. 3.12. In fig. 3.11, both the standalone RANS and LES-constrained RANS agree with the LES predictions of c_f and c_h in the incoming boundary layer region. However, while the standalone RANS diverges from the LES predictions near the separation bubble (marked by $c_f \leq 0$), the LES-constrained RANS is perfectly able to replicate the LES data at this nominal condition. We note that this is a crucial advantage for MFSA, where LES-constrained RANS is able to reproduce LES baseline data, whereas standalone and frozen RANS cannot. Further, fig. 3.12 demonstrates the effective reconstruction of the separation bubble, another key feature of the SBLI flow configuration. As part of MFSA, LES-constrained RANS equations




Figure 3.11: c_f and c_h curves show the extent to which the LES-constrained RANS (blue line) replicates LES data (black dots) at the base condition, whereas the standalone RANS (red line) shows systematic errors when compared to LES.




Figure 3.12: The bubble profile, marked by the local bubble height (h_b) shows the extent to which LES-constrained RANS (blue line) replicates LES data (black dots) at the base condition, whereas the standalone RANS (red line) highlighting the systematic errors between LES and RANS.

are a RANS-like low-fidelity system that matches LES at some base condition, and now can be run at different input parameters to estimate the parametric sensitivity.

3.2.2 Distribution of the forcing term

The forcing term that is subtracted from standalone RANS model is the residual of the LES data on the standalone RANS system at the nominal conditions i.e., $\mathbf{f} = M \left(\overline{\mathbf{v}}^{\text{LES}}, \Pi^{\text{LES}}; \mathbf{a}_b \right)$. The LES data exactly satisfies the mean conservation laws (cf. eqn. (2.3)). Standalone RANS using auxiliary variables Π models the extra correlation terms (cf. eqn. (2.4)), which will be never be perfect and the misfit in these models and its non-linear interaction with other terms is seen in the equation is what is seen in the residual of LES data.

Since the conservation of mass in eqn. (2.3) has no additional correlation term the residual of LES data should be 0. However, given the numerical and statistical error in LES

data and different numerical schemes used by LES and RANS solver this may not be the case, exactly. For the conservation of momentum in eqn. (2.3), the residual of LES data, governed by the misfit in the model is,

$$f_{\rho \mathbf{u}} = -\frac{\delta_r}{\rho_e U_e} \left[\frac{\partial}{\partial x_j} \left(\bar{\rho}^{\text{LES}} \left(2\nu_t^{\text{LES}} \tilde{S}_{ij}^{\text{LES}} + R_{ij}^{d \text{LES}} \right) \right) \right]. \quad (3.1)$$

For the conservation of energy equation the residual reads

$$\begin{aligned} f_{\rho e_0} = & -\frac{\delta_r}{\rho_e U_e c_p T_e} \left[\frac{\partial}{\partial x_j} \left(\bar{\rho}^{\text{LES}} \left(\frac{\nu_t^{\text{LES}} c_p}{\text{Pr}_t^{\text{LES}}} \frac{\partial \tilde{T}^{\text{LES}}}{\partial x_j} + h_{Tj}^{\text{LES}} \right) + \bar{\rho}^{\text{LES}} \left(2\nu_t^{\text{LES}} \tilde{S}_{ij}^{d \text{LES}} + R_{ij}^{d \text{LES}} \right) \tilde{u}_k^{\text{LES}} \right) \right. \\ & \left. + \frac{\partial}{\partial x_j} \left(-\overline{\sigma_{jk} u_k''}^{\text{LES}} + \frac{1}{2} \bar{\rho}^{\text{LES}} \overline{u_j'' u_k'' u_k''}^{\text{LES}} \right) \right]. \end{aligned} \quad (3.2)$$

The LES data satisfies all the other terms of the standalone RANS equation exactly and what is left is the misfit in the model for the correlation terms. In the chosen standalone RANS model few correlation terms are not modeled, and hence they only appear as part of the LES residual at the baseline conditions. In order to compute these forcing terms the inferred fields for ν_t and Pr_t from the LES data at baseline conditions are used. Hence, the residual for the auxiliary variables is also computed. This is done by computing the left-hand side in the eqn. (2.9), using $\tilde{\mathbf{u}}^{\text{LES}}$ and $\tilde{\nu}^{\text{LES}}$, where the latter term is related to ν_t^{LES} through eqn. 2.10.

In fig. 3.13, the absolute value of the forcing function $\mathbf{f}$ (residual of baseline LES data) are visualized for each component of the mean conservation laws. We only look for forcing terms (LES residual) larger than $O(10^{-4})$ to isolate the part of the residual larger than the




Figure 3.13: Distribution of the forcing function used in the proposed multi-fidelity sensitivity algorithm framework. The forcing function is the residual of the LES data on the standalone RANS equation. The absolute value for each component is contrasted in log scale.

statistical noise in the LES signal. For the conservation of mass, the residual of LES data is concentrated over the top of the separation bubble, a key flow feature having different predictions from LES and standalone RANS. This mismatch is highlighted in the LES residual on the standalone RANS equations and also persists in the residual of the conservation of momentum and energy equations. In the pre-shock (upstream of the separation bubble) and post-shock (downstream of the bubble) the presence of modeling errors in the residual of conservation of momentum and energy equation are highlighted, especially close to the wall. These modeling errors are absent in the residuals in the conservation of mass equations, which has no correlation terms to be modeled, as discussed previously. It is interesting to note that larger modeling-based residuals present in the residual for the u_2 component of the conservation of momentum equations, highlighting the fact that the model arguably works well for the streamwise component equation, but is not able to capture the influence of Reynolds stress in the wall-normal component equation.

3.2.3 Grid assessment for MFSA

In fig. 3.14, the c_f curve of the standalone RANS and LES-constrained RANS runs are compared with baseline LES data on two grids. First being the LES grid itself, and the second grid being the “half” grid which is LES grid coarsened by a factor of two everywhere. The c_f which determines important flow features such as separation bubble is used as a metric to decide upon grid requirements. For standalone RANS, the half-grid seems to be good enough to provide grid independent results, albeit still being unable to match accuracy of LES results. For LES-constrained RANS simulations, the same underlying LES grid is needed




Figure 3.14: Comparison of standalone RANS (top) and LES-constrained RANS (bottom) with LES data overlaid as black dots on two grids - the LES grid and the LES grid coarsened by a factor of two everywhere, the “half” grid.

to reproduce the LES data itself.

3.3 Sensitivity estimation

In this section we assess how well standalone RANS and LES-constrained RANS (MFSA) can replicate the LES parametric sensitivities. In particular, we measure the sensitivity of QoIs J for perturbations in the parameter set $\mathbf{a}$. Here, J comprises of c_f , c_h and wall-normal profiles of u_1 and T at different streamwise stations, and $\mathbf{a}$ contains the wall temperature (T_w) and shock angle (β). Typically, J depends on $\bar{\mathbf{v}}(\mathbf{a})$ and $\mathbf{a}$. Eventually, we can compute the sensitivity with respect to each parameter $a_i \in \mathbf{a}$ as

$$\frac{dJ}{da_i} = \frac{\partial J}{\partial a_i} + \frac{\partial J}{\partial \bar{\mathbf{v}}} \cdot \frac{d\bar{\mathbf{v}}}{da_i}. \quad (3.3)$$

We note that computing $d\bar{\mathbf{v}}/da_i$ is the most computationally intensive part of the process. The simplest way to compute this term is by running multiple simulations at different a_i values, and then using a finite-difference method in the parametric space. This is discussed in detail for both LES and RANS (standalone and LES-constrained) simulations below.

3.3.1 LES perturbation setup

Since LES is quite expensive, the parametric sensitivity is computed only for one perturbation in each direction, using finite-differences. Here, the perturbation sizes in $\mathbf{a}$ are 10% and 0.6% in T_w and β , respectively. The corresponding sensitivities are considered as benchmark for the subsequent RANS calculation. These perturbations are imposed via the boundary conditions on the LES flow domain.

Applying finite-differences to the QoIs obtained from LES involves subtraction and then scaling by the difference in a_i (cf. eqn. (1.1)). This leads to accumulation of the statistical error, with the error for dJ/da_i being $\sqrt{2}\varepsilon_J/\Delta a_i$, with ε_J being the statistical error in QoI J . The statistical error is assumed to be the same at the perturbed and the base state, since there are same number of snapshots to compute the ensemble average from each LES. Multiplying the statistical error with a factor t_α , the confidence interval for any QoI or its derivative can be written as $J \pm t_\alpha \varepsilon_J$ and $dJ/da \pm \sqrt{2}t_\alpha \varepsilon_J/\Delta a_i$, respectively (cf. [Arias-Ramírez & Larsson, 2022](#)). In this work, t_α is chosen to illustrate the 95% confidence interval.

3.3.2 RANS and MFSA perturbation setup

RANS (standalone and LES-constrained), being computationally cheap, allows for many more perturbed states to be run. For T_w , the perturbations are of the size 1%, −0.5%, 0.25%, −0.1%, 0.01%, −0.001% and for β , the perturbation sizes are 0.2%, −0.15%, 0.085%, −0.06%. The sensitivity to each parameter is found by fitting a curve to it (e.g. a parabola) and evaluating its derivative.

We similarly enforce the perturbations through the boundary conditions. However, for T_w , we additionally need to modify the inflow T profile according to changes in T_w . At the inflow the temperature profile is modified as

$$T_{\text{inflow}}^p = T_{\text{inflow}}^b + \left. \frac{\partial f}{\partial T_w} \right|_b (T_w^p - T_w^b), \quad (3.4)$$

where the function $f(T_w, u_1)$ is taken from [Zhang et al. \(2014\)](#) as

$$\begin{aligned} \frac{T}{T_e} = f(T_w, u) = & \frac{T_w}{T_e} + \left(\frac{T_r}{T_e} - \frac{T_w}{T_e} \right) \left(0.2 \left(\frac{u_1}{U_e} \right)^2 + 0.8 \frac{u_1}{U_e} \right) \\ & + \left(1.0 - \frac{T_r}{T_e} \right) \left(\frac{u_1}{U_e} \right)^2. \end{aligned} \quad (3.5)$$

Here, superscript p denotes a perturbed case and b denotes the base case sampled from LES. We note that eqn. (3.4) is used only for the inflow T profile. All the other quantities specified at the inflow are sampled from the LES base state.

3.3.3 Results: wall quantities

We begin the sensitivity assessment by considering the wall quantities c_f and c_h . The first row in fig. 3.15 and 3.16 illustrates the baseline LES profiles. The vertical lines indicate the separation ($c_f = 0$ and $u_1 dc_f/dx_1 < 0$) and reattachment ($c_f = 0$ and $u_1 dc_f/dx_1 > 0$) locations.

Next, the second and third row in fig. 3.15 and 3.16 show the sensitivities of the wall quantities. At the separation and reattachment points, both dc_f/dT_w and $dc_f/d\beta$ are < 0 . This indicates that the separation location moves upstream, and the reattachment moves downstream, i.e., that the size of the separation bubble increases as we increase T_w and β . This is consistent with previous work in the SBLI literature ([Bernardini et al., 2016](#); [Volpiani et al., 2018](#)).

In the incoming boundary layer, both MFSA and standalone RANS show systematic errors in predicting sensitivities w.r.t T_w . This is however not noticed for sensitivities w.r.t




Figure 3.15: Comparison of base LES sensitivities of c_f w.r.t. T_w and β against standalone RANS and MFSA. The dashed vertical lines mark separation and reattachment locations. The gray region marks the 95% confidence interval of the LES predictions.




Figure 3.16: Comparison of base LES sensitivities of c_h w.r.t. T_w and β against standalone RANS and MFSA. The dashed vertical lines mark separation and reattachment locations. The gray region marks the 95% confidence interval of the LES predictions.

β since the incoming boundary layer is unaffected by the impinging shock. Just before the separation, the standalone RANS and MFSA c_f and c_h sensitivity predictions diverge from the LES sensitivity, both in terms of magnitude and minima location. The mismatch is much smaller for the proposed MFSA compared to standalone RANS.

Inside the separation bubble, both standalone RANS and LES predict a positive overshoot of c_f which the proposed method fails to capture. However, standalone RANS and LES still disagree in the location and magnitude of this offshoot. This mismatch between MFSA and other predictions propagates downstream after reattachment, following which these profiles start agreeing with the LES sensitivities to within statistical errors.

For c_h sensitivity inside the separation bubble, both standalone RANS and MFSA show varying mismatches from the LES predictions. Just before reattachment, the sensitivity w.r.t T_w predicted by MFSA differs severely from the LES predictions. However, w.r.t β , the sensitivity prediction by MFSA is qualitatively similar to the LES predictions, albeit with a mismatch. In this region, standalone RANS has a contrastingly different behavior for the sensitivity w.r.t both β and T_w . In the post-shock region w.r.t T_w , all the predictions start agreeing again. However w.r.t β , we note that proposed MFSA performs remarkably well compared to standalone RANS, being as good as LES predictions to within the statistical error.

The MFSA c_f sensitivity predictions are agnostic to the choice of the Pr_t model. However, the c_h sensitivity predictions are affected by that choice. We note that there is no clear better choice.

3.3.4 Results: wall normal profiles

Next, we analyze the sensitivity of wall normal profiles of u_1 and T . In figs. 3.17 and 3.19, we show the baseline data in the first row followed by the corresponding sensitivities in the next two rows. The left and right columns show the results for the incoming boundary layer ($16\delta_r$ upstream of the impingement location) and the post reattachment region ($5\delta_r$ downstream of the impingement location), respectively.

Using MFSA we get a larger mismatch in the sensitivity of u_1 w.r.t. T_w in the buffer layer region, as compared to standalone RANS and LES. The performance however improves as we move towards the freestream where the MFSA sensitivity lies within the confidence interval of the LES prediction. Even the standalone RANS sensitivity prediction, which predicts an offshoot towards the freestream, lies at the edge of the confidence interval of the LES sensitivity prediction. The sensitivity of u_1 w.r.t. β is zero down to machine precision. All the predictions agree to this within statistical errors.

In the post-shock region, both the predicted sensitivities agree with the LES predictions near the wall, followed by deviation away from LES towards the freestream. Here, the MFSA diverges more as compared to standalone RANS. We note that the choice of Pr_t model does not significantly alter the MFSA predictions of the u_1 sensitivities.

For the sensitivity of the T profile w.r.t. T_w , we note that the LES sensitivity is well reproduced by all the methods, except for MFSA with uniform Pr_t in the post reattachment region. The sensitivities w.r.t β by all methods lie in the LES confidence interval. For T profiles, we note that the choice of Pr_t model has an influence on the eventual MFSA sensitivity prediction.




Figure 3.17: Comparison of sensitivities of u_1 profiles normal to wall in the incoming boundary layer, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.




Figure 3.18: Comparison of sensitivities of u_1 profiles normal to wall in the post-shock reattachment region, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.




Figure 3.19: Comparison of sensitivities of T profiles normal to the wall in the incoming boundary layer, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.




Figure 3.20: Comparison of sensitivities of T profiles normal to wall in the post-shock reattachment region, w.r.t T_w and β against standalone RANS and MFSA. The gray region marks the 95% confidence interval of the LES predictions.

3.3.5 Results: output pdfs

We finally demonstrate how the proposed MFSA enables UQ sampling. To this end, we assume that T_w follows a normal distribution with a mean of $0.5T_r$ and standard deviation of $0.005T_r$ (see first row in fig. 3.21). We focus on scalar QoIs such as the separation location (x_s), the reattachment location (x_r), and spatially averaged wall quantities (c_f and c_h) defined as

$$\bar{C}_{f,h} = \frac{\int_{x_a}^{x_b} c_{f,h} dx}{\int_{x_a}^{x_b} dx}, \quad x_{a,b} = x_s \pm 5\delta_r. \quad (3.6)$$

Obtaining a full output pdf is computationally prohibitive for LES, so we ran 3 additional cases at $T_w = 0.52T_r, 0.45T_r, 0.47T_r$, in addition to $0.5T_r$ and $0.55T_r$ already used to compute LES sensitivities, followed by fitting parabolic response surfaces for the QoIs. These response surfaces were then sampled to produce the benchmark LES pdfs. This process was repeated for the RANS cases (standalone and MFSA) to compute response surfaces that were then sampled.

The resulting pdfs are shown in fig. 3.21, which shows a clear systematic discrepancy between the standalone RANS and LES. In contrast, the MFSA pdf are perfectly centered around the LES ones, by construction. On the downside, the MFSA overpredicts the width of all pdfs, more so than does standalone RANS.

3.4 Concluding remarks

In this chapter, a computationally affordable and intrusive multi-fidelity method that enables sensitivity estimation and UQ exercises for LES is examined. The proposed method




Figure 3.21: The pdf of scalar QoIs are compared between LES, standalone RANS and MFSA.

is tested for a shock/boundary layer interaction problem and compared to both the “true” sensitivity estimated from LES and the sensitivity computed by a standalone RANS method. The proposed method does provide some qualitative improvements in the sensitivity predictions, especially around the separated flow region, but the standalone RANS offers better performance in fully developed boundary layers for momentum-related sensitivities. Methods to improve the accuracy of MFSA are needed to be explored. The work discussed here carries forward the discussion from [Arias-Ramírez \(2022\)](#), where an algebraic RANS closure model was used for MFSA, while in this work the SA-RANS model is used. SA-RANS model has a transport equation for one of its auxiliary variables and is known to be more reliable (cf. [Wilcox, 1998](#)). However, extending the algorithm to a better closure does not improve the MFSA results. This suggests that the path to improvement of accuracy may not just be higher order closure but maybe recalibration of RANS models for the constrained use within the MFSA. Alternative idea could be to explore other inference framework which links the LES data to RANS model.

Overall, the study serves as a proof-of-concept that the proposed multi-fidelity sensitivity algorithm can be a useful addition to the turbulence simulation toolbox, and also that there is ample room for modeling improvements in future work.

Chapter 4: Inference of turbulent diffusivities from LES

For the proposed method in the thesis, the RANS auxiliary variables at the base LES state, Π^{LES} , are needed. The auxiliary variables are used to find the turbulent diffusivities ν_t and λ_t , which in turn help to model the turbulent fluxes like Reynolds stresses R_{ij} and heat flux q_{Tj} respectively, as per the Boussinesq approximation in eqn. (2.5) and (2.6). The method to first infer ν_t^{LES} and λ_t^{LES} from LES data and then backtrack Π^{LES} is briefly discussed in sec. 2.4.1. In this chapter ideas pertaining to inference of Π^{LES} and its influence on the accuracy of the sensitivity are explored in further detail. For the purposes of inference, all the data is sampled from LES baseline results.

4.1 Inference cost functions

The main component of the inference process is to first find ν_t^{LES} and λ_t^{LES} . Finding Π^{LES} after that is relatively straight forward, usually an algebraic equation is only needed to be inverted, based on the choice of the RANS model. For this thesis, SA-RANS model is used, which has the auxiliary variables are $\tilde{\nu}$ and Pr_t . Equation (2.10) and $\text{Pr}_t = \nu_t/\lambda_t$ link these auxiliary variables to the turbulent diffusivities. The main effort is concentrated on

finding the turbulent diffusivities as they directly model the LES data.

In addition, the constrained RANS equations in eqn. (2.8), where the residual of the LES data is subtracted from the standard set of equations, always ensures that for any value of Π^{LES} , ergo any value of ν_t^{LES} and λ_t^{LES} , the base case can be reconstructed from the forced RANS system. Since this is true for any inferred value of turbulent diffusivities, inferring these quantities could possibly have an impact on eventual accuracies of sensitivity predictions.

With that aim, alternative ideas are explored in this section. In section 2.4.1, error tensor $E_{ij}(\nu_t) = R_{ij}^d + 2\nu_t S_{ij}^d$, and error vector $e_k(\lambda_t) = q_{T_k} + \lambda_t g_k$ were defined. The most apparent choice is to minimize some norm of these quantities. For this, the cost functions

$$\mathcal{J}_{M1}^{\nu_t} = \int_{\Omega} \frac{E_{ij}E_{ij}}{U_R^4} d\Omega, \text{ and } \mathcal{J}_{M1}^{\lambda_t} = \int_{\Omega} \frac{e_k e_k}{U_R^2 T_R^2} d\Omega, \quad (4.1)$$

look at the integrated non-dimensional norm over the domain of interest. The subscript “R” is used to denote reference values of velocity (U_R) and temperature (T_R). These are just scalar numbers, at the reference location. For this work, it is the free stream value at the RANS inlet as in fig. 3.4 and sec. 3.1.2. Minimizing the cost functions in eqn. (4.1) gives the ν_t and λ_t fields that have the least modeling error.

Upon looking at the averaged conservation laws in eqn. (2.3), the terms to be modeled (R_{ij} and q_{T_j}) appear in form of their spatial divergence. This motivates the alternative cost functions

$$\mathcal{J}_{M2}^{\nu_t} = \int_{\Omega} \frac{\delta_R^2}{U_R^4} \frac{\partial (\bar{\rho} E_{ik})}{\partial x_k} \frac{\partial (\bar{\rho} E_{ij})}{\partial x_j} d\Omega, \text{ and } \mathcal{J}_{M2}^{\lambda_t} = \int_{\Omega} \frac{\delta_R^2}{U_R^2 T_R^2} \frac{\partial (c_p \bar{\rho} e_k)}{\partial x_k} \frac{\partial (c_p \bar{\rho} e_j)}{\partial x_j} d\Omega. \quad (4.2)$$

The ν_t and λ_t fields minimizing this cost ensures that the LES data has minimal discrepancy with RANS equations, as far as these terms are concerned. Here, δ_R is used as a length scale to account for length dimensions coming from taking a spatial derivative.

The two forms of cost functions in eqn. (4.1) and (4.2) are referred to as “Method 1” and “Method 2”, represented by M1 and M2, respectively.

4.1.1 Rescaling of terms

The terms used to infer ν_t and λ_t under the M2 paradigm are scaled by density ρ and specific heat c_p , as required by the conservation laws for momentum and energy, respectively. Hence a shift of variables from $\nu_t \rightarrow \mu_t = \bar{\rho}\nu_t$ and $\lambda_t \rightarrow \kappa_t = c_p\bar{\rho}\lambda_t$ is done to make the inference process convenient. The error tensor E_{ij} is now redefined as $E_{ij}(\mu_t) := \bar{\rho}R_{ij}^d + 2\mu_t S_{ij}^d$ and error vector e_k is rescaled to $e_k = h_{T_k} + \kappa_t g_k$, $h_{T_k} = c_p\bar{\rho}q_{T_j} = c_p\bar{\rho}\widetilde{T''u_k''}$. This scaling of variables becomes useful in variable density compressible flows, which is our application case and if the fluid is not calorifically perfect. The M1 and M2 cost functions are also accordingly scaled.

4.1.2 Inference equations

The now scaled M1 cost function is

$$\mathcal{J}_{M1}^{\mu_t} = \int_{\Omega} \frac{E_{ij}E_{ij}}{\rho_R^2 U_R^4} d\Omega, \text{ and } \mathcal{J}_{M1}^{\kappa_t} = \int_{\Omega} \frac{e_k e_k}{(c_p \rho_R U_R T_R)^2} d\Omega. \quad (4.3)$$

Applying the Euler-Lagrange equation i.e., setting the variation $\delta \mathcal{J}_{M1}^{\mu_t}$ and $\delta \mathcal{J}_{M1}^{\kappa_t}$ to 0, gives the equation leading to optimal solution. The equations read

$$\begin{aligned} \frac{4}{\rho_R^2 U_R^4} \left(\bar{\rho} R_{ij}^d S_{ij}^d + 2\mu_t S_{pq}^d S_{pq}^d \right) &= 0, \text{ and} \\ \frac{2}{(c_p \rho_R U_R T_R)^2} (h_{T_k} g_k + \kappa_t g_j g_j) &= 0. \end{aligned} \quad (4.4)$$

The optimal μ_t field obtained from eqn. (4.4) exactly reproduces the the production term appearing in the transport equation of R_{ij} (cf. [Wilcox et al., 1998](#)). The eqn. (4.4) for M1 is a set of uncoupled algebraic equations that produces point wise estimates of μ_t and κ_t which minimizes the modeling error in the Boussinesq equations (cf. eqns. (2.5) and (2.6)). These fields are easy to obtain as a post processing step on the same grid the LES calculations were performed. The derivation of the Euler-Lagrange equations can be found in [Appendix A](#).

The M2 approach minimizes the alternative cost functions

$$\mathcal{J}_{M2}^{\mu_t} = \int_{\Omega} \frac{\delta_R^2}{\rho_R^2 U_R^4} \frac{\partial E_{ij}}{\partial x_j} \frac{\partial E_{ik}}{\partial x_k} d\Omega, \text{ and } \mathcal{J}_{M2}^{\kappa_t} = \int_{\Omega} \frac{\delta_R^2}{(c_p \rho_R U_R T_R)^2} \frac{\partial e_k}{\partial x_k} \frac{\partial e_j}{\partial x_j} d\Omega. \quad (4.5)$$

Applying the Euler-Lagrange equations to these cost functions i.e., setting the variation $\delta \mathcal{J}_{M2}^{\mu_t}$ and $\delta \mathcal{J}_{M2}^{\kappa_t}$ to 0, the equations are

$$\begin{aligned} \frac{\delta_R^2}{\rho_R^2 U_R^4} \left[-8 \frac{\partial}{\partial x_j} \left(S_{kj} S_{ki} \frac{\partial \mu_t}{\partial x_i} \right) - 8 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial S_{jk}}{\partial x_k} \right) S_{ij} \right) \mu_t - 4 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial (\bar{\rho} R_{jk})}{\partial x_k} \right) S_{ij} \right) \right] &= 0, \\ \text{and} \\ \frac{\delta_R^2}{(c_p \rho_R U_R T_R)^2} \left[-2 \frac{\partial}{\partial x_j} \left(g_j g_i \frac{\partial \kappa_t}{\partial x_i} \right) - 2 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial g_k}{\partial x_k} \right) g_i \right) \kappa_t - 2 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial h_k}{\partial x_k} \right) g_i \right) \right] &= 0. \end{aligned} \quad (4.6)$$

The R_{ij} and S_{ij} fields used in eqn. (4.6) are deviatoric (trace free). The superscript d denoting that is dropped for brevity. Detailed derivation of eqn. (4.6) can be found in Appendix A. In contrast to a set of algebraic equations in eqn. (4.4), eqn. (4.6) is a set of uncoupled elliptic PDEs. This makes computing μ_t and κ_t from M2 costlier as compared to M1 as now PDEs are needed to be solved. Here, M2 minimizes the error in the terms as they appear (spatial divergence of turbulent fluxes) in the RANS conservation laws as compared to the modeling error itself as is the case in M1. For statistically 1-D flow configuration such as a channel flow, the solution to eqn. (4.4) (M1) satisfies eqn. (4.6) (M2) as well. The boundary condition and numerical methods to solve them will be discussed in the forthcoming sections.

4.2 Augmenting the inference process

The inference process becomes ill-posed where $|S| = \sqrt{S_{ij}S_{ij}} \rightarrow 0$ and $|g| = \sqrt{g_i g_i} \rightarrow 0$. Under these conditions, there is no velocity or thermal gradient against which the turbulent fluxes can be modeled and the Boussinesq approximation (see eqn. (2.5) and eqn. (2.6)) become ill-defined. The inference process, both based on M1 and M2, needs to be regularized in order for the process to be robust in the entire domain. This is done by augmenting the cost functions in eqn. (4.3) and (4.5). Furthermore, in the presence of insufficient averaging, the mean temperature, velocity and turbulent fluxes will have statistical noise which eventually will lead to choppy inferred fields. It is worthwhile to note that by definition we need infinite number of samples to get exact mean fields, which is impractical, so there is always some amount of statistical noise in the averaged data. The

error is proportional to the square root of the number of ensembles available, as is the case in a typical Monte Carlo sampling exercise (cf. sec. 1.2). The inference cost functions are further augmented to ensure that the inferred fields are sufficiently smooth in spite of imperfect averaging.

4.2.1 Regularization term

The focus of this work is on wall-bounded flows. Situations where velocity and temperature gradients vanish typically arises in the freestream region of the flow. This motivates finding these regions, marking them with an indicator function χ and enforcing the freestream value at those locations. The indicator functions are defined based on the norms of S_{ij} and g_j as,

$$\chi_s = \begin{cases} 1 & |S|^2 < \epsilon_s U_R^2 / \delta_R^2 \\ 0 & |S|^2 \geq \epsilon_s U_R^2 / \delta_R^2 \end{cases}, \text{ and } \chi_g = \begin{cases} 1 & |g|^2 < \epsilon_g T_R^2 / \delta_R^2 \\ 0 & |g|^2 \geq \epsilon_g T_R^2 / \delta_R^2 \end{cases}. \quad (4.7)$$

The choice of the freestream values is based on the choice of the underlying RANS model. In the freestream, these turbulent parameters are usually set to zero. However, this work we deal with the SA model, which forces quantities like turbulent viscosities to be of the same order of magnitude as that of their molecular counterparts. Since, the eventual goal of the inference process is to find turbulent viscosities that will be used in the SA paradigm, the freestream values $\mu_t^\infty = 1.5\mu_R$ and $\kappa_t^\infty = 1.5c_p\mu_R/\text{Pr}$. Here the subscript R denotes the reference value used in the viscosity-temperature power law and Pr is the Prandtl number of the fluid.

The cost functions ensuring the desired behavior are

$$\mathcal{J}_r^{\mu_t} = \int_{\Omega} \frac{\chi_s}{(\rho_R U_R \delta_R)^2} (\mu_t - \mu_t^{\infty})^2 d\Omega, \text{ and } \mathcal{J}_r^{\kappa_t} = \int_{\Omega} \frac{\chi_g}{(c_p \rho_R U_R \delta_R)^2} (\kappa_t - \kappa_t^{\infty})^2 d\Omega. \quad (4.8)$$

The Euler-Lagrange equations of which are

$$\frac{2\chi_s}{(\rho_R U_R \delta_R)^2} (\mu_t - \mu_t^{\infty}) = 0, \text{ and } \frac{2\chi_g}{(c_p \rho_R U_R \delta_R)^2} (\kappa_t - \kappa_t^{\infty}) = 0, \quad (4.9)$$

respectively. The eqn. (4.9) are algebraic equations, similar to the standalone M1 equations (cf. eqn. (4.4)). They can be integrated into either (M1 or M2) inference process without any extra computational cost.

4.2.2 Smoothness constraint

To ensure smoothness, the gradient of the desired fields are required to be relatively small. Choppiness, arising from imperfect averaging, usually results in large unphysical gradients. The natural cost function from this standpoint is

$$\mathcal{J}_s^{\mu_t} = \int_{\Omega} \frac{1}{(\rho_R U_R)^2} \frac{\partial \mu_t}{\partial x_k} \frac{\partial \mu_t}{\partial x_k} d\Omega, \text{ and } \mathcal{J}_s^{\kappa_t} = \int_{\Omega} \frac{1}{(c_p \rho_R U_R)^2} \frac{\partial \kappa_t}{\partial x_j} \frac{\partial \kappa_t}{\partial x_j} d\Omega. \quad (4.10)$$

These cost functions yield the the Laplace's equation, which for each variables reads as,

$$\frac{-1}{(\rho_R U_R)^2} \frac{\partial^2 \mu_t}{\partial x_k \partial x_k} = 0, \text{ and } \frac{-1}{(c_p \rho_R U_R)^2} \frac{\partial^2 \kappa_t}{\partial x_k \partial x_k} = 0, \quad (4.11)$$

for μ_t and κ_t , respectively. Standalone M2 equations are a set of PDEs, so smoothness constraints can be added at no additional cost from a numerical solution standpoint. However for M1, the standalone equations are a pointwise estimate, so applying the smoothness constraint makes the inference process slightly expensive. The detailed derivation of the cost functions can be found in Appendix A.

4.3 Effective inference equations

The overall cost function to infer the model parameters φ (μ_t or κ_t) for any method M (M1 or M2) is,

$$\mathcal{J}^\varphi = \mathcal{J}_M^\varphi + \mathcal{J}_r^\varphi + \omega_s \mathcal{J}_s^\varphi . \quad (4.12)$$

Here, ω_s is the relative weight of the smoothing cost function, and ϵ (both for $|S|$ and $|g|$) is used to find the indicator maps χ . They are the hyperparameters for the inference process.

For M1, the resulting Euler-Lagrange equation is obtained by adding eqns. (4.4), (4.9) and (4.11), in similar spirit to eqn. (4.12). This stems from the fact that taking the variation of the cost function $\delta\mathcal{J}$ is a linear process. For φ , the structure of the regularized-and-smoothness-augmented M1 equation is

$$-a^\varphi \frac{\partial^2 \varphi}{\partial x_k \partial x_k} + B^\varphi \varphi + F^\varphi = 0 , \quad (4.13)$$

with

$$\begin{aligned}
a^{\mu_t} &= \frac{2\omega_s}{(\rho_R U_R)^2}, \quad a^{\kappa_t} = \frac{2\omega_s}{(c_p \rho_R U_R)^2}, \\
B^{\mu_t} &= \frac{8S_{pq}^d S_{pq}^d}{(\rho_R U_R^2)^2} + \frac{2\chi_s}{(\rho_R U_R \delta_R)^2}, \quad B^{\kappa_t} = \frac{2g_j g_j}{(\rho_R U_R T_R)^2} + \frac{2\chi_g}{(c_p \rho_R U_R \delta_R)^2}, \text{ and} \\
F^{\mu_t} &= \frac{4\bar{\rho} R_{ij}^d S_{ij}^d}{(\rho_R U_R^2)^2} - \frac{2\mu_t^\infty \chi_s}{(\rho_R U_R \delta_R)^2}, \quad F^{\kappa_t} = \frac{2h_{T_k} g_k}{(\rho_R U_R T_R)^2} - \frac{2\kappa_t^\infty \chi_g}{(c_p \rho_R U_R \delta_R)^2}.
\end{aligned} \tag{4.14}$$

For M2, the PDE structure is very similar, except for difference in the leading order term and takes the form

$$-\frac{\partial}{\partial x_j} \left(A_{jk}^\varphi \frac{\partial \varphi}{\partial x_k} \right) + B^\varphi \varphi + F^\varphi = 0, \tag{4.15}$$

with

$$\begin{aligned}
A_{jk}^{\mu_t} &= 8 \left(\frac{\delta_R}{\rho_R U_R^2} \right)^2 S_{ji} S_{ki} + \frac{2\omega_s}{(\rho_R U_R)^2} \delta_{jk}, \quad A_{jk}^{\kappa_t} = 2 \left(\frac{\delta_R}{c_p \rho_R U_R T_R} \right)^2 g_j g_k + \frac{2\omega_s}{(c_p \rho_R U_R)^2} \delta_{jk}, \\
B^{\mu_t} &= -8 \left(\frac{\delta_R}{\rho_R U_R^2} \right)^2 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial S_{jk}}{\partial x_k} \right) S_{ij} \right) + \frac{2\chi_s}{(\rho_R U_R \delta_R)^2}, \quad B^{\kappa_t} = -2 \left(\frac{\delta_R}{c_p \rho_R U_R T_R} \right)^2 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial g_k}{\partial x_k} \right) g_i \right) \\
&\quad + \frac{2\chi_g}{(c_p \rho_R U_R \delta_R)^2}, \text{ and} \\
F^{\mu_t} &= \frac{-4}{(\rho_R U_R^2)^2} \left(\frac{\partial}{\partial x_i} \left(\frac{\partial (\bar{\rho} R_{jk})}{\partial x_k} \right) S_{ij} \right) - \frac{2\mu_t^\infty \chi_s}{(\rho_R U_R \delta_R)^2}, \quad F^{\kappa_t} = \frac{-2}{(\rho_R U_R T_R)^2} \left(\frac{\partial}{\partial x_i} \left(\frac{\partial h_k}{\partial x_k} \right) g_i \right) \\
&\quad - \frac{2\kappa_t^\infty \chi_g}{(c_p \rho_R U_R \delta_R)^2}.
\end{aligned} \tag{4.16}$$

Both eqn. (4.13) and eqn. (4.15) have a very similar structure, both being linear and

elliptic PDEs similar to the Helmholtz equations. The main difference between the two is leading order diffusion term. For M1 i.e., eqn. (4.13), it is the Laplacian which means there is uniform isotropic diffusion. However for M2, there is anisotropic diffusion where the directions depend on diffusion matrices $S_{ji}S_{ki}$ for μ_t and g_jg_k for κ_t . For the latter, the diffusion matrix is degenerate i.e., is of rank 1 and diffusion only happens in the direction g_j . This becomes a challenging problem to track numerically (cf. Crouseilles *et al.*, 2015).

The equations are solved on the LES domain (cf. fig. 3.4) with homogeneous boundary conditions - homogeneous Dirichlet on the wall and homogeneous Neumann everywhere else. The RANS domain which is a subset of the LES domain samples the μ_t and κ_t values to compute the auxiliary variables Π . At the wall, since there is no slip boundary conditions and isothermal temperature conditions, there is no fluctuation in these variables and hence, no turbulent flux. This motivates the homogeneous Dirichlet boundary condition at the wall. All the other LES boundaries are far away from the RANS domain so the natural boundary condition suffices there. Near the LES inflow and outflow boundaries, the data is anyway not physically relevant, since inflow takes some streamwise distance to become an equilibrium boundary layer and there is a sponge layer near the outflow.

4.4 Numerical solution of the inference equations

Similar to the RANS solver, finite-volume based discretization is used with second order schemes to calculate derivatives. For any elliptic PDE, which is a solution to an optimization problem, finite-element based discretization is a more natural choice. However, finite-volume framework was chosen to utilize the already built underlying RANS framework.

The overall equation for both the methods M1 and M2 are very similar with the nature of diffusion being the only difference. The equations for M1 are a special case for the M2 equations from a numerical standpoint. For the sake of brevity, the discussion here focuses on the M2 equations. For the M2 equations, after discretization, the linear system reads $D^\varphi \bar{\varphi} = -\bar{F}^\varphi$, with $\bar{\varphi}$ and $\bar{F}^\varphi$ being finite dimensional vectors containing cell-averaged quantities and the linear operator

$$D^\varphi = \frac{\delta}{\delta x_j} \left(A_{jk}^\varphi \frac{\delta \bar{\varphi}}{\delta x_j} \right) + \bar{B}^\varphi,$$

acting on $\bar{\varphi}$. Here, δ represents the discretization of the derivative operator. In the classical cell-centered finite-volume approach, constructing the matrix form from the linear operator D_{num}^φ with anisotropic diffusion is not straight-forward (cf. sec. 2.5.3). This limits the linear systems methods that can be used. The Krylov subspace method do not need the explicit matrix form of the linear operator, rather the action of the linear operator suffices the purpose (cf. Saad, 2003). This however was not able to converge and the methods stalled after a few thousand iterations. To solve the system $D^\varphi \bar{\varphi} = -\bar{F}^\varphi$, we instead propose the iteration method

$$\begin{aligned} \tilde{D}^\varphi \bar{\varphi}^{n+1} + D^\varphi \bar{\varphi}^n + \bar{F}^\varphi - \tilde{D}^\varphi \bar{\varphi}^n &= 0 \\ \tilde{D}^\varphi (\Delta \bar{\varphi}^n) &= R_n^\varphi = - \left(D^\varphi \bar{\varphi}^n + \bar{F}^\varphi \right), \quad \Delta \bar{\varphi}^n = \bar{\varphi}^{n+1} - \bar{\varphi}^n \\ \Delta \bar{\varphi}^n &= - \left[\tilde{D}^\varphi \right]^{-1} R_n^\varphi, \quad \bar{\varphi}^{n+1} = \bar{\varphi}^n + \alpha^\varphi \Delta \bar{\varphi}^n \end{aligned} \tag{4.17}$$

If $\text{S.R.}(\tilde{D}^{\varphi^{-1}} D^\varphi) < 1$, where S.R. is the spectral radius of the operator, the proposed

iterations should converge (cf. [Saad, 2003](#)). The linear operator $\tilde{D}^\varphi$ should have isotropic diffusion so that its matrix form can be constructed. The natural choice becomes

$$\tilde{D}^\varphi = \frac{\delta}{\delta x_j} \left(A_{ii} \frac{\delta \overline{\varphi}}{\delta x_j} \right) + \overline{B}^\varphi, \quad (4.18)$$

where the isotropic diffusion is taken to be the trace of the anisotropic diffusion matrix. For M1 equations, the diffusion is isotropic so $\tilde{D}^\varphi$ and D^φ are the same and the solution can be obtained in one iteration. Algebraic multigrid method implemented in the library pyAMG is used to obtain $[\tilde{D}^\varphi]^{-1}$ (cf. [Bell et al., 2022](#)). The proposed iterations (cf. eqn. (4.17)) are called the outer iterations and the iterations of the algebraic multigrid method to obtain $[\tilde{D}^\varphi]^{-1}$, at each outer iteration are the inner iterations. For M2, it takes about 200-300 outer iterations depending on the underlying grid and some choice of hyperparameters for the solution to converge to machine precision.

4.4.1 Post processing

From a physical point of view, the turbulent inferred parameters must be positive i.e. $\mu_t, \kappa_t \geq 0$. This is not done as a part of inference process, but is rather done after the inferred fields are obtained. The fields are bounded away from zero by $\varphi = \max(\varphi, 0)$, where φ could be either μ_t or κ_t .

The objective of the inferring these fields is to eventually obtain effective “LES” values for the auxiliary variables Π for any RANS models (cf. sec. 2.4.1). The auxiliary variables have a theoretical near-wall behavior, $\tilde{\nu} \sim h$, $k \sim h$ and $\text{Pr}_t \sim O(1)$. This implies for $\mu_t/\overline{\rho} \sim h^3$ and $\kappa_t/(c_p \overline{\rho}) \sim h^3$, with h being the wall-normal height. To get these exact near-




Figure 4.1: Contours of μ_t near the separation bubble are compared for M1 and M2 inference methods. The black lines mark a streamline near the edge of the incoming boundary layer (top) and separation bubble (bottom), respectively.

wall behaviors, the underlying LES grid needs to match the DNS near-wall requirements, which is quite overbearing. Since this behavior not guaranteed from the underlying LES data, this is enforced after the variables Π are computed and proper near-wall-behavior is prescribed to those variables.

4.5 Comparison of inferred fields from methods M1 and M2

The inferred μ_t fields from the two methods are shown in Fig. 4.1. The regions where $\mu_t < 0$, are clipped to be zero (cf. sec. 4.4.1). Since this region of negative values occurs near the freestream, the value was raised from zero to their freestream values, to maintain continuity. Similar clipping is done for κ_t , and then Pr_t is computed. The Pr_t fields are further clipped to be between 0.1 and 2.5 for reasons discussed in sec. 3.1.2. This helps avoid “divide-by-zero” difficulties as well when computing Pr_t . Typically, most RANS models use $\text{Pr}_t = 0.9$ (cf. Wilcox *et al.*, 1998). Hence, the inferred Pr_t from LES data helps assess the




Figure 4.2: The first row shows the $\bar{\rho}$ contours from base case LES. Each succeeding row illustrates the μ_t and Pr_t obtained from the M1 and M2 inference methods at different locations marked by vertical dashed lines in the $\bar{\rho}$ contours. These stations represent the attached boundary layer region, near-separation region and the post-shock reattachment region.

validity of this assumption. From a RANS-modeling standpoint, these variables delineate the model efficacy the most. The results here are based on the choice of $\omega_s = 10^{-3}$ and $\epsilon = \epsilon_s = \epsilon_g = 10^{-3}$ for both M1 and M2. This choice is further discussed in sec. 4.6.

In the attached boundary layer region, (second row in fig. 4.2) both the methods produce near-identical profiles for μ_t and Pr_t . These numbers are much more important near the wall and in the buffer layer ($y/\delta_r \lesssim 0.6$). Note that the reference value for δ is taken from RANS inflow profile (cf. fig. 3.4). The predicted μ_t profiles start from zero at the wall and move super-linearly towards their peak value in the buffer layer, where most of the turbulent production happens, and eventually they die out to their freestream values. Typically, all RANS models set $\text{Pr}_t = 0.9$, both the methods here indicate that a non-uniform value is probably a better choice. Near the wall, they have a value close to the typical model value of 0.9, but the value keeps decaying as we move away from the wall. Near the free stream edge, the Pr_t is needed to be clipped. This happens as both κ_t and μ_t decay to the freestream values at different rates which leads to extreme values in the Pr_t (cf sec. 3.1.2). An equilibrium boundary layer is approximately a 1-D flow, hence both M1 and M2 agree on their μ_t and Pr_t predictions.

Next, we look at the region just before separation, where the flow has a 2D structure, given how the streamlines turn over the bubble. The μ_t profiles predicted by both the methods differ quite a lot, especially in the buffer layer region. The profiles start similarly at the wall but reach different values in the peak region. The peak value prediction for M2 is half of that for M1. The Pr_t profiles similarly do not agree with each other. The near-wall value predicted from M1 is $\approx 30\%$ more than that for the M2 prediction. Also, M1 required clipping (for the upper bound). M1 predicts near uniform value in the middle of

the boundary layer, whereas M2 has a slowly decaying nature. Towards the freestream these values do not hold any significant meaning.

In the post-shock reattachment region, both M1 and M2 values agree for the most part. The larger μ_t peak values as compared to the pre-shock region indicates turbulence amplification after the shock. For μ_t , both methods start similarly at the wall, but go towards their peak values at different rates in the buffer layer, where M2 reaches its peak much faster. The M1 profile has a lower peak compared to the M2 profile. Both the methods predict a contrasting near-wall behavior for Pr_t , where M2 suggests a value > 1 while M1 suggests < 1 . In the middle of the boundary layer, both the methods suggest similar behavior. The plateau region in the Pr_t profile for M2 indicates the fluid Pr value, suggesting that in this region both μ_t and κ_t were negative and raised to their freestream values.

Most of the differences between μ_t fields from M1 and M2 values can be described by the peak μ_t value at each streamwise location. This is sketched in fig. 4.1, accordingly. Fig. 4.1 suggests that in the incoming boundary layer, both M1 and M2 have a similar peak, with M2 experiencing a sharp dip compared to M1 just before separation. In the post-shock, region both M1 and M2 fields experience an amplification by the shock, which is slightly larger for M2 as compared to M1.

Overall, there is not much of a difference between the M1 and M2 predictions wherever there is attached near-equilibrium flow. The major difference arises just before separation location. This difference will also be useful to assess the dependency of sensitivity results on the inferred parameters.




Figure 4.3: Influence of the hyperparameter ϵ on the inferred μ_t profiles at a station in the incoming boundary layer (top row) and the post-shock reattachment (bottom row) regions.

4.6 Choice of hyperparameters and grid on the inferred profiles

In this section, the influence of the hyperparameters ϵ , which decides the indicator map and ω_s , which assigns the weight to the smoothness penalization term is examined. The discussion is similar for both the fields μ_t and κ_t across both the methods. For brevity, only plots for μ_t from M1 are shown at a station in the incoming boundary layer and the post-shock reattachment region. In fig. 4.3, we see that the choice of ϵ alters the approach




Figure 4.4: Influence of the hyperparameter ω_s on the inferred μ_t profiles at a station in the incoming boundary layer (top row) and the post-shock reattachment (bottom row) regions.

to the the freestream value of the inferred field. The larger the value of ϵ , the earlier the freestream values start from within the boundary layer. For smaller values of ϵ the freestream values start from way outside the boundary layer thickness, especially in the post-shock reattachment region. The value of $\epsilon = 10^{-3}$ transitioned to freestream values closest to the actual boundary layer height. The results shown here use $\omega_s = 10^{-3}$ as the other hyperparameter.

In fig. 4.4, the influence of the relative weight of the smoothness penalization is

illustrated with the other hyperparameter being $\epsilon = 10^{-3}$ for this discussion. For very low values of ω_s , the transition towards free stream values is very sharp. Larger values of ω also end up smearing the peak values of the inferred quantities inside buffer layer. The value of $\omega_s = 10^{-3}$ provides minimal smearing with smooth profiles and freestream transition. The value of ω_s also has an influence on the eventual convergence of the proposed outer iterations (cf. eqn. (4.17)), larger the value of ω_s faster the outer iterations converge. This becomes quite important for outer iterations related to κ_t which do not converge if ω_s is lower than a threshold.

Since we are inferring the fields μ_t and κ_t , or eventually Pr_t , by solving PDEs numerically, the grid resolution becomes another hyperparameter to investigate. The solution one obtains from numerically solving the PDE should be grid independent. The grid used for the LES calculation forms the basic grid and then we repeat the process on finer grids to assess grid convergence. Note that this is done only for the M2 method as it has a PDE for just the inference process itself (cf. eqn. (4.6)). Whereas, M1 becomes a PDE only upon adding the smoothness constraint, otherwise it is a simple pointwise calculation. For this comparison, the grids are only refined in the wall-normal direction, as the grid requirements for a RANS problem are coarser in the streamwise direction and require near-wall refinement, in comparison to LES grids. This is tabulated below.

The grid G1 is the LES grid itself. Grids G2, G3 are developed to analyze the grid dependency of the inference process. They are non-uniform grids in the wall-normal direction with a grid stretching factor of 1%. In the fig. 4.5, both G2 and G3 profiles agree with each other. These profiles disagree with the LES grid (G1) inferred profiles in terms of the peak value. This discrepancy is less for the post-shock reattached boundary layer as




Figure 4.5: Influence of the wall normal grid distribution on the inferred μ_t fields at a station in the incoming boundary layer (top row), and the post-shock reattached boundary layer.

Grid name	$\Delta y_w / \Delta y_{w,LES}$
G1	1
G2	0.3
G3	0.15

Table 4.1: Sequence of the 3 grids used for computation of M2 viscosity, describing the near-wall refinement as compared to the LES grid

compared to that in the incoming boundary layer. Possibly this could be due to the boundary layer amplification due to the shock relaxes the wall-normal grid requirements. The results discussed in sec. 4.5 used the results from grid G3.

4.7 Simulations using M1- and M2-inferred data

4.7.1 Results: frozen RANS simulations at base state

In this section, the performance of the M1- and M2-inferred fields with frozen RANS equations (cf. eqn. (3.2)) are contrasted with baseline LES results. The flow configuration is same as discussed in sec. ?? and sec. 3.1.2.

First, the streamlines of the flow-field are analyzed. In fig. 4.6, first the streamlines in freestream and the outer part of the incoming boundary layer are considered. In the freestream region, the flow turning by the incident shock is captured equally well by all the methods. Differences are however observed post the reflected shock interaction. This could be explained by the separation location, which is different for each case. In the outer part of the boundary layer, matching observations are made. The more interesting part being, when the LES streamlines start experiencing the effect of the bubble and turn upward, the

frozen RANS streamlines (for both cases) goes down, probably providing more momentum to be lost at the wall. Overall, the M2 predictions are more closer to the LES results. Here, only the uniform $Pr_t = 0.9$ case is considered, since it was previously observed that the choice of Pr_t has little impact on the overall results.

A congruent narrative is seen in fig. 4.7 with the visualization of the streamlines closer to the separation bubble. The M2 frozen RANS separation bubble is closer to that of LES data, whereas M1 frozen RANS produces a very shallow bubble. Both the frozen RANS methods miss out on the long upstream mouth present in the LES data. When the LES flow-field starts experiencing the presence of the bubble, the streamlines start turning upwards. However, for both frozen RANS, an opposite effect is seen in that region, with their streamlines turning downwards. Arguably, the M2 frozen RANS framework is closer to the LES data, but the qualitative behavior is similar for both choices in the inference framework. Probably, this is due to the limitations of the Boussinesq approximation, especially near the separation region.

The c_f and c_h curves with M2-inferred fields are closer to LES data as compared to M1 fields, as seen in fig. 4.8. However, the qualitative behavior for both the frozen RANS cases is very similar, both experiencing an increase of c_f just before separation in contrast to LES. Consequently, the separation location predictions by frozen RANS are more downstream as compared to LES, but all the methods predict a similar reattachment location. This overshoot of c_f could arguably be explained by the downward turning of streamlines in the frozen RANS cases (cf. fig. 4.7), which brings extra momentum towards the wall. A similar narrative is seen for c_h as well. Both the M1- and M2-based frozen RANS differ from the baseline LES data. They overpredict the heat transfer in the incoming boundary layer,




Figure 4.6: The streamlines in the freestream ($y \gtrsim 2.75\delta_r$) and the outer part of boundary layer ($y \lesssim 2\delta_r$) are compared for inference choices with LES (black curves). The M1-inferred frozen RANS streamlines are the red curves and M2-inferred frozen RANS streamlines are the blue curves. Here, both frozen RANS cases are with fixed $\text{Pr}_t = 0.9$. The brown line without any arrow indicates the bubble profile.




Figure 4.7: The streamlines near the separation region and the bubble profiles are compared for inference choices with LES (black curves). The M1-inferred frozen RANS streamlines are the red curves and M2-inferred frozen RANS streamlines are the blue curves. The brown line in each plot is the corresponding separation bubble. Here, both frozen RANS cases are with fixed $\text{Pr}_t = 0.9$.




Figure 4.8: Comparison of c_f and c_h profiles for frozen RANS with inferred fields from M1 and M2 methods with the baseline LES data.

and underpredict it in the post-shock reattached region. Arguably, upstream of the LES separation bubble, M2 is closer to LES data as compared to M1. For both c_h and c_f curves, the choice of Pr_t seems to have very little impact on the values of these wall quantities, as alluded before.

In fig. 4.9 and fig. 4.10, the wall-normal profiles of streamwise velocity u_1 and temperature T are examined. The velocity u_1 in the incoming boundary layer predicted by both M1 and M2 cases matches very well with the LES data. Just before the separation region, differences between frozen RANS and LES are noticeable, especially in the buffer




Figure 4.9: The wall-normal profiles of the streamwise velocity u_1 comparing the differences for each inference methods at separate stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer, near separation (second row: left to right), inside the bubble and post-shock reattachment (third row: left to right). The black line shows LES data. The red and magenta lines are M1 frozen RANS, with fixed Pr_t and variable Pr_t respectively. The blue and green lines are M2 frozen RANS, with fixed Pr_t and variable Pr_t respectively.

layer. Arguably, the extra streamwise momentum observed here is congruent with c_f overshoot observed before separation for frozen RANS. For these two regions, using either M1 or M2 (with uniform Pr_t and non-uniform Pr_t), results in no observable difference. Inside the bubble, visible differences between M1- and M2-inferred results are seen, with the latter being closer to LES profiles and does a better job of flow-field reconstruction. In the post-shock reattachment region, all the frozen RANS cases deviate from each other marginally. The overall frozen RANS prediction in this region is still qualitatively similar to LES results itself. Probably, the errors in this region are propagated from the separation bubble upstream.

Similar observations are made for the temperature profiles as far as comparison between frozen RANS with M1 and M2 is concerned. However, both M1 and M2 frozen RANS predictions never agree with the LES data. Frozen RANS, for both the cases, overpredicts the temperature in all the regions.

4.7.2 Results: MFSA sensitivity predictions

In the section, the sensitivity process for MFSA discussed in sec. 3.3 is repeated with the M2 inferred fields. The results are then compared with the LES and MFSA sensitivity with M1 inferred fields discussed in sec. 3.3. The process of constraining the RANS equations (cf. eqn. (2.8)), ensures that the LES data will be replicated at the specified condition, irrespective of the inferred values. Comparing the sensitivities obtained from these two different inference frameworks helps to gauge their influence on the eventual sensitivity prediction. In Sec. 4.5 highlighted differences between the M1 and M2 inferred Π variables.




Figure 4.10: The wall-normal profiles of the temperature T comparing the differences for each inference methods at separate stations along the streamwise direction marked as dashed lines in first row. These mark different regions which are - incoming boundary layer, near separation (second row: left to right), inside the bubble and post-shock reattachment (third row: left to right). The black line shows LES data. The red and magenta lines are M1 frozen RANS, with fixed Pr_t and variable Pr_t respectively. The blue and green lines are M2 frozen RANS, with fixed Pr_t and variable Pr_t respectively.




Figure 4.11: Comparison of base LES sensitivities of c_f w.r.t. T_w and β against MFSA prediction using M1- and M2-inferred variables. The dashed vertical lines mark separation and re-attachment locations. The gray region marks the 95% confidence interval of the LES predictions.




Figure 4.12: Comparison of base LES sensitivities of c_h w.r.t. T_w and β against MFSA prediction using M1- and M2-inferred variables. The dashed vertical lines mark separation and re-attachment locations. The gray region marks the 95% confidence interval of the LES predictions.

Recall, the differences between the two inference methods was dominantly observed just before the separation region (cf. fig. 4.1).

Since M2 directly minimizes the error in the terms as they appear in the conservation equations, it was expected to provide an improvement in the sensitivity prediction. Fig. 4.11 and 4.12 show a comparison between sensitivity prediction for c_f and c_h (cf. eqn. (3.3a) and (3.3b)) w.r.t T_w and β using the M1 and M2 inferred parameters. Here, both the choices of - (i) uniform $\text{Pr}_t = 0.9$ and, (ii) variable Pr_t as inferred from the LES data for both the methods are explored. In the incoming boundary layer, where differences between inferred parameters was small (cf. sec. 4.5), no differences are seen in the sensitivity prediction from both the methods and both the QoI, c_f and c_h . A similar trend is noticed in the post-shock region as well, where the boundary layers starts to regain equilibrium.

However, in the separated region ($c_f \leq 0$) the sensitivity for both c_f and c_h shows a reduction in accuracy for sensitivity prediction with M2 inferred variables as compared to the one with M1 inferred variables. The results with M2 variables show a similar qualitative behavior to the results obtained from M1 variables, having a similar location of the peak values (both negative and positive). It is the magnitude of these peak values that differs a lot. M2 variables overpredict these peaks (in terms of magnitude) compared to the prediction from the M1 variables, which in turn overpredicts these peaks when compared with LES sensitivities.

In fig. 4.13 and fig. 4.14 comparison between MFSA sensitivity with M1 and M2 inference method is made for wall-normal profiles at a streamwise location just before separation (marked in fig. 4.2). At this location, largest variation among the M1 and M2 inferred fields is seen. In congruence with the observations in the sensitivity of wall




Figure 4.13: Comparison of sensitivities of u_1 wall-normal profiles w.r.t T_w and β , just before separation, between M1- and M2-inferred fields. The gray region marks the 95% confidence interval of the LES predictions.




Figure 4.14: Comparison of sensitivities of T wall-normal profiles w.r.t T_w and β , just before separation, between M1- and M2-inferred fields. The gray region marks the 95% confidence interval of the LES predictions.

quantities such as c_f and c_h , M1 inferred fields performs better than M2 inferred sensitivity.

In spite of M2 sensitivity results behaving contrary to expectations, they throw light on the fact that the inferred parameters can have an influence on the MFSA sensitivity predictions. Improving the inference process by altering the cost function can improve the sensitivity prediction by MFSA. This could be a possible future direction for improvement of MFSA results.

4.8 Concluding remarks

The process of inference has an impact on the accuracy of the sensitivity predicted by the MFSA. The M2 inference framework, which models the term directly appearing in the conservation equations was envisioned to give a better accuracy for sensitivity prediction, however that was not the case. They do a better job at computing the closest flow-field to LES at baseline conditions without any forcing. In spite of the relatively poor performance of M2, it does motivate finding better ways to infer the RANS auxiliary variables from the LES data. One way could be to minimize the residual of the LES data i.e., find Π^{LES} that minimizes $\|M(\bar{\mathbf{v}}^{\text{LES}}, \Pi^{\text{LES}}; \mathbf{a}_b)\|$. This is exactly what M2 framework does, but just for the momentum equation. Including the energy equation would lead to a coupled system of equations for μ_t and κ_t . Alternatively, the auxiliary equations can be solved with the LES data to obtain inferred fields. However, that process may still require constraining the RANS solver as well. Although computationally expensive, the fields can be inferred such that there is no need for constraining the RANS equations (cf. [Hayek et al., 2018](#)). Here the

cost function reads,

$$\|\bar{\mathbf{v}} - \bar{\mathbf{v}}^{\text{LES}}\|_{\Omega} \quad \text{s.t.} \quad M(\bar{\mathbf{v}}, \Pi; \mathbf{a}_b) .$$

The Π^{LES} obtained from such a framework will lead to a standalone RANS system being as close to LES data as possible, at the base conditions. The only downside of this is its computational cost. It may need many RANS simulations to be performed. Since the main objective of this work is to find the sensitivity at a very small fraction of the LES cost, this method in the current setup does not fit the bill.

However, having said that, maybe finding “best” inferred fields may not be the correct direction forward. Through construction, the MFSA equations (cf. eqn. (2.8)) ensures that the baseline LES data will be reproduced by MFSA for any choice of Π^{LES} . The research focus should be spent on finding better predictive models for accounting changes in Π under perturbations. The inferred fields at base condition does influence the accuracy, but arguably, most of the differences between LES and MFSA sensitivities results come from the RANS model behavior. Improved inference methods at perturbed states can be used to re-calibrate the RANS models to improve MFSA results.

Chapter 5: Summary

The work presented in the thesis assesses the feasibility and accuracy of the proposed multi-fidelity sensitivity algorithm (MFSA) to estimate the parametric sensitivity of an LES calculation. Calculating the sensitivity of LES in a computationally affordable way is a challenging task. The approach taken in this thesis is to find a RANS system that replicates the LES information at the base state, and uses it to find the change of flow state w.r.t. any perturbations to an input parameter. The process of finding an effective RANS system for LES requires inferring the RANS model values from LES data and also constraining the RANS system to replicate LES data. Here, the RANS system is more intrusively altered to ensure it matches LES data. This is in contrast to the classical multi-fidelity methods where each method is used as a black box without one model influencing the usage of another.

Initially, the keystone idea was perceived to be the inference of RANS model variables (cf. chapter 4). It was assumed that the frozen turbulence assumption along with optimally inferred RANS variables would suffice. The bulk of the effort in the thesis was spent there in trying to come up with well-posed cost functions that provide optimal and smooth RANS variables. However, as the sister work of the thesis, [Arias-Ramírez \(2022\)](#) showed that frozen turbulence assumption is not sufficient and the RANS model parameters are also needed to be perturbed along with the flow state, requiring a RANS-like closure model for

the perturbed RANS model parameters. The focus shifted to assessing the feasibility of the method with a known RANS model, the SA model in this work. The next challenge was to ensure that the RANS replicates the LES data. No modeling would ensure that all the LES data satisfies a RANS model exactly. So a constraining term was added to ensure that a RANS model always replicates the data for any choice of inferred model variables (cf. chapter 3). Alternatively, [Arias-Ramírez \(2022\)](#) used linearized RANS equations about LES data to have a similar effect. [Arias-Ramírez \(2022\)](#) used an algebraic RANS model, a natural question from there on was if, using a higher order RANS model can improve the accuracy? The results in this work in conjunction with [Arias-Ramírez \(2022\)](#) suggests that simply using higher order closures may not suffice.

The results discussed in the thesis indicate that the proposed MFSA, in its current form, can provide qualitative improvements in the sensitivity estimation as compared to that of a standalone RANS system. This is especially true for flow regimes such as separation where RANS models are known to lack accuracy. However, standalone RANS performs reasonably well in flow regimes such as wall attached flows, for which they have been calibrated. In conclusion, the proposed MFSA probably offers a potentially feasible method to predict the sensitivity of an LES base flow, but requires modifications to improve its accuracy. The MFSA is able to compute the parametric sensitivity at a fraction of a cost compared to that of running an LES, which was one of the goals of the thesis. Also, inferring “best” RANS model variables from LES data at base condition seems to be insufficient when considering sensitivity results.

5.1 Future directions

The proposed MFSA needs modifications to the process to improve the accuracy of the estimated sensitivity. The results from different inference processes (cf. sec. 4.7.2) show that the MFSA accuracy is influenced by the inference process, which could be a pathway to improve the accuracy. A few ideas related to the same are listed in sec. 4.8.

Other possible avenues could be to modify the coefficients that appear in the transport equations of the RANS auxiliary variables. For instance, eqn. 2.9 contains many coefficients that have been tailored to produce sufficiently accurate standalone results. These coefficients could now be re-calibrated such that they are better at predicting the sensitivity instead.

A lot of RANS-based closures rely on the Boussinesq approximation (cf. eqn. 2.5) which have been known to work great for certain flow regimes but not all of them. Probably using the transport equation for the full Reynolds stress tensor R_{ij} could work better.

If there is a situation where there are $O(10)$ or $O(100)$ input parameters, one could run a couple more LESs at states other than the base state and using the additional LES information to calibrate models. This information could then be used to re-calibrate the model to estimate the full sensitivity at a substantially affordable computational cost. For instance, in the SBLI example (cf. chapter 3) in this thesis, we could have used another LES at a perturbed T_w condition and used that to calibrate the SA model and then computed the sensitivity w.r.t β . It would be more expensive than the current proposed method but could be potentially more accurate. The cost advantage would be more prominent when the number of input parameters is large.

If multiple LESs are being run, then data-driven modeling approaches such machine

learning or reduced order model approaches such as proper orthogonal decomposition (POD) could also be used as lower-fidelity prediction models for sensitivity. Alternatively, better inference methods can be used to re-calibrate the RANS model coefficients which could improve the performance of the method, if perturbed LESs are being run.

All the methods in discussion together or in a standalone setup could assumably help MFSA improve the accuracy of its sensitivity prediction for LES calculations.

Appendix A: Calculation of the Euler-Lagrange equations

In this chapter, the derivations of the Euler-Lagrange equations for cost functionals in chapter 4 are discussed. The Euler-Lagrange equations are used to obtain stationary points of a functional acting on an infinite dimensional space. Functionals are mathematical objects that take an element of a vector space and return a scalar value. The cost functions $\mathcal{J}$ that were discussed in chapter 4 are functionals acting on an infinite dimensional space of functions. Vector spaces such as $\mathbb{R}^n$ are finite dimensional vector spaces of dimension n . Functional spaces on the other hand are of infinite dimension. For instance, consider the space of analytical functions C^∞ . Any function $f \in C^\infty$ can be expressed in the form of a Taylor series, so this space C^∞ has the basis $\{x^i\}_{i=0}^\infty$, which as infinite number of elements so the underlying functional space is of infinite dimension. This difference of dimensionality is what makes optimizing for functionals acting on function spaces different from the classical finite dimensional optimization techniques.

To find the optimal point for a finite dimensional functional $\mathcal{F}$ acting on $\mathbb{R}^n$, first the stationary point is found by setting the $\nabla \mathcal{F} = 0$ and from this set the optimal is found. If $\mathcal{F}$ is convex then the stationary point will also be the minima. For functional J on the functional space a similar approach is used, where instead of the gradient there is the variation $\delta \mathcal{J}$. Consider a field $\phi(\mathbf{x})$ on the domain $\Omega \subseteq \mathbb{R}^3$. Here $\mathbf{x}$ is a point in the domain Ω . A general

cost functional $\mathcal{J}$ for such a field ϕ is

$$\mathcal{J} = \int_{\Omega} L(\phi, \nabla_x \phi, \mathbf{x}) d\Omega. \quad (\text{A.1})$$

The field ϕ that optimizes the cost functional satisfies the PDE given by,

$$\frac{\partial L}{\partial \phi} - \frac{\partial}{\partial x_j} \left(\frac{\partial L}{\partial (\partial \phi / \partial x_j)} \right) = 0, \quad (\text{A.2})$$

with appropriate boundary conditions (cf. [Goldstein et al., 2002](#); [Liberzon, 2011](#)). Here we use the Einstein indicial notation where repeated indices imply summation. For all the cases in chapter 4, ϕ could be μ_t or κ_t and Ω is the LES domain.

A.1 M1 inference equations

The M1 equations (cf. eqn. 4.3) have $L = E_{ij}E_{ij}$ for μ_t and $L = e_k e_k$ for κ_t . Note that for the purposes of the calculation, the R_{ij} term has been scaled with $\bar{\rho}$, and superscript d denoting the deviatoric part is dropped for brevity. The scaling terms used to non-dimensionalize the cost are not used here as they are constants, they just pre-multiply the equation.

For μ_t :

$$L = E_{ij}E_{ij} = R_{ij}R_{ij} + 4\mu_t R_{ij}S_{ij} + 4\mu_t^2 S_{ij}S_{ij}$$

Plugging in the E-L equations (eqn. A.2),

$$\begin{aligned}\frac{\partial L}{\partial \mu_t} &= 4R_{ij}S_{ij} + 8\mu_t S_{ij}S_{ij} \\ \frac{\partial L}{\partial (\partial \mu_t / \partial x_j)} &= 0.\end{aligned}$$

The overall equation reads,

$$4R_{ij}S_{ij} + 8\mu_t S_{ij}S_{ij} = 0.$$

For κ_t :

$$L = e_k e_k = h_{Tk} h_{Tk} + 2\kappa_t g_k h_{Tk} + \kappa_t^2 g_k g_k$$

Plugging in the E-L equations (eqn. A.2),

$$\begin{aligned}\frac{\partial L}{\partial \kappa_t} &= 2g_k h_{Tk} + 2\kappa_t g_k g_k \\ \frac{\partial L}{\partial (\partial \kappa_t / \partial x_j)} &= 0.\end{aligned}$$

The overall equation reads,

$$2g_k h_{Tk} + 2\kappa_t g_k g_k = 0.$$

A.2 M2 inference equations

The M2 equations (cf. eqn. 4.5) have $L = B_k B_k$ for μ_t , where $B_k = \partial E_{kj} / \partial x_j$, and $L = b^2$ for κ_t , where $b = \partial e_j / \partial x_j$.

For μ_t :

$$L = B_k B_k \quad B_k = \frac{\partial R_{ki}}{\partial x_i} + 2\mu_t \frac{\partial S_{ki}}{\partial x_i} + 2 \frac{\partial \mu_t}{\partial x_i} S_{ki}$$

Plugging in the E-L equations (eqn. A.2),

$$\begin{aligned} \frac{\partial L}{\partial \mu_t} &= 2B_k \frac{\partial B_k}{\partial \mu_t} = 4B_k \frac{\partial S_{ki}}{\partial x_i} \\ \frac{\partial L}{\partial (\partial \mu_t / \partial x_j)} &= 2B_k \frac{\partial B_k}{\partial (\partial \mu_t / \partial x_j)} = 4B_k S_{kj} . \end{aligned}$$

The overall equation reads,

$$\begin{aligned} 4B_k \frac{\partial S_{kj}}{\partial x_j} - \frac{\partial}{\partial x_j} (4B_k S_{kj}) &= 0 , \\ \implies -4 \frac{\partial B_k}{\partial x_j} S_{kj} &= 0 , \\ \implies -8 \frac{\partial}{\partial x_j} \left(\frac{\partial \mu_t}{\partial x_i} S_{ki} \right) S_{kj} - 8 \frac{\partial S_{ki}}{\partial x_i} S_{kj} \frac{\partial \mu_t}{\partial x_j} \\ &\quad - 8 \frac{\partial}{\partial x_j} \left(\frac{\partial S_{ki}}{\partial x_i} \right) S_{kj} \mu_t - 4 \frac{\partial}{\partial x_j} \left(\frac{\partial R_{ki}}{\partial x_i} \right) S_{kj} = 0 . \end{aligned}$$

The first term can be re-written as,

$$\frac{\partial}{\partial x_j} \left(\frac{\partial \mu_t}{\partial x_i} S_{ki} \right) S_{kj} = \frac{\partial}{\partial x_j} \left(S_{kj} S_{ki} \frac{\partial \mu_t}{\partial x_i} \right) - \frac{\partial S_{ki}}{\partial x_i} S_{kj} \frac{\partial \mu_t}{\partial x_j}.$$

Combing all the terms together we get,

$$-8 \frac{\partial}{\partial x_j} \left(S_{kj} S_{ki} \frac{\partial \mu_t}{\partial x_i} \right) - 8 \frac{\partial}{\partial x_j} \left(\frac{\partial S_{ki}}{\partial x_i} \right) S_{kj} \mu_t - 4 \frac{\partial}{\partial x_j} \left(\frac{\partial R_{ki}}{\partial x_i} \right) S_{kj} = 0$$

For κ_t :

$$L = b^2 \quad b = \frac{\partial h_{Ti}}{\partial x_i} + \kappa_t \frac{\partial g_i}{\partial x_i} + 2 \frac{\partial \kappa_t}{\partial x_i} g_i$$

Plugging in the E-L equations (eqn. [A.2](#)),

$$\begin{aligned} \frac{\partial L}{\partial \kappa_t} &= 2b \frac{\partial b}{\partial \kappa_t} = 2b \frac{\partial g_i}{\partial x_i} \\ \frac{\partial L}{\partial (\partial \kappa_t / \partial x_j)} &= 2b \frac{\partial b}{\partial (\partial \kappa_t / \partial x_j)} = 2b g_j. \end{aligned}$$

The overall equation reads,

$$\begin{aligned} 4b \frac{\partial g_j}{\partial x_j} - \frac{\partial}{\partial x_j} (4b g_j) &= 0, \\ \implies -2 \frac{\partial b}{\partial x_j} g_j &= 0, \\ \implies -2 \frac{\partial}{\partial x_j} \left(\frac{\partial \kappa_t}{\partial x_i} g_i \right) g_j - 2 \frac{\partial g_i}{\partial x_i} g_j \frac{\partial \kappa_t}{\partial x_j} \\ &\quad - 2 \frac{\partial}{\partial x_j} \left(\frac{\partial g_i}{\partial x_i} \right) g_j \kappa_t - 2 \frac{\partial}{\partial x_j} \left(\frac{\partial h_{Ti}}{\partial x_i} \right) g_j = 0. \end{aligned}$$

The first term can be re-written as,

$$\frac{\partial}{\partial x_j} \left(\frac{\partial \kappa_t}{\partial x_i} g_i \right) g_j = \frac{\partial}{\partial x_j} \left(g_j g_i \frac{\partial \kappa_t}{\partial x_i} \right) - \frac{\partial g_i}{\partial x_i} g_j \frac{\partial \kappa_t}{\partial x_j}.$$

Combing all the terms together we get,

$$-2 \frac{\partial}{\partial x_j} \left(g_j g_i \frac{\partial \kappa_t}{\partial x_i} \right) - 2 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial g_k}{\partial x_k} \right) g_i \right) \kappa_t - 2 \left(\frac{\partial}{\partial x_i} \left(\frac{\partial h_{Tk}}{\partial x_k} \right) g_i \right) = 0$$

A.3 Regularization term

In absence of $|S|$ and $|g|$, the Boussinesq approximation (cf. eqn. 2.5 and eqn. 2.6) is ill-defined. An augmentation to the cost function is provided (cf. eqn. 4.8). For this augmentation the $L = \chi_s (\mu_t - \mu_t^\infty)$ for μ_t and $L = \chi_g (\kappa_t - \kappa_t^\infty)$ for κ_t

For μ_t :

$$\begin{aligned} \frac{\partial L}{\partial \mu_t} &= 2\chi_s (\mu_t - \mu_t^\infty), \\ \frac{\partial L}{\partial (\partial \mu_t / \partial x_j)} &= 0. \end{aligned}$$

Plugging in the E-L equations (eqn. A.2), the overall equation reads,

$$2\chi_s (\mu_t - \mu_t^\infty) = 0$$

For κ_t :

$$\begin{aligned}\frac{\partial L}{\partial \kappa_t} &= 2\chi_g (\kappa_t - \kappa_t^\infty) , \\ \frac{\partial L}{\partial (\partial \kappa_t / \partial x_j)} &= 0 .\end{aligned}$$

Plugging in the E-L equations (eqn. A.2), the overall equation reads,

$$2\chi_g (\kappa_t - \kappa_t^\infty) = 0$$

A.4 Smoothness constraint

The smoothness cost function penalizes the gradient of the fields. This has $L = \partial \mu_t / \partial x_k \partial \mu_t / \partial x_k$ for μ_t and $L = \partial \kappa_t / \partial x_k \partial \kappa_t / \partial x_k$ for κ_t .

For μ_t :

$$\begin{aligned}\frac{\partial L}{\partial \mu_t} &= 0 \\ \frac{\partial L}{\partial (\partial \mu_t / \partial x_j)} &= -2 \frac{\partial \mu_t}{\partial x_j} .\end{aligned}$$

Plugging in the E-L equations (eqn. A.2), the overall equation reads,

$$-2 \frac{\partial^2 \mu_t}{\partial x_j \partial x_j} = 0$$

For κ_t :

$$\begin{aligned}\frac{\partial L}{\partial \kappa_t} &= 0 \\ \frac{\partial L}{\partial (\partial \kappa_t / \partial x_j)} &= -2 \frac{\partial \kappa_t}{\partial x_j} .\end{aligned}$$

Plugging in the E-L equations (eqn. [A.2](#)), the overall equation reads,

$$-2 \frac{\partial^2 \kappa_t}{\partial x_j \partial x_j} = 0$$

It is interesting to note that the Laplacian operator is an output of minimizing the gradient of a field over the domain.

Bibliography

- ALLMARAS, STEVEN R & JOHNSON, FORRESTER T 2012 Modifications and clarifications for the implementation of the spalart-allmaras turbulence model. In *Seventh international conference on computational fluid dynamics (ICCFD7)*, , vol. 1902. ICCFD7-1902 Big Island, Hawaii.
- ARIAS-RAMÍREZ, WALTER 2022 A multi-fidelity approach to sensitivity estimation in large eddy simulations. PhD thesis, University of Maryland, College Park.
- ARIAS-RAMÍREZ, WALTER & LARSSON, JOHAN 2022 Grid sufficiency in large eddy simulations as a hypothesis test. *International Journal of Computational Fluid Dynamics* **36** (3), 260–264.
- BELL, NATHAN, OLSON, LUKE N. & SCHRODER, JACOB 2022 PyAMG: Algebraic multigrid solvers in python. *Journal of Open Source Software* **7** (72), 4142.
- BERNARDINI, M., ASPROULIAS, I., LARSSON, J., PIROZZOLI, S. & GRASSO, F. 2016 Heat transfer and wall temperature effects in shock wave turbulent boundary layer interactions. *Phys. Rev. Fluids* **1**, 084403.
- BLONIGAN, PATRICK J & WANG, QIQI 2018 Multiple shooting shadowing for sensitivity analysis of chaotic dynamical systems. *Journal of Computational Physics* **354**, 447–475.
- BODONY, DANIEL J 2006 Analysis of sponge zones for computational fluid mechanics. *Journal of Computational Physics* **212** (2), 681–702.
- BRUNTON, STEVEN L, NOACK, BERND R & KOUMOUTSAKOS, PETROS 2020 Machine learning for fluid mechanics. *Annual review of fluid mechanics* **52**, 477–508.
- BUSH, ROBERT H, CHYCZEWSKI, THOMAS S, DURAISAMY, KARTHIKEYAN, EISFELD, BERNHARD, RUMSEY, CHRISTOPHER L & SMITH, BRIAN R 2019 Recommendations for future efforts in rans modeling and simulation. In *AIAA scitech 2019 forum*, p. 0317.
- CHANDRAMOORTHY, NISHA, FERNÁNDEZ, PABLO, TALNIKAR, CHAITANYA & WANG, QIQI 2017 The ensemble adjoint method for chaotic cfd: A feasibility study. In *23rd AIAA Computational Fluid Dynamics Conference, Denver, USA*.

- CHANDRAMOORTHY, N, WANG, ZN, WANG, Q & TUCKER, P 2018 Toward computing sensitivities of average quantities in turbulent flows. In *Proceedings of the Summer Program*, p. 1.
- CLEMENS, NOEL T & NARAYANASWAMY, VENKATESWARAN 2014 Low-frequency unsteadiness of shock wave/turbulent boundary layer interactions. *Annual Review of Fluid Mechanics* **46**, 469–492.
- CROUSEILLES, NICOLAS, KUHN, MATTHIEU & LATU, GUILLAUME 2015 Comparison of numerical solvers for anisotropic diffusion equations arising in plasma physics. *Journal of Scientific Computing* **65**, 1091–1128.
- DHAMANKAR, NITIN S, BLAISDELL, GREGORY A & LYRINTZIS, ANASTASIOS S 2018 Overview of turbulent inflow boundary conditions for large-eddy simulations. *Aiaa Journal* **56** (4), 1317–1334.
- DI RENZO, MARIO, OBEROI, NIKHIL, LARSSON, JOHAN & PIROZZOLI, SERGIO 2022 Crossflow effects on shock wave/turbulent boundary layer interactions. *Theoretical and Computational Fluid Dynamics* **36** (2), 327–344.
- FIDKOWSKI, KRZYSZTOF J & DARMOFAL, DAVID L 2011 Review of output-based error estimation and mesh adaptation in computational fluid dynamics. *AIAA journal* **49** (4), 673–694.
- GAITONDE, DATTA V 2015 Progress in shock wave/boundary layer interactions. *Progress in Aerospace Sciences* **72**, 80–99.
- GAITONDE, DATTA V & ADLER, MICHAEL C 2023 Dynamics of three-dimensional shock-wave/boundary-layer interactions. *Annual Review of Fluid Mechanics* **55**.
- GERACI, GIANLUCA, ELDRED, MICHAEL S & IACCARINO, GIANLUCA 2017 A multifidelity multilevel monte carlo method for uncertainty propagation in aerospace applications. In *19th AIAA non-deterministic approaches conference*, p. 1951.
- GOLDSTEIN, H., POOLE, C.P. & SAFKO, J.L. 2002 *Classical Mechanics*. Addison Wesley.
- HAYEK, MICHAEL E, WANG, QIQI & LASKOWSKI, GREGORY M 2018 Adjoint-based optimization of rans eddy viscosity model for u-bend channel flow. In *2018 AIAA Aerospace Sciences Meeting*, p. 2091.
- HOSDER, SERHAT, WALTERS, ROBERT & BALCH, MICHAEL 2007 Efficient sampling for non-intrusive polynomial chaos applications with multiple uncertain input variables. In *48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, p. 1939.
- JAMESON, ANTONY 1988 Aerodynamic design via control theory. *Journal of scientific computing* **3** (3), 233–260.

- JOHNSEN, ERIC, LARSSON, JOHAN, BHAGATWALA, ANKIT V, CABOT, WILLIAM H, MOIN, PARVIZ, OLSON, BRITTON J, RAWAT, PRADEEP S, SHANKAR, SANTHOSH K, SJÖGREEN, BJÖRN, YEE, HELEN C *ET AL.* 2010 Assessment of high-resolution methods for numerical simulations of compressible turbulence with shock waves. *Journal of Computational Physics* **229** (4), 1213–1237.
- KHAN, MD RAIHAN ALI & HASAN, ABM TOUFIQUE 2021 Design and evaluation of generic bump for flow control in a supersonic inlet isolator. *Journal of Fluids Engineering* **143** (5), 051207.
- KIM, JEONGLAE, BODONY, DANIEL J & FREUND, JONATHAN B 2014 Adjoint-based control of loud events in a turbulent jet. *Journal of Fluid Mechanics* **741**, 28–59.
- LARSSON, JOHAN 2018 Grid-adaptation for chaotic multi-scale simulations as a verification-driven inverse problem. In *2018 AIAA Aerospace Sciences Meeting*, p. 0371.
- LARSSON, JOHAN, KUMAR, VEDANT, OBEROI, NIKHIL, RENZO, MARIO DI & PIROZZOLI, SERGIO 2022 Large-eddy simulations of idealized shock/boundary-layer interactions with crossflow. *AIAA Journal* **60** (5), 2767–2779.
- LEA, DANIEL J, ALLEN, MYLES R & HAINE, THOMAS WN 2000 Sensitivity analysis of the climate of a chaotic system. *Tellus A: Dynamic Meteorology and Oceanography* **52** (5), 523–532.
- LEVEQUE, RANDALL J 2002 *Finite volume methods for hyperbolic problems*, , vol. 31. Cambridge university press.
- LIBERZON, DANIEL 2011 *Calculus of variations and optimal control theory: a concise introduction*. Princeton university press.
- LUCHINI, PAOLO & BOTTARO, ALESSANDRO 2014 Adjoint equations in stability analysis. *Annual Review of fluid mechanics* **46**, 493–517.
- NAJM, HABIB N 2009 Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics. *Annual review of fluid mechanics* **41**, 35–52.
- NI, ANGXIU & WANG, QIQI 2017 Sensitivity analysis on chaotic dynamical systems by non-intrusive least squares shadowing (nilss). *Journal of Computational Physics* **347**, 56–77.
- OBERKAMPF, WILLIAM L & TRUCANO, TIMOTHY G 2002 Verification and validation in computational fluid dynamics. *Progress in aerospace sciences* **38** (3), 209–272.
- PEHERSTORFER, BENJAMIN, WILLCOX, KAREN & GUNZBURGER, MAX 2018 Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *Siam Review* **60** (3), 550–591.
- POPE, STEPHEN B 2000 *Turbulent flows*. Cambridge university press.

- QIAN, ELIZABETH, PEHERSTORFER, BENJAMIN, O'MALLEY, DANIEL, VESSELINOV, VELIMIR V & WILLCOX, KAREN 2018 Multifidelity monte carlo estimation of variance and sensitivity indices. *SIAM/ASA Journal on Uncertainty Quantification* **6** (2), 683–706.
- RUMSEY, CHRISTOPHER 2013 2-d zero pressure gradient flat plate verification case. <https://turbmodels.larc.nasa.gov/flatplate.html>.
- SAAD, YOUSEF 2003 *Iterative methods for sparse linear systems*. SIAM.
- SAGAUT, PIERRE 2006 *Large eddy simulation for incompressible flows: an introduction*. Springer Science & Business Media.
- SCHLICHTING, HERMANN & GERSTEN, KLAUS 2016 *Boundary-layer theory*. springer.
- SHU, CHI-WANG 1999 High order eno and weno schemes for computational fluid dynamics. *High-order methods for computational physics* pp. 439–582.
- SIVIA, DEVINDERJIT & SKILLING, JOHN 2006 *Data analysis: a Bayesian tutorial*. OUP Oxford.
- ŚLIWIAK, ADAM A & WANG, QIQI 2022 Space-split algorithm for sensitivity analysis of discrete chaotic systems with multidimensional unstable manifolds. *SIAM Journal on Scientific Computing* **44** (5), A3290–A3316.
- ŚLIWIAK, ADAM A & WANG, QIQI 2023 Approximating linear response of physical chaos. *Nonlinear Dynamics* **111**, 1835–1869.
- SLOTNICK, JEFFREY P, KHODADOUST, ABDOLLAH, ALONSO, JUAN, DARMOFAL, DAVID, GROPP, WILLIAM, LURIE, ELIZABETH & MAVRIPLIS, DIMITRI J 2014 Cfd vision 2030 study: a path to revolutionary computational aerosciences. *Tech. Rep.*.
- SPALART, PHILIPPE & ALLMARAS, STEVEN 1992 A one-equation turbulence model for aerodynamic flows. In *30th aerospace sciences meeting and exhibit*, p. 439.
- TOOSI, SIAVASH 2019 Error estimation, grid selection and convergence verification in large eddy simulation. PhD thesis, University of Maryland, College Park.
- TOOSI, SIAVASH & LARSSON, JOHAN 2020 Towards systematic grid selection in les: Identifying the optimal spatial resolution by minimizing the solution sensitivity. *Computers & Fluids* **201**, 104488.
- VOLPIANI, PEDRO S., BERNARDINI, MATTEO & LARSSON, JOHAN 2018 Effects of a nonadiabatic wall on supersonic shock/boundary-layer interactions. *Phys. Rev. Fluids* **3**, 083401.
- WALZ, ALFRED 1969 *Boundary layers of flow and temperature*. MIT press.
- WANG, JIAN-XUN, WU, JIN-LONG & XIAO, HENG 2017 Physics-informed machine learning approach for reconstructing reynolds stress modeling discrepancies based on dns data. *Physical Review Fluids* **2** (3), 034603.

- WANG, QIQI 2013 Forward and adjoint sensitivity computation of chaotic dynamical systems. *Journal of Computational Physics* **235**, 1–13.
- WANG, QIQI 2014 Convergence of the least squares shadowing method for computing derivative of ergodic averages. *SIAM Journal on Numerical Analysis* **52** (1), 156–170.
- WILCOX, DAVID C 1998 *Turbulence modeling for CFD*, , vol. 2. DCW industries La Canada, CA.
- WILCOX, DAVID C *ET AL.* 1998 *Turbulence modeling for CFD*, , vol. 2. DCW industries La Canada, CA.
- WU, JINLONG, XIAO, HENG, SUN, RUI & WANG, QIQI 2019 Reynolds-averaged navier–stokes equations with explicit data-driven reynolds stress closure can be ill-conditioned. *Journal of Fluid Mechanics* **869**, 553–586.
- WU, XIAOHUA 2017 Inflow turbulence generation methods. *Annual Review of Fluid Mechanics* **49**, 23–49.
- ZHANG, YOU-SHENG, BI, WEI-TAO, HUSSAIN, FAZLE & SHE, ZHEN-SU 2014 A generalized reynolds analogy for compressible wall-bounded turbulent flows. *Journal of Fluid Mechanics* **739**, 392–420.