ABSTRACT

Title of Proposal: A MULTI-FIDELITY APPROACH TO SENSITIVITY ESTIMATION IN LARGE EDDY SIMULATIONS

Walter Arias-Ramírez
Doctor of Philosophy, 2022

Dissertation directed by: Professor Johan Larsson
Department of Mechanical Engineering

An approach to compute approximate sensitivities in a large eddy simulation (LES) is proposed and assessed. The multi-fidelity sensitivity analysis (MFSA) solves a linearized mean equation, where the mean equation is based on the LES solution. This requires closure modeling which makes the computed sensitivities approximate. The closure modeling is based on inferring the eddy viscosity from the LES data and in predicting the change in turbulence (or the perturbed eddy viscosity) using a simple algebraic model. The method is assessed for the flow over a NACA0012 airfoil at a fixed angle of attack, with the Reynolds number as the varying parameter and the lift, drag, skin friction, and pressure coefficients as the quantities-of-interest. The results show the importance of accurate closure modeling, specifically that treating the eddy viscosity as "frozen" is insufficiently accurate. Also, predictions obtained using the algebraic model for closing the perturbed eddy viscosity are closer to the true sensitivity than results obtained using the fully RANS-based method which is the state-of-the-art and most common method used in industry. The proposed method aims to complement, rather than replace, the current state-of-the-art method in situations in which sensitivities with higher fidelity are required.

A MULTI-FIDELITY APPROACH TO SENSITIVITY ESTIMATION IN LARGE EDDY SIMULATIONS

by

Walter Arias-Ramírez

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
2022

Advisory Committee:
Professor Dr. Johan Larsson, Advisor/Chair
Professor Dr. Amir Riaz,
Professor Dr. James D. Baeder, Dean's representative
Professor Dr. James H. Duncan
Professor Dr. Kenneth T. Kiger

To the loving memory of my father
Álvaro Arias Pacheco.
Where would I be without your love?

Soy
Soy lo que dejaron
Soy toda la sobra de lo que se robaron
Mano de obra campesina para tu consumo
Soy el desarrollo en carne viva
Mi piel es de cuero, por eso aguanta cualquier clima
Soy un mojado cruzando la frontera
Soy la fotografía de un desaparecido
Soy un falso positivo militar
Soy aquel por quien lloran las madres de Soacha
Soy un indigena Kogui defendiendo la Sierra Nevada
Soy un niño forzosamente reclutado por la guerrilla
Soy Aureliano Buendía frente al pelotón de fusilamiento
Soy Colombia contra Alemania anotándote en el 92'
Soy Herrera, Nairo y Egan ganándote en la montaña
Soy Maradona contra Inglaterra anotándote dos goles
Soy la sangre dentro de tus venas
Soy lo que me enseñó mi padre
Soy América latina
Un pueblo sin piernas, pero que camina
Aquí estamos de pie
(Parts extracted from Latinoamérica by Calle 13)

Acknowledgments

I owe my gratitude to all the people who have made this thesis possible.

I would like to thank Dr. Larsson for his academic orientation and for allowing me to be part of his research projects. I am indebted to the agencies that sponsored my research. I would like to thank the National Science Foundation for supporting this research through the career grant CBET-143633. I would like to thank the Predictive Science Academic Alliance Program III for supporting this research. I would like to thank the administrators of the supercomputers I used during this research: the University of Maryland's Deepthought2, Johns Hopkins' MARCC (especially to Dr. Combariza), and most importantly, the Livermore's computing system (especially for the ruby and borax machines). Without all of them, I could not have been able to finish this thesis. I would like to thank Dr. Wolf and Dr. Olson for supporting me to start this Ph.D.

Empathy was, perhaps, the best companion during this journey and it can only be felt and shared fully by peers who also are going, or went, throughout the same experiences. For that, I would like to thank my friends and colleagues for livening up the strenuous, challenging, and burdensome experiences that a Ph.D. entails. Thank you, Emma, Abishek, and Nikhil, your direct and indirect contributions to this document are invaluable. Thank you, Erick, Luis, Ruben, Marta, Sergio, Andrés, Daniel, David, Nelson, Victoria, Ali, Vedant, John, Markus, and Raphael for sharing with me during the last 5 years. I will miss the meetings at lunchtime on campus and the Wednesday nights in this town. I would like to thank Guillermo Díaz Salamanca and his team for el Cocuyo and el Tren de la Tarde. These radio shows were the perfect companion during the long nights. Finally, I would like to thank my family for their support and patience throughout this journey.

Table of Contents

List of Tables

List of Figures

List of Algorithms

Definitions

$\mathbf{a}$: set of parameters in the design space $a_1, a_2, ...a_n$.

$\mathbf{v}$: represent the solution fields velocity u_i, and pressure p. $\mathbf{v} = [u_i, p]$.

$\delta\mathbf{v}$: perturbed solution fields: velocity δu_i, and pressure δp.

c: airfoil chord.

d: wall normal distance.

l_z airfoil spanwise length.

J_j: j-th quantity of interest.

ν: kinematic viscosity.

ν_t: eddy viscosity.

δ: boundary layer thickness.

Re: Reynolds number based on an unitary chord (c=1).

θ: angle of attack with respect to the horizontal plane.

R: auto-correlation coefficient.

τ: delayed time.

τ_w: wall shear stress.

u_τ: friction velocity.

$q_\infty = \frac{1}{2}\rho u^2$; dynamic pressure.

$\mathcal{C}_f = \tau_w / q_\infty$; skin friction coefficient.

$\mathcal{C}_p = (p_w - p_\infty)/q_\infty$; pressure coefficient.

D: drag force.

L: is the lift force.

$A = l_z c$; is the wing area.

$\mathcal{C}_D = D/q_\infty A$; instantaneous aerodynamic drag coefficient.

$\mathcal{C}_L = L/q_\infty A$; instantaneous aerodynamic lift coefficient.

n_b: number of batches in which the signal of the aerodynamic coefficients is divided with the intention to obtain independent samples of these coefficients.

n_i: number of samples in batch i.

$t_b = \Delta t n_i$: time per batch.

$\widehat{\mathcal{C}_{D,i}}$: aerodynamic drag coefficient mean of batch i. $\widehat{\mathcal{C}_{D,i}} = \frac{1}{n_i}\sum_{j=1}^{n_i} \mathcal{C}_{D,j}$.

$\widehat{\mathcal{C}_{L,i}}$: aerodynamic lift coefficient mean of a batch i. $\widehat{\mathcal{C}_{L,i}} = \frac{1}{n_i}\sum_{j=1}^{n_i} \mathcal{C}_{D,j}$.

$\overline{\mathcal{C}_D}$: aerodynamic drag coefficient mean of the batches mean. $\overline{\mathcal{C}_D} = \frac{1}{n_b}\sum_{i=1}^{n_b} \widehat{\mathcal{C}_{D,i}}$

$\overline{\mathcal{C}_L}$: aerodynamic lift coefficient mean of the batches mean. $\overline{\mathcal{C}_L} = \frac{1}{n_b}\sum_{i=1}^{n_b} \widehat{\mathcal{C}_{L,i}}$

$S_{\widehat{\mathcal{C}_D}}$: sample standard deviation of the means of the batches in which the aerodynamic drag coefficient signal is divided. $S_{\widehat{\mathcal{C}_D}} = \frac{1}{n_b-1}\sum_{i=1}^{n_b} \left(\widehat{\mathcal{C}_{D,i}} - \overline{\mathcal{C}_D}\right)$

$S_{\widehat{\mathcal{C}_L}}$: sample standard deviation of the means of the batches in which the aerodynamic lift coefficient signal is divided. $S_{\widehat{\mathcal{C}_L}} = \frac{1}{n_b-1}\sum_{i=1}^{n_b} \left(\widehat{\mathcal{C}_{L,i}} - \overline{\mathcal{C}_L}\right)$

$S_{\overline{\mathcal{C}_D}}$: sample standard errors for the aerodynamic drag coefficient.

$S_{\overline{\mathcal{C}_L}}$: sample standard errors for the aerodynamic lift coefficient.

α: level of significance.

Operators

$\mathcal{N}(\mathbf{v};\mathbf{a})$: LES model. It solves for $\mathbf{v}(\mathbf{a})$ with parameters $\mathbf{a}$.

$\mathcal{R}(\mathbf{v},\nu_t;\mathbf{a})$: RANS model. It solves for $\mathbf{v}(\mathbf{a})$ and $\nu_t(\mathbf{a})$. It uses the eddy viscosity constitutive model for the Reynolds stress tensor.

$\mathcal{L}(\delta\mathbf{v},\delta\nu_t;\mathbf{v},\nu_t,\mathbf{a},\delta\mathbf{a})$: linearized RANS model. It solves for $\delta\mathbf{v}(\mathbf{v},\nu_t,\mathbf{a},\delta\mathbf{a})$ and $\delta\nu_t(\mathbf{v},\nu_t,\mathbf{a},\delta\mathbf{a})$ using the base flow fields $\mathbf{v}(\mathbf{a})$ and the inferred eddy viscosity $\nu_t(\mathbf{a})$ from $\mathcal{N}(\mathbf{v};\mathbf{a}) = 0$.

Subscripts

∞: Free–stream quantity

Superscripts

$\overline{\langle\cdot\rangle}$: mean quantity

$\widetilde{\langle\cdot\rangle}$: filtered quantity

Chapter 1: Introduction

1.1 Computational methods in fluid mechanics

Computational fluid dynamics (CFD) has established itself as a useful technique in modern scientific research and engineering. The suite of CFD tools used by researchers and engineers consists of physical models that take different assumptions and approximations, and can be broadly categorized as either "high-fidelity" or "low-fidelity". The essential characteristics of these approaches, and how they are currently being used, are sketched in Fig. 1.1. High-fidelity models, like large eddy simulation (LES), make use of fewer assumptions than low-fidelity models and more accurately capture the physical phenomena, but at a higher computational cost. Low-fidelity models, like Reynolds-averaged Navier Stokes (RANS) and unsteady RANS (URANS), have been very successful in academia and in industry because they are sufficiently accurate and computationally affordable, and have become the go-to tools for the analysis of many engineering problems. Furthermore, since the pioneering work of Jameson in 1988 [5], the engineering design process has been revolutionized by integrating low-fidelity CFD tools, like the adjoint-RANS, into the decision-making process [6, 7]. For example, significantly fewer designs were physically built and tested for the Boeing 777 than for the earlier 757 and 767 programs, and without the inverse design CFD capabilities, the resulting design would not have achieved customer satisfaction [6].

Currently, RANS-based models are the standard tools used to support the engineering design and optimization processes. Although these models are sufficient to produce designs that are close to optimal for many purposes within the design process, there are certain hydrodynamic phenomena, especially near the edges of the operational envelope, that they cannot predict due to their inherent approximations and assumptions. Consequently, these models cannot accurately predict certain phenomena; for example, an airfoil near aerodynamic stall [8].

Model / Use of CFD	Analysis: single prediction with specific conditions	Design: beyond single prediction: optimization, and UQ
Low-fidelity: Potential flow, RANS, URANS	Well-established	State-of-the-art: Adjoint-RANS
High-fidelity: DES, LES, DNS	State-of-the-art	Under-explored and computationally expensive

(a)

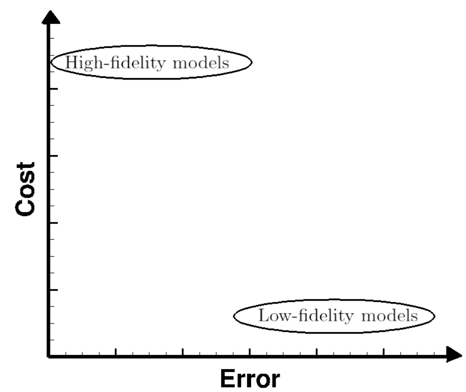

(b)

Figure 1.1: (a) Current penetration of CFD in research and industry. DES stands for detached eddy simulation and UQ stands for uncertainty quantification; (b) cost vs error for each model.

On the other hand, high-fidelity models, such as LES, have gained ground as a viable technique for flow predictions and, because of advancements in computing processing times, are becoming the preferred tool to generate physical insight in complex problems such as flow separation, aeroacoustics, and mixing, among others. LES is a turbulence modeling technique that resolves the larger flow- and geometry-dependent scales while modeling the smaller scales, which exhibit nearly universal behavior. In contrast, RANS models most or all of the turbulent activity. If one can resolve the larger scales that are of engineering interest, and model the key interactions between the larger and smaller scales of the flow accurately enough to capture the correct dynamics of the larger scales, one should be able to achieve a higher accuracy than RANS and much closer to a direct numerical simulation (DNS) but at a much lower cost. For this reason, LES offers a good trade-off and would represent a good alternative to predict the flow physics involved in conditions near the edges of the operational envelope.

However, despite the increasing interest in high-fidelity models, their usage has been restricted primarily to academia. In industry, their use is slowly being adopted, but only for analysis and rarely in the actual design process. The main obstacle preventing engineers from taking full advantage of high-fidelity tools in the design process is that they must be able to not only assess a single design, but must be able to incorporate high-fidelity tools into the iterative redesign and reassessment process. That is, in order to be able to make decisions in the design process, engineers need more than a single prediction from LES, they need to be able to integrate LES with the tools of sensitivity analysis. Consequently, as the design process converges, higher accuracy in predictions are required. This creates opportunities for using the strengths of LES.

The current research is sponsored by the Predictive Science Academic Alliance Program (PSAAP), in which one of the main goals is to create affordable methods for uncertainty quantification (UQ) of chaotic and turbulence-resolving simulations, like LES. One of the paths to enable UQ with LES, and the one chosen in this study, is by using sensitivity analysis. For that reason, the main objective of this work is to enable a computationally affordable sensitivity analysis using LES.

To define what is sensitivity, first consider J_j to be the j-th quantity-of-interest (QoI) and N_J the number of QoIs. Also, consider a_i to be i-th random parameter and N_a the number of random parameters. The sensitivity of J_j is defined as its gradient with respect to all uncertain parameters, a_i, i.e. $\partial J_j/\partial a_i$. In engineering design, some quantities-of-interest are of practical value (e.g., drag, lift, pitching moment, skin friction, and pressure profiles from a flow past an airfoil). Derivatives of the quantities-of-interest with respect to parameters are useful not only in uncertainty quantification, but also in numerical optimization applications such as machine learning, among others.

1.2 Uncertainty quantification in computational fluid dynamics

In the context of uncertainty quantification, methods to propagate uncertainty applied to CFD can be divided into two main fields: probabilistic and non-probabilistic [9]. Probabilistic methods focus on the computation of full statistics of the QoI with the ultimate goal being to use the classical statistic estimator and the variance to obtain the expected value (i.e. the model's output). Probabilistic methods can be separated into two sub-categories: sampling-based and non-sampling-based methods. A brief description of the probabilistic methods currently used in the context of LES is shown in section 1.2.2. The most known non-probabilistic methods are gradient-based algorithms that aim to compute the sensitivity of the QoI without finding its probability density function [10] and they are described in section 1.2.1.

1.2.1 Sensitivity-based methods for uncertainty quantification

The most well-known non-probabilistic method for uncertainty quantification is sensitivity analysis. Consider J_j to be the j-th QoI and N_J the number of QoIs, also consider a_i to be i-th random parameter and N_a the number of random parameters. The sensitivity of the j-th QoI, J_j, is defined as its gradient with respect to all uncertain parameters, a_i. Sensitivity analysis represents an alternative approach to the sampling-based methods. Instead of sampling the parametric space, the method computes the local gradient $\partial J_j / \partial a_i$. Although knowledge of $\partial J_j / \partial a_i$ does not provide information about the statistics of J_j (e.g., the probability distribution function), it does allow for an estimation of the variance of J_j, ${\sigma_{J_j}}^2$, through the "error-propagation formula" commonly used by experimentalists. For the most general case, the variance of J_j can be approximated as

$$\sigma_{J_j}{}^2 = \sum_{i=1}^{N_a}\sum_{k=1}^{N_a}\left[\sigma^2_{a_i a_k}\frac{\partial J_j}{\partial a_i}\frac{\partial J_j}{\partial a_k}\right], \tag{1.1}$$

where ${\sigma^2}_{a_i a_k}$ is the covariance between parameters a_i and a_k. The difficulty comes in the computation of the sensitivity gradient $\partial J_j/\partial a_i$.

The sensitivity can be computed in three different ways. The first and most common method is finite-differencing in the parameter space. In this method, the value of J_j is computed at two nearby points in the parametric space $J_j(a_i)$ and $J_j(a_i + \Delta a_i)$, in which one parameter at a time is manually perturbed. This method costs $N_a + 1$ simulations of the governing model, one simulation at the base condition and one at the perturbed condition for each parameter.

The second method to predict the sensitivity is called the tangent linear or forward method. It solves the linearized equations for small perturbations of the governing model. By using this approach, the exact derivatives of the objective function can be calculated (in the limit of infinitesimal perturbations). The cost to obtain the gradient of J_j using this method is N_a simulations of the linear model plus one simulation of the governing model.

The third method is also based on the linearized equations of the governing model, however, it derives and solves for the adjoint of the linear equations instead. Since the adjoint equation is similar to the basic flow problem (e.g., it has convection/diffusion just like the original problem, but "backward"), the computational cost of solving the adjoint is the same as the tangent equation. This method has enabled adjoint-based shape optimization, uncertainty estimation, grid-adaptation, and flow control in many CFD applications over the last 30 years [5,11]. The main reason for such success is that to compute the sensitivity gradient $\partial J_j/\partial a_i$, one only needs a single additional simulation for each J_j, regardless of the value of N_a. In other words, the adjoint method requires only N_J additional simulations to obtain N_J gradients.

Overall, when computing the sensitivity there are situations in which there are more random parameters a_i than QoIs J_j and vice-versa. The tangent linear method is more suited for situations in which there are fewer random parameters a_i than QoIs J_j ($N_a < N_J$), while the adjoint method is optimal if the opposite is true ($N_a > N_J$).

1.2.1.1 Sensitivity analysis in chaotic systems

Despite the great appeal of the adjoint method, its application to date has been for inviscid, laminar, or RANS-modeled flows. This is because the linearized and adjoint equations of time-averaged objective functions of chaotic dynamical systems diverge due to the "butterfly effect" [12].

Potential solutions to circumvent this divergence have been proposed. One of the potential methods to compute derivatives of long-time-averaged functions of chaotic dynamical systems is the least square shadowing (LSS), developed by Wang [13]. The LSS method establishes a least square optimization problem, to calculate a "tangent" solution to the original that does not exhibit exponential growth (the shadow trajectory). This solution is then used for sensitivity analysis. Blonigan *et. al.* [14] used LSS in a two-dimensional chaotic flow around an airfoil, and found that the cost to solve LSS in such a case is at least four orders of magnitude higher than that of the baseline CFD analysis. A non-intrusive least square shadowing (NILSS) method was proposed as a variant of the original LSS. The method requires only minor modifications to existing solvers to work, and the cost is proportional to the number of positive Lyapunov exponents [15]. However, it has been proven that methods based on shadow trajectories have a systematic error, which can be non-zero if the connecting map between the base and shadowing trajectory is not differentiable [16]. More recently, the space-split sensitivity or S3 method was proposed [17]. This method is another potential trajectory-based, ergodic-averaging method to differentiate statistics in chaotic systems. For a 1-D chaotic case the rate of convergence was found to be $\sqrt{n}$ where n is the number of samples.

1.2.2 Probabilistic methods for uncertainty quantification

Probabilistic methods can be separated into two sub-categories: sampling-based and non-sampling-based methods. To carry out sampling-based methods, one needs only a reliable deterministic simulation code that represents a physical model, and that can be run repetitively at different parameter values. For this reason, sampling-based methods are favored in practical applications because they allow engineers to use their existing solvers as a "black box". The most well-known probabilistic, sampling-based method is the standard Monte Carlo (MC) sampling. This method computes the expectation and the variance by performing independent and random samples of the random variables, where every parameter has been perturbed in a normally distributed fashion. Besides being non-intrusive, this method is independent of the dimensionality of the parameter space, i.e. it does not have the "curse of dimensionality". However, the convergence rate of the MC method is governed by the Central Limit Theorem, being of the order of $\sqrt{n}$ (where n is the amount of independent samples). This convergence rate makes it prohibitively expensive, requiring $\mathcal{O}(10^3)$ to $\mathcal{O}(10^4)$ independent samples to estimate the statistics of the QoI [9, 10, 18], something unaffordable in the context of LES.

In the realm of probabilistic, non-sampling-based methods, the primary method is the Galerkin Polynomial Chaos (PC) method. Based on the work of Wiener (1938), the PC method propagates the uncertainty through a model by using Galerkin projection to reformulate the uncertain variables in the governing equations onto a stochastic space spanned by a set of orthogonal multi-variate polynomials, $\Psi_i(a)$, that are functions of a random variable, a [19, 20]. That is, each uncertain variable in the model is represented using an infinite series called the polynomial chaos expansion (PCE). For example, for the Navier-Stokes equations, velocity and pressure are considered stochastic processes, and they are represented as

$$u(\mathbf{x}, t, a) = \sum_{i=0}^{p} u_i(\mathbf{x}, t)\Psi_i(a_1, a_2, ..., a_d) \tag{1.2}$$

where each u_i is deterministic and is denoted as the random mode i of the velocity. Here, p is the truncation order of the polynomials, Ψ, and d is the number of random dimensions of a. One of the main advantages of reformulating the problem using eq (1.2) is that finding the PC representation of the model output requires running the model only once. In relatively simple cases, the PC method proves to converge faster than the standard MC method [21]. Also, in some specific cases, the PC method theoretically is proven to converge exponentially [21]. However, the PC method can encounter major limitations in more complex problems. It suffers from the "curse of dimensionality"; for example, for a polynomial chaos expansion as shown in eq (1.2), the number of terms in the resulting PCE is $(d + p)!/p!d!$ [20]. In addition, its implementation requires the modification of the deterministic code, which may be inconvenient for many complex computational problems. Orszag and Bissonete [22] concluded that truncated PCE may be unsuitable for predicting the physics of high-Reynolds number flows, in particular because the nonlinearities propagate energy into the higher-order terms. Wiener-Hermite expansions fail to represent the turbulence energy cascade [23]. Consequently, the PC method is not suitable as a tool to be used in complex turbulent problems.

1.2.2.1 Stochastic expansion representation methods

There is a suite of probabilistic and sampling-based methods that circumvent the intrusivity disadvantage of the PC method. They allow for the use of already available and reliable deterministic solvers. One is the stochastic collocation (SC) method, proposed by Mathelin and Hussaini [18]. The other is the non-intrusive polynomial chaos expansion (NIPC) method developed by Hosder, Walters, and Perez [19]. Both methods construct an approximated representation of a metric of interest in the random space. The main difference between them is that, whereas the SC method forms interpolation functions for known coefficients, the NIPC method estimates coefficients for known orthogonal polynomial basis functions. In SC, collocation points are chosen in the random space, and they have associated quadrature weights or sparse grids (approaches based on random sampling are not suitable). In particular, Mathelin and Hussaini [18] chose the Gauss-Legendre points and weights and, with Lagrange interpolation, the probability distribution of the solution is constructed. In NIPC, the basis functions are obtained from the Askey family of hypergeometric orthogonal polynomials. To reduce the nonlinearity of the expansion and improve convergence, the polynomial bases are chosen such that their orthogonality weighting functions match the probability density function of the uncertain parameters, up to a constant. NIPC and SC methods provide a significant gain in efficiency over Monte-Carlo sampling for low-dimensional systems. However, as is the case with the PC method, these two methods suffer from the curse of dimensionality. The exponential rise in the number of quadrature points makes these approaches inefficient for high-dimensional problems [20]. Furthermore, it is shown that for simple design problems, the deterministic model has to be evaluated on the order of many hundreds to a couple of thousands of times, something that is unaffordable in the context of LES.

1.2.2.2 Multi-fidelity methods in uncertainty quantification

In many situations, multiple models are available to predict the same output, or QoI, with varying levels of accuracy and varying computational costs [9, 24]. An obvious solution to reduce computational cost in uncertainty quantification is to make use of multiple deterministic models with differing levels of fidelity; these are called multi-fidelity methods. A multi-fidelity method combines outputs from computationally cheap, low-fidelity models with outputs from the high-fidelity models. This combination can be a good tradeoff between cost and accuracy, leading to significant savings in CPU time and providing unbiased estimators of the statistics of the high-fidelity model output [9]. Several multi-fidelity methods have been proposed in the last few years, where the main differences lie in the low-fidelity model and what Peherstorfer, Willcox, and Gunzburger [9] call the model management strategy. Various types of low-fidelity models and model management strategies are briefly described below. A detailed description can be found in [9].

The types of low-fidelity models have been classified by Peherstorfer, Willcox, and Gunzburger into three categories: simplified low-fidelity models, projection-based low-fidelity models, and data-fit low-fidelity models. Simplified low-fidelity models, as the name suggests, are simplifications derived from the high-fidelity model, which are used by taking advantage of domain expertise and in-depth knowledge of the implementation details of the high-fidelity model. In CFD, RANS represents the low-fidelity model of DNS and LES. Coarse-grid approximations and early stopping criteria are also part of the simplified low-fidelity models. Projection-based low-fidelity models are derived from high-fidelity models by mathematically exploiting the problem structure. One common method is proper orthogonal decomposition (POD), which utilizes snapshots of the high-fidelity model to construct a basis for the low-dimensional subspace. Data-fit low-fidelity models compute an interpolation or a regression of the high-fidelity realizations. In these types of low-fidelity models, classical Lagrange polynomials can be used to derive data-fit models. Also, the "Kriging" interpolation method or Gaussian process regression is being widely used in CFD. Model management in multi-fidelity methods defines how different models are employed when using the method. Three main strategies were distinguished: adaptation, fusion, and filtering.

In short, sampling-based uncertainty quantification methods have been under continuous development, where the main objective in recent years has been to reduce the computational cost of sampling while maintaining accuracy. The multi-fidelity approach may reduce the cost to only a few hundred times the cost of a single simulation [25]. While impressive, this is still too much to be useful in practice in the context of LES.

1.3 Objective

It is apparent that the main obstacle for integrating high-fidelity models such as LES into the engineering design process is the unaffordable computational cost of its sensitivity. The sampling-based, multi-fidelity, and multi-level approach is the state-of-the-art technique, however, it is too computationally expensive for processes that require a fast turnaround. Traditional sensitivity analysis of chaotic dynamical systems diverges, disqualifying it as an option. The current state-of-the-art methods for sensitivity analysis try to find a solution to the exact adjoint problem by solving the shadow trajectory. However, the computational cost still makes them unaffordable. For that reason, the objective of this work is to perform a first assessment and feasibility analysis of a new - and computationally affordable - methodology to approximate the sensitivity of a QoI from a large eddy simulation. Since the proposed method is physics-based, it is assumed that the accuracy will be strongly dependent on the flow, the random parameters, and the QoIs.

The proposed multi-fidelity sensitivity analysis (MFSA) method represents the first attempt of estimating the sensitivity of chaotic and turbulence-resolving simulations by leveraging linearized low-fidelity models to approximate the change of the QoI due to small perturbations in the problem parameters. This study considers a relevant but geometrically simple case (flow over an airfoil) with a single random parameter (Reynolds number) and four QoIs (lift, drag, skin friction, and pressure coefficients).

Traditionally, the outcome of an LES or DNS is the prediction of the QoIs, for instance, the aerodynamic lift coefficient. The outcome of the proposed MFSA is not only the prediction of the aerodynamic coefficient, but also the sensitivity of the QoI at the extra cost of only one RANS simulation. The proposed method aims to complement, rather than replace, the fully RANS-based method, which is the state-of-the-art in industry, in situations where sensitivities with higher fidelity are required.

1.4 Assessment methodology and outline

The proposed MFSA method is presented in Chapter 2. However, in order to perform a valid assessment, one needs to compare these predictions against a benchmark. Currently, the only feasible way to compute the sensitivity of J_j from an LES of a turbulent flow is by using finite-differencing (FD) in the random space. Therefore, in this work, the sensitivity obtained using FD of two different large eddy simulations is established as the true value and used to quantify the accuracy of the MFSA's predictions. Similarly, the sensitivity of J_j is computed using FD in the random space, but using only RANS. This will allow a direct comparison against the current state-of-the-art methodology in engineering. However, before computing the finite-difference sensitivity using LES, a series of steps need to be completed, which are discussed next.

First, one needs to choose a test case. A turbulent flow over a NACA 0012 profile is chosen to assess the proposed method. This profile is chosen because of the numerous studies and data published for different angles of attack and Reynolds numbers (see, for example, [26–30]) and because it is relevant for industrial applications and engineering design. The aerodynamic drag and lift, the skin friction and pressure coefficients, $\mathcal{C}_D$, $\mathcal{C}_L$, $\mathcal{C}_f$, and $\mathcal{C}_p$, respectively, are defined as the QoIs in this work, and the uncertain design parameter is the Reynolds number, Re. Table 1.1 shows the conditions chosen. A flow past an airfoil with a Reynolds number of 4×10^5 and zero incidence is chosen to verify the in-house solver's results against other wall-resolved LES in literature; this case is named and referenced as Base Case 1 from now on. A flow past an airfoil at the same Reynolds number but at an angle of attack of $5°$ is chosen to assess the suitability of the proposed method and to compare its accuracy against the current state-of-the-art method in industry; this case is named and referenced Base Case 2 from now on. This angle of attack is chosen to ensure that both aerodynamic coefficients are larger than zero in a situation in which the flow around the airfoil remains attached. A moderate Reynolds number is chosen in order to have computationally affordable wall-resolved large eddy simulations.

Table 1.1: Cases of study

Base Case	θ	Re
1	$0°$	4×10^5
2	$5°$	4×10^5

Second, perform a verification of the solvers, which were developed in-house. Chapters 3, 4, and 6 show the mathematical formulations, numerical methods and the verification processes for the LES, RANS, and linearized RANS solvers, respectively.

Third, carry out a grid convergence study and a statistical convergence study of the time-averaged QoIs. Section 3.5 presents LES results of a turbulent flow over a NACA 0012 airfoil, including a grid convergence study using hypothesis testing, and the convergence of the statistical parameters. Section 4.4 presents RANS solutions of a turbulent flow over a NACA 0012 airfoil, including a standard grid convergence study.

Finally, the finite-difference sensitivity can be computed. Chapter 5 illustrates the computation of (what is defined as) the *true* sensitivities of the QoIs. An important question that arises when trying to compute the sensitivity using FD is what is the optimal ΔRe to use? Section 5.3 describes the procedure followed by this study to estimate the optimal perturbation.

Chapter 7 analyzes the predictions from the proposed methodology and systematically assess the possible sources of error. Finally, Chapter 8 presents the conclusions, contributions and novelties of the work, and suggestions for future work.

Chapter 2: Proposed methodology

This chapter describes the proposed multi-fidelity sensitivity analysis (MFSA) method for the special case of an incompressible flow with constant properties.

2.1 Mathematical background

Consider the time-averaged QoI, J_j

$$J_j(\mathbf{v}, \mathbf{a}) = \int_T \int_\Omega \zeta\left(\mathbf{v}, \mathbf{a}\right) d\mathcal{V} dt + \int_T \int_{\partial\Omega} \xi\left(\mathbf{v}, \mathbf{a}\right) d\mathcal{S} dt \, , \tag{2.1}$$

where ζ and ξ represent functions that operate on the volume and boundaries of the domain, $\mathbf{a}$ is a vector of design parameters $a_1, a_2, ...a_n$, and $\mathbf{v}$ is the solution vector which satisfies the governing equation

$$\mathcal{N}(\mathbf{v}; \mathbf{a}) = 0 \, . \tag{2.2}$$

In the particular case of the incompressible Navier-Stokes equations, $\mathbf{v}$ contains the velocity and pressure fields ($\mathbf{v} = \{u_i, p\}$) while $\mathcal{N}$ is the mass and momentum conservation equations.

If one perturbs (infinitesimally) the parameter vector $\mathbf{a}$ by adding $\delta\mathbf{a}$, equation (2.2) becomes

$$\mathcal{N}(\mathbf{v} + \delta\mathbf{v}; \mathbf{a} + \delta\mathbf{a}) = 0 \, , \tag{2.3}$$

where $\mathbf{v} + \delta\mathbf{v}$ is the new solution for the perturbed parameter vector $\mathbf{a} + \delta\mathbf{a}$. The objective function for the new solution is

$$J_j(\mathbf{v}+\delta\mathbf{v},\mathbf{a}+\delta\mathbf{a}) \approx J_j(\mathbf{v},\mathbf{a})+\int_T\int_\Omega \left[\frac{\partial\zeta}{\partial\mathbf{v}}\delta\mathbf{v}+\frac{\partial\zeta}{\partial\mathbf{a}}\delta\mathbf{a}\right] d\mathcal{V}dt+\int_T\int_{\partial\Omega}\left[\frac{\partial\xi}{\partial\mathbf{v}}\delta\mathbf{v}+\frac{\partial\xi}{\partial\mathbf{a}}\delta\mathbf{a}\right] d\mathcal{S}dt\,. \tag{2.4}$$

Using angular brackets $\langle\cdot,\cdot\rangle$ to denote the inner products, the sensitivity of J_j can then be defined as

$$\delta J_j = \left\langle\frac{\partial\zeta}{\partial\mathbf{v}},\delta\mathbf{v}\right\rangle_\Omega + \left\langle\frac{\partial\zeta}{\partial\mathbf{a}},\delta\mathbf{a}\right\rangle_\Omega + \left\langle\frac{\partial\xi}{\partial\mathbf{v}},\delta\mathbf{v}\right\rangle_{\partial\Omega} + \left\langle\frac{\partial\xi}{\partial\mathbf{a}},\delta\mathbf{a}\right\rangle_{\partial\Omega}, \tag{2.5}$$

which is the infinitesimal difference between the QoI of a perturbed solution compared to the base. However, in order to find δJ_j it is necessary to find the perturbed solution $\delta\mathbf{v}$. For that, the equation (2.3) may be linearized as

$$\mathcal{N}(\mathbf{v}+\delta\mathbf{v};\mathbf{a}+\delta\mathbf{a}) \approx \mathcal{N}(\mathbf{v};\mathbf{a}) + \frac{\partial\mathcal{N}}{\partial\mathbf{v}}\delta\mathbf{v} + \frac{\partial\mathcal{N}}{\partial\mathbf{a}}\delta\mathbf{a} + \mathcal{O}(\delta\mathbf{v}^2,\delta\mathbf{a}^2) = 0\,, \tag{2.6}$$

obtaining the linear system

$$\mathcal{L}\delta\mathbf{v} = f\delta\mathbf{a}, \quad \text{where} \quad \mathcal{L} \equiv \frac{\partial\mathcal{N}}{\partial\mathbf{v}}, \quad \text{and} \quad f \equiv \frac{\partial\mathcal{N}}{\partial\mathbf{a}}\,, \tag{2.7}$$

where f represents the forcing term due to the perturbation of $\mathbf{a}$. If a solution of the system proposed in equation (2.7) can be found, one can compute δJ_j using equation (2.5), and find the sensitivity

$$\frac{\partial J_j}{\partial a_i} = \lim_{\delta a_i\to 0}\frac{\delta J_j}{\delta a_i}\,. \tag{2.8}$$

Here, the computed sensitivity would be exact in the limit of an infinitesimal perturbation $\delta\mathbf{a}$.

However, it is well known that this definition of the sensitivity is not meaningful for chaotic problems since the linearized dynamics are inherently unstable, meaning that one cannot compute the perturbed solution $\delta \mathbf{v}$ in a stable way by solving eq. (2.7). The sensitivity δJ_j could instead be approximated using a finite difference in the parameter space, i.e., by solving the LES equation for a slightly perturbed parameter $\mathbf{a} + \Delta \mathbf{a}$, computing the QoI at the perturbed condition from eq. (2.1), and then approximating δJ_j using a finite difference. This is stable since the averaging is performed prior to the differentiation. It would require one LES solution for each parameter, which is computationally infeasible for large numbers of parameters.

2.2 A proposed multi-fidelity sensitivity analysis (MFSA)

The proposed methodology will be described for the case of an incompressible flow with constant properties, for which the LES governing equations are

$$\begin{aligned} &\frac{\partial u_j}{\partial x_j} = 0\,, \\ &\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} + \frac{\partial p}{\partial x_i} - \frac{\partial}{\partial x_j}\left[2\nu S_{ij} - \tau_{ij}(u_k)\right] = 0\,, \end{aligned} \tag{2.9}$$

where the full LES-resolved solution is $\mathbf{v} = \{u_i, p\}$ and $\tau_{ij}(u_k)$ is a subgrid model; the LES equation can then be denoted by $\mathcal{N}(\mathbf{v}(\mathbf{a}); \mathbf{a}) = 0$.

Many QoIs in practical applications depend only on the mean solution $\overline{\mathbf{v}}$, and not on the instantaneous solution $\mathbf{v}$. The QoI J_j could then be re-define as

$$J_j(\overline{\mathbf{v}}, \mathbf{a}) = \int_\Omega \overline{\zeta}\,(\overline{\mathbf{v}}, \mathbf{a})\, d\mathcal{V} + \int_{\partial\Omega} \overline{\xi}\,(\overline{\mathbf{v}}, \mathbf{a})\, d\mathcal{S}\,, \tag{2.10}$$

where $\overline{(\cdot)}$ is the ensemble averaged operator, $\overline{\zeta}$ and $\overline{\xi}$ are re-defined to operate on the mean solution $\overline{\mathbf{v}}$. By using eq. (2.10), one can compute a point estimate of $J_j(\overline{\mathbf{v}}(\mathbf{a}), \mathbf{a})$.

The sensitivity of this definition is

$$\delta J_j = \int_\Omega \left[\frac{\partial \overline{\zeta}}{\partial \overline{\mathbf{v}}} \delta \overline{\mathbf{v}} + \frac{\partial \overline{\zeta}}{\partial \mathbf{a}} \delta \mathbf{a} \right] d\mathcal{V} + \int_{\partial\Omega} \left[\frac{\partial \overline{\xi}}{\partial \overline{\mathbf{v}}} \delta \overline{\mathbf{v}} + \frac{\partial \overline{\xi}}{\partial \mathbf{a}} \delta \mathbf{a} \right] d\mathcal{S} \, . \tag{2.11}$$

This definition of the sensitivity is exact. As said before, one can approximate δJ_j using a finite difference in the parameter space. However, this approach has a cost of N_a, which is not affordable. Trying to compute directly the instantaneous perturbed field would require a linearization of a chaotic system which is known to diverge when integrated over a long time period [13].

The key idea of the proposed method is to compute the perturbed mean solution $\delta \overline{\mathbf{v}} = \{\delta \overline{u_i}, \delta \overline{p}\}$ using a lower-fidelity turbulence modeling approach, rather than by LES. The mean LES solution $\overline{\mathbf{v}} = \{\overline{u}_i, \overline{p}\}$ satisfies the mean governing equations

$$\begin{aligned} \frac{\partial \overline{u}_j}{\partial x_j} &= 0 \, , \\ \frac{\partial \overline{u}_i}{\partial t} + \overline{u}_j \frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{p}}{\partial x_i} - \frac{\partial}{\partial x_j} \left(2\nu \overline{S}_{ij} - \overline{\tau}_{ij} - \overline{u_i' u_j'} \right) &= 0 \, , \end{aligned} \tag{2.12}$$

to within averaging and numerical errors. This can be linearized to

$$\begin{aligned} \frac{\partial \delta \overline{u}_j}{\partial x_j} &= 0 \, , \\ \frac{\partial \delta \overline{u}_i}{\partial t} + \overline{u}_j \frac{\partial \delta \overline{u}_i}{\partial x_j} + \delta \overline{u}_j \frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \delta \overline{p}}{\partial x_i} - \frac{\partial}{\partial x_j} \left(2\nu \delta \overline{S}_{ij} - \delta \overline{\tau}_{ij} - \delta \overline{u_i' u_j'} \right) &= 0 \, . \end{aligned} \tag{2.13}$$

These equations are still exact, but suffer from a closure problem just like standard RANS equations. Specifically, the perturbed mean Reynolds and subgrid stress terms need to be modeled, which is accomplished in this work using an eddy viscosity approach.

The eddy viscosity hypothesis for the full Reynolds stress tensor is

$$\overline{u_i' u_j'} = -2\nu_t \overline{S}_{ij} + \frac{2k}{3} \delta_{ij} \, , \tag{2.14}$$

where ν_t is the eddy viscosity and $k = \overline{u'_k u'_k}/2$ is the turbulent kinetic energy. Linearization yields

$$\delta\overline{u'_i u'_j} = -2\nu_t \delta\overline{S}_{ij} - 2\overline{S}_{ij}\delta\nu_t + \frac{2\delta k}{3}\delta_{ij}\,. \tag{2.15}$$

The perturbed mean subgrid stress is absorbed into the model of the perturbed Reynolds stress (or, equivalently, we neglect it), and also neglect the perturbed turbulent kinetic energy. This requires the specification of the eddy viscosity ν_t and its perturbation $\delta\nu_t$. Using an existing RANS turbulence model for ν_t is not appealing since that would produce an eddy viscosity that is inconsistent with the LES mean solution. One of the approaches, and the one used in this thesis, is to infer ν_t by minimizing the error between the deviatoric part of the Reynolds stress tensor given by the high-fidelity model (or LES), and the deviatoric part of Reynolds stress tensor computed using the eddy viscosity assumption as (for more details see [31])

$$\nu_t = -\frac{\overline{S}_{ij}\overline{u'_i u'_j}}{2\overline{S}_{kl}\overline{S}_{kl}}. \tag{2.16}$$

This then yields equation (2.13) for infinitesimal perturbations of the mean equations as

$$\begin{aligned} \frac{\partial \delta\overline{u}_j}{\partial x_j} &= 0\,, \\ \frac{\partial \delta\overline{u}_i}{\partial t} + \overline{u}_j\frac{\partial \delta\overline{u}_i}{\partial x_j} + \delta\overline{u}_j\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \delta\overline{p}}{\partial x_i} - \frac{\partial}{\partial x_j}\left(2(\nu+\nu_t)\delta\overline{S}_{ij} + 2\overline{S}_{ij}\delta\nu_t\right) &= 0\,. \end{aligned} \tag{2.17}$$

The linearized equation (2.17) can then be denoted by

$$\mathcal{L}(\delta\overline{\mathbf{v}}, \delta\nu_t; \overline{\mathbf{v}}, \nu_t, \mathbf{a}, \delta\mathbf{a}) = 0,$$

the eddy viscosity perturbation, $\delta\nu_t$, is either assumed to be zero (the "frozen turbulence" assumption) or modeled. In this work, the eddy viscosity perturbation is modeled using Prandtl's zero-equation model

$$\nu_t = \kappa^2 d^2 \left|\overline{S}\right| D(d^+, A^+)\,, \tag{2.18}$$

where d is the wall-normal distance and D is the Van Driest damping function which depends on $d^+ = u_\tau d/\nu$ and the model constant A^+. Without going into details of this model (it will be described in Chapter 6), in the particular case where the Reynolds number is the only perturbed parameter (the case in this thesis), the eddy viscosity perturbation is given by

$$\delta\nu_{t,Pr} = \nu_t \left[\frac{\overline{S}_{ij}\delta\overline{S}_{ij}}{\overline{S}_{kl}\overline{S}_{kl}} + \frac{d\,\mathrm{sgn}(\tau_w)\exp\left[\frac{-d^+}{A^+}\right]\delta\tau_w}{\nu A^+\sqrt{D|\tau_w|}} - \frac{d^+\exp\left[\frac{-d^+}{A^+}\right]\delta\nu}{\nu A^+\sqrt{D}} \right]. \tag{2.19}$$

The proposed methodology represents a multi-fidelity method for sensitivity analysis, where the low-fidelity model (linearized-RANS in this particular case) is leveraged to reduce the computational cost of computing the sensitiviy due to small perturbations, but recoursing to the high-fidelity model to improve the accuracy of the prediction. As previously said, the tangent equation method is more suited for situations in which there are fewer random parameters a_i than QoIs J_j, ($N_a < N_J$). For example, in a flow past an airfoil in which there is interest in knowing how the uncertainty or small changes in the Reynolds numbers will affect the drag, lift, skin-friction and pressure coefficients, the forward (or tangent) approach could estimate the sensitivities of each of these QoI's with respect to Reynolds numbers by solving N_a simulations of the linear model plus one simulation of the governing model. Furthermore, since the adjoint method is derived from the linearized RANS, the proposed method represents a feasibility study for a multi-fidelity sensitivity analysis in which the adjoint could be used to compute the gradient $\partial J_j/\partial a_i$.

The proposed methodology combines two different fidelity models, LES and RANS, which may have different grid-requirements. Therefore, an interpolation would be necessary to solve the linearized problem. The interpolation is code dependent, but model constraints on the interpolated field must be enforced, e.g. the divergence-free constraint. The algorithm 1 shows the general steps of the proposed multi-fidelity sensitivity analysis.

Algorithm 1 Multi-fidelity sensitivity analysis (MFSA).

Solve the high-fidelity (LES) model at the nominal condition, $\mathcal{N}(\mathbf{v}(\mathbf{a});\mathbf{a}) = 0$, and compute ensemble average fields $\overline{\mathbf{v}}$ and $\overline{u_i'u_j'}$.

Infer the eddy viscosity field from the LES solution at the nominal condition.

if grid requirements between models differ **then**

  Interpolate the base solution onto the low-fidelity (linearized RANS) grid.

  Enforce the model-specific constraints on the interpolated solution, e.g. divergence-free and the proper order in which each variable approaches the wall.

for $i = 1$ to N_a **do**

  Solve the linearized RANS problem for the specific perturbed parameter δa_i, i.e., solve $\mathcal{L}(\delta\overline{\mathbf{v}}, \delta\nu_t; \overline{\mathbf{v}}, \nu_t, \mathbf{a}_i, \delta a_i) = 0$, which includes the calculation of the perturbation $\delta\nu_t$.

  Compute the sensitivity estimate $\partial J_j/\partial a_i$.

Chapter 3: Large eddy simulation

3.1 Mathematical formulation

This thesis deals with the three-dimensional incompressible and constant property flow of Newtonian fluids. Under these assumptions, the dynamics of such flow is described by the Navier-Stokes equations

$$\frac{\partial u_i}{\partial x_i} = 0 \quad , \tag{3.1}$$

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{\partial p}{\partial x_i} + 2\nu \frac{\partial S_{ij}}{\partial x_j} \quad , \tag{3.2}$$

where u_i is the instantaneous velocity, p is the instantaneous pressure divided by the constant density, S_{ij} is the rate of strain tensor, and ν is the kinematic viscosity. These equations, and their solutions, are formally defined in an infinite dimensional (function) space. For numerical simulations, the equations are represented discretely in a finite dimensional phase space. If solved directly, i.e. in a DNS, the full range of length scales - from the integral scale down to the Kolmogorov scale - must be resolved. The main concept of LES is to reduce the computational complexity (i.e. the number of degrees of freedom, or the dimension, of the system) by resolving only the flow dependent, large scales. This is typically accomplished by applying a low-pass filter to the Navier-Stokes equations that removes the smaller, supposedly universal, scales.

The filtered three-dimensional incompressible and density constant Navier-

Stokes equations can be derived by applying a low-pass filter with filter width $\widetilde{\Delta}$ to equations (3.1) and (3.2) (assuming that the filtering operation commutes with differentiation), resulting in the equations

$$\frac{\partial \widetilde{u}_j}{\partial x_j} = 0, \tag{3.3}$$

$$\frac{\partial \widetilde{u}_i}{\partial t} + \frac{\partial \widetilde{u_i u_j}}{\partial x_j} = -\frac{\partial \widetilde{p}}{\partial x_i} + 2\nu \frac{\partial \widetilde{S_{ij}}}{\partial x_j} \quad , \tag{3.4}$$

where $\widetilde{u}_i$ and $\widetilde{p}$ are the low-pass filtered (or resolved) velocity and pressure, respectively.

The filtering process, or the "coarse-graining"process, is in fact just a formal way of removing (or highly attenuating) the length-scales below the filter-width $\widetilde{\Delta}$. In practice, it can be done either implicitly when generating the grid or explicitly by application of a low-pass filter to the equation and then using a grid that is fine enough to solve the coarse-grained equations.

Due to the nonlinearity of the convective term in the Navier-Stokes equations, the term $\widetilde{u_i u_j}$ is not a function of the resolved velocity. As a result, equation (3.4) is not in closed form. The classical approach, and the one taken in this thesis, is to re-arrange this term by noting that

$$\widetilde{u_i u_j} = \widetilde{u}_i \widetilde{u}_j + \widetilde{u_i u_j} - \widetilde{u}_i \widetilde{u}_j \tag{3.5}$$

and, using this re-arragement, equation (3.4) becomes

$$\frac{\partial \widetilde{u}_i}{\partial t} + \frac{\partial \widetilde{u}_i \widetilde{u}_j}{\partial x_j} = -\frac{\partial \widetilde{p}}{\partial x_i} + 2\nu \frac{\partial \widetilde{S_{ij}}}{\partial x_j} - \frac{\partial \tau_{ij}}{\partial x_j} \quad , \tag{3.6}$$

where $\tau_{ij} \equiv \widetilde{u_i u_j} - \widetilde{u}_i \widetilde{u}_j$ is called the sub-grid scale (SGS) stress term. This residual stress term arises due to the lack of commutation between multiplication and filtering, and cannot be expressed in terms of the filtered (or resolved) velocity field $\widetilde{u}_i$ and, hence, it must be modeled.

3.1.1 Closure model

The following constitutive relation closure model is used in this study

$$\tau_{ij} - \frac{\tau_{kk}}{3}\delta_{ij} = -2\nu_{\mathrm{SGS}}\widetilde{S}_{ij} \quad , \tag{3.7}$$

where ν_{SGS} is the SGS eddy viscosity. The resolved rate-of-strain tensor, $\widetilde{S}_{ij}$, has zero trace and, hence, only the deviatoric (or anisotropic) part of the residual stress tensor can be modeled. The isotropic part $(\tau_{kk}/3)\delta_{ij}$ is absorbed into a modified pressure in equation (3.6). In this study, the SGS eddy viscosity model developed by Vreman [32] with a constant $c = 0.03$ is used to compute ν_{SGS}.

3.2 Numerical methods

The numerical methods used in this study are standard in the field of CFD. They are described here for the purposes of completeness. The filtered three-dimensional incompressible and density constant Navier-Stokes equations are solved in conservative form using an in-house code framework, called *Tortuga*, based on the finite-volume method. The numerical scheme used in this study for the spatial discretization of the convective and diffusive terms is fourth-order central. The solver uses a collocated grid arrangement in which the approach developed by Zang *et al.* [33] is used to ensure strong coupling between the pressure and velocity fields. The spatially discretized system is integrated in time using a fractional step method to ensure that the continuity constraint is satisfied. The time marching is a mixed implicit/explicit scheme that follows the method developed by Wray [34], where the wall normal diffusion is treated in a Crank-Nicolson like fashion and the remaining terms are treated explicitly by a three-stage Runge-Kutta scheme. The incompressible solver uses the Hypre library [35] to solve the Poisson system.

3.3 Modular verification of *Tortuga* operators

All solvers in *Tortuga* were built in a modular fashion based on the object-oriented concept. With this concept as the architectural premise, the *Tortuga* framework was built around operators that helped to construct the numerical methods discussed above. A good practice is to perform *modular verification*, in which the results for each module are compared against an analytical solution. Additionally, a grid convergence study should be carried out to check the order-of-convergence. To illustrate this approach, the gradient of a scalar operator, *gradC2Cscalar*, using second-order schemes is examined below. The gradient of a three-dimensional function, $F(x, y, z)$, is computed both analytically and numerically, and the $\mathcal{L}_2$ norm of the difference between both results is used to check the convergence. The function is defined as

$$
\begin{aligned}
F(x, y, z) &= \cos(k_X X)\cos(k_Y Y)\cos(k_Z Z)\,, \\
X(x, y) &= x\cos(\theta) + y\sin(\theta)\,, \\
Y(x, y) &= y\cos(\theta) - x\sin(\theta)\,, \\
Z(z) &= z\,,
\end{aligned}
$$

where

$$k_X = \frac{2\pi}{1.5}\ ,\ k_Y = \frac{2\pi}{1.2}\ ,\ k_Z = \frac{2\pi}{0.7}\,,$$

are the wavenumbers in the rotated X, Y, and Z directions. The analytical gradient is

$$\begin{aligned}
\frac{\partial F}{\partial x} &= -k_X \frac{\partial X}{\partial x} \cos(k_Z Z) \sin(k_X X) \cos(k_Y Y) - k_Y \frac{\partial Y}{\partial x} \cos(k_Z Z) \cos(k_X X) \sin(k_Y Y)\,, \\
\frac{\partial F}{\partial y} &= -k_X \frac{\partial X}{\partial y} \cos(k_Z Z) \sin(k_X X) \cos(k_Y Y) - k_Y \frac{\partial Y}{\partial y} \cos(k_Z Z) \cos(k_X X) \sin(k_Y Y)\,, \\
\frac{\partial F}{\partial z} &= -k_Z \sin(k_Z Z) \cos(k_X X) \cos(k_Y Y)\,.
\end{aligned}$$

Figures 3.1(a) to (c) show a comparison between the results obtained analytically and numerically for the gradient of $F(x, y, z)$ for a grid of 64 volumes in each direction. In Fig. 3.1(d) one can see that the operator has a second-order convergence. The same process was carried out to verify the whole set of operators in the *Tortuga* framework.

3.4 Verification: LES of a turbulent channel flow at $Re_\tau \approx 545$

A turbulent and incompressible flow in plane channels is used as a first verification case for the large eddy simulation solver. The flow is driven by a body force that keeps the bulk velocity constant. The channel has a width of $2H$, the streamwise length is $L_x/H = 10$, and the spanwise length is $L_z/H = 3$. The Reynolds number based on friction velocity is set at $Re_\tau \approx 545$. A grid convergence study is carried out to verify that the solution converges to the benchmark DNS of Del Alamo and Jimenez [1]. Table 3.1 shows the mesh resolution at the wall in viscous units for three different grids. Figure 3.2 compares *Tortuga* and benchmark results. One can see that as one refines the grid, the results converge to the DNS solution.

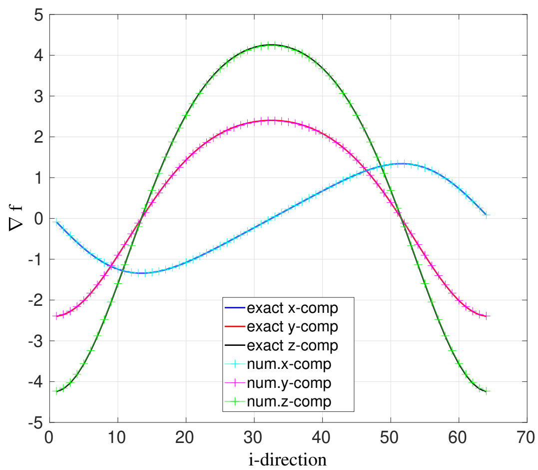

(a) Gradient of $F(x, y, z)$ in i direction

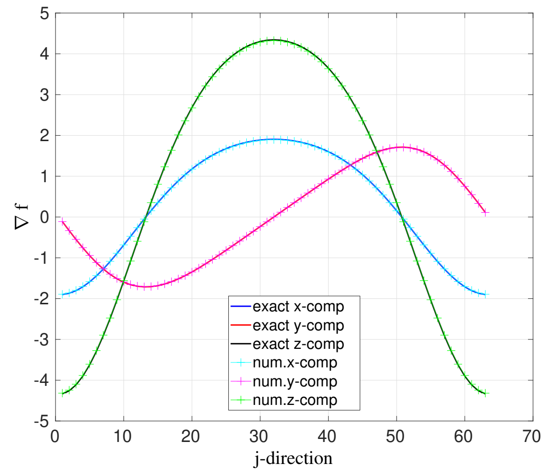

(b) Gradient of $F(x, y, z)$ in j direction

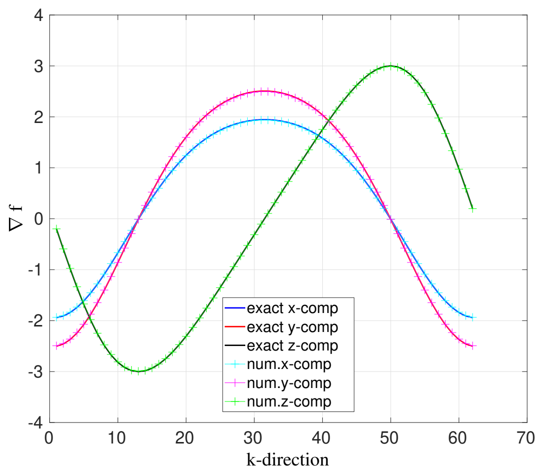

(c) Gradient of $F(x, y, z)$ in k direction

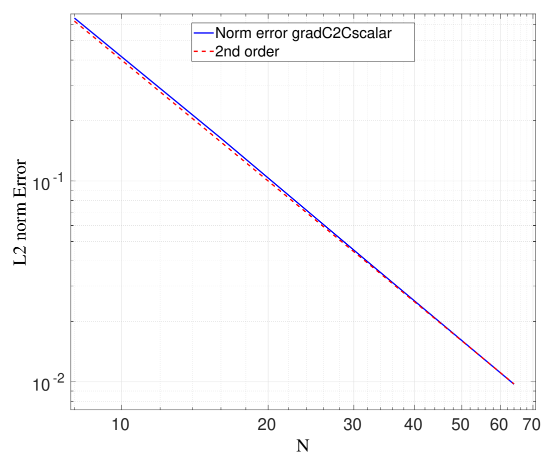

(d) Grid convergence

Figure 3.1: Verification of the gradient of a scalar operator.

3.5 Large eddy simulation of the turbulent flow past a NACA 0012 profile

The turbulent flow over a NACA 0012 airfoil is chosen to perform the feasibility study of the proposed multi-fidelity sensitivity analysis (MFSA) method. This flow was chosen because it is a relevant but geometrically simple case to analyze. There is sufficient data of this profile to which the in-house solver can be verified. Also, this case offers a transparent and straigthforward way to perform sensitivity analysis of the aerodynamic drag, lift, and the skin friction and pressure coefficients

Table 3.1: Grid convergence of a large eddy simulation of a channel flow at $Re_\tau \approx$ 545.

Grid	(dx^+, dy^+, dz^+)
$gch1$= 102×44×56	(51.06, 4.53, 27.9)
$gch2$= 121×50×88	(44.59, 3.39, 18.39)
$gch3$= 216×72×176	(25.01, 3.22, 9.20)

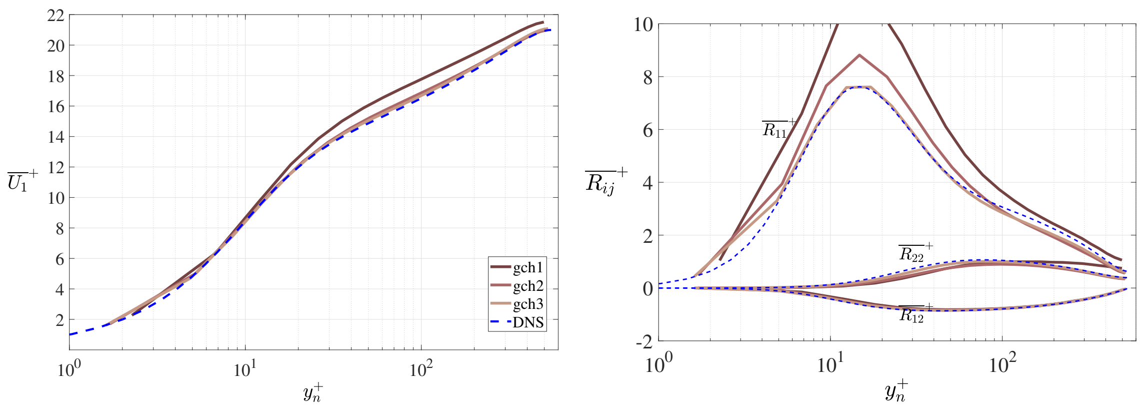

(a) (b)

Figure 3.2: Large eddy simulation of a channel flow at $Re_\tau \approx 545$; (a) convergence of the mean-velocity in inner-scaled units for grids in Table 3.1; (b) convergence of Reynolds stress tensor components in inner-scaled units for grids in Table 3.1. The dotted blue lines correspond to the DNS solution of Del Alamo and Jimenez [1] for the same setup. The other colors from brightest to darkest correspond to three grids.

due to changes on a design parameter, like the Reynolds number. Table 3.2 shows the conditions chosen. Base Case 1 is chosen to verify the in-house solver's results against other wall-resolved LES in literature (see section 3.5.4). Base Case 2 is chosen to assess the suitability of the proposed method and to compare its accuracy against the fully RANS-based method, which is the current state-of-the-art in industry, in a situation in which the flow around the airfoil remains attached. A moderate Reynolds number is chosen in order to have computationally affordable wall-resolved large eddy simulations.

3.5.1 Computational details

The present study uses a symmetric body-fitted O-grid block to simulate a turbulent flow past an airfoil. All LES cases shown in this study use a computational domain of $L_\eta = 100c$ (Fig. 3.3(a)), a spanwise width of $L_z = 0.1c$ for Base Cases 1 and 2. The NACA 0012 trailing edge is rounded as shown in Fig. 3.3(b), and unless specified, the rounded trailing edge is the default condition when running a turbulent flow past an airfoil. To ensure transition, the boundary layers are tripped on both sides of the airfoil using steady suction over the region $0.05 < x/c < 0.075$, and blowing over the region $0.075 < x/c < 0.1$, as shown in Fig. 3.4(a). The magnitude of suction and blowing is constant, with $|U_{blowing}| = |U_{suction}| = 0.03U_\infty$. The tripping is applied in the spanwise direction over the regions $0.01 < z/c < 0.04$ and $0.05 < z/c < 0.09$. Other configurations using uniform tripping in the spanwise direction were tested. However, partial span tripping demonstrated better results of boundary layer transition compared to full span tripping. A similar mechanism with the same purpose was used in the study of Wolf *et al.* [36]. Figure 3.4(b) shows the instantaneous visualization of the flow past a NACA 0012 using the referenced tripping mechanism. One can see the coherent vortical structures identified using the Q-criterion defined as $Q = \frac{1}{2}(\Omega_{ij}\Omega_{ij} - S_{ij}S_{ij})$, where S_{ij} and Ω_{ij} are the symmetric and anti-symmetric part of the velocity gradient, respectively.

Table 3.2: Cases of study

Base Case	θ	Re
1	$0°$	4×10^5
2	$5°$	4×10^5

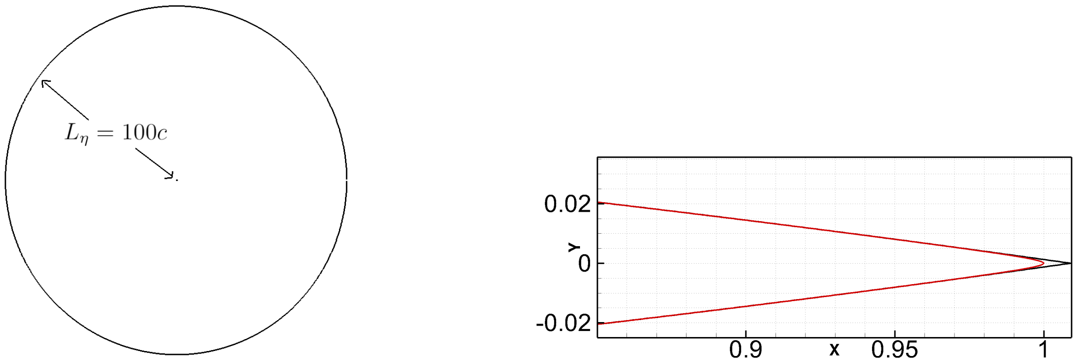

(a) (b)

Figure 3.3: (a) Computational domain, c is the airfoil chord; (b) rounded NACA 0012 trailing edge; The black solid line is the standard NACA 0012 profile, the red solid line is the rounded NACA 0012 profile used throughout this study.

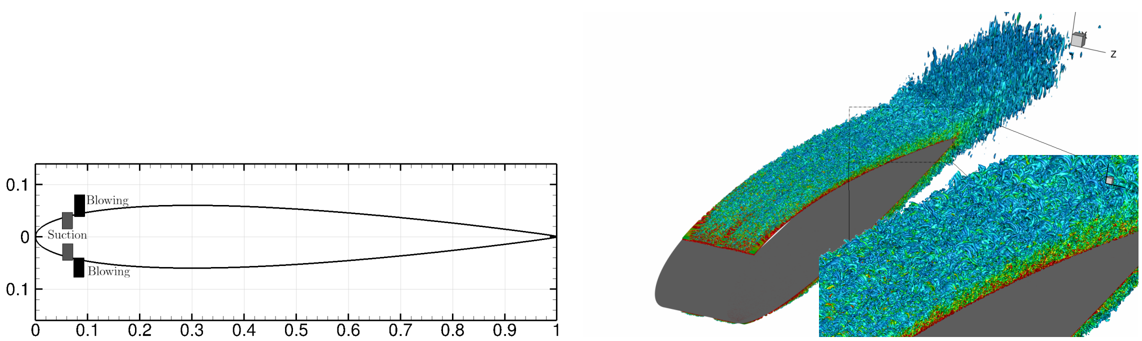

(a) (b)

Figure 3.4: (a) Sketch of the boundary layer tripping; (b) instantaneous visualization of the NACA 0012 case showing coherent vortical structures identified through the Q-criterion colored by vorticity magnitude.

3.5.2 Statistical convergence of the QoIs and computation of confidence intervals using the batch method

The aerodynamic drag and lift coefficients, $\mathcal{C}_D$, and $\mathcal{C}_L$, are functions of the instantaneous flow field when predicted using large eddy simulations. However, the proposed method makes use of ensemble average fields rather than instantaneous ones. Computing ensemble averages for $\mathcal{C}_D$ and $\mathcal{C}_L$ would require solving several LES, but, since the study cases under consideration are assumed to be ergodic dynamical systems, the average of the time series of the coefficients is assumed to be equivalent to the ensemble average. In this work, the time signals of $\mathcal{C}_D$ and $\mathcal{C}_L$ are treated as random variables under the framework of probability theory, which provides the sufficient theoretical background to compute the statistical parameters. In particular, this work is interested in the sample mean, the standard error and the confidence interval for a QoI.

The LES outputs (of interest for this study) are continuous time signals of $\mathcal{C}_D$ and $\mathcal{C}_L$. However, to properly obtain expected values and confidence intervals from these continuous time signals, a series of steps need to be followed. In this section, the procedure for calculating statistical parameters and confidence intervals, in general, is described for a random variable and for the difference of two random variables. Next, the statistics for Base Cases 1 and 2 are shown.

3.5.2.1 Procedure to compute statistics and confidence intervals using the batch method

Consider two time-varying quantities $\mathcal{X}_1(t)$ and $\mathcal{X}_2(t)$ whose time series are shown in Fig. 3.5. These two signals were artificially created using a random function with the same mean and standard deviation, $\mu_{\mathcal{X}_1} = \mu_{\mathcal{X}_2} = 0$ and $S_{\mathcal{X}_1} = S_{\mathcal{X}_2} = 0.1$. However, $\mathcal{X}_1$ and $\mathcal{X}_2$ were created with different autocorrelation coefficient values. Specifically, the autocorrelation at time lag $\tau = 1$ for $\mathcal{X}_1$ is 0.01 and for $\mathcal{X}_2$ is 0.9. In other words, two samples of $\mathcal{X}_1$ taken at $\tau = 1$ time units apart are independent of each other, whereas two samples of $\mathcal{X}_2$ taken at $\tau = 1$ time units apart show a strong correlation.

In real-life applications, it is common to have a sample, or time series, of a QoI whose population distribution is unknown. Luckily, under the framework of probability theory, one can calculate the statistics of a sample with unknown distribution using the central limit theorem (CLT). The CLT states that when independent random variables are added, their properly normalized sum tends toward a normal

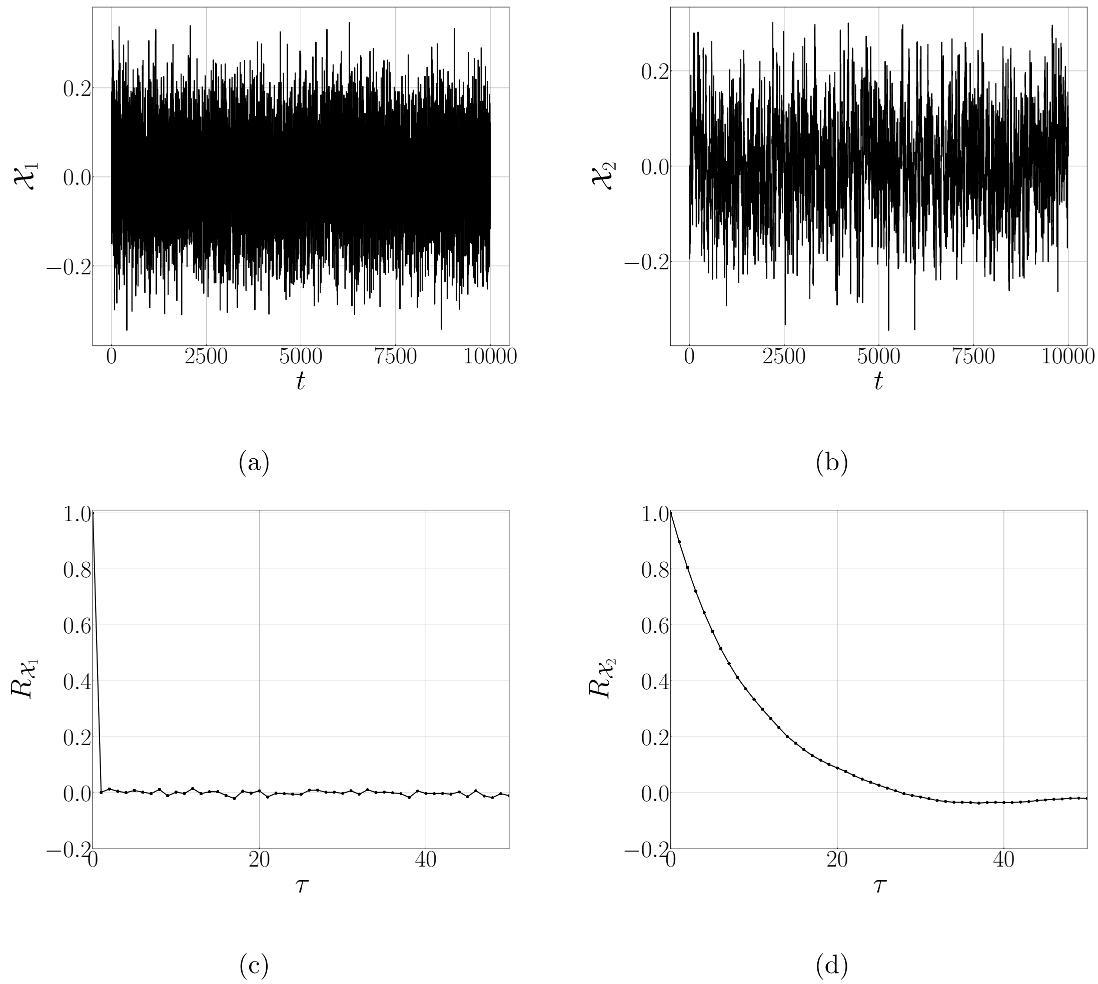

Figure 3.5: Sample time signals (top) with their auto-correlation coefficients (bottom).

distribution even if the original variables themselves are not normally distributed. In the example from above, since the distributions of $\mathcal{X}_1$ and $\mathcal{X}_2$ are unknown, the CLT can be leveraged to compute the statistical parameters needed. The computation of confidence intervals requires the standard error to be estimated, i.e., the standard deviation of the sample mean. This estimation depends on the degree of autocorrelation in the signal. The autocorrelation coefficient is defined as

$$R_{\mathcal{X}}(\tau) = \frac{\frac{1}{T-\tau}\int_0^{T-\tau} \mathcal{X}'(t)\mathcal{X}'(t+\tau)dt}{\frac{1}{T}\int_0^T \mathcal{X}'^2(t)dt} \quad , \tag{3.8}$$

where $\mathcal{X}'(t)$ is the signal fluctuation (zero mean), T is the length of the signal, τ

is the lag time. Figure 3.5 shows the autocorrelation coefficient for $\mathcal{X}_1$ and $\mathcal{X}_2$. As mentioned previously, $\mathcal{X}_1$ is sufficiently independent at $\tau = 1$ while $\mathcal{X}_2$ has an autocorrelation of $R_{\mathcal{X}_2}(1) = 0.9$ and only shows sufficient independence at $\tau \geq 20$. If one were to use the law of large numbers (i.e., that the standard error is the standard deviation divided by the square-root of the number of samples) for the signal $\mathcal{X}_1$ one could use every single observation as an independent sample. However, to apply the law of large numbers for signal $\mathcal{X}_2$, it would be necessary to use observations that are separated by at least 20 units. Observations from experiments or LES look more like the signal $\mathcal{X}_2$, in which there is a high degree of autocorrelation.

To facilitate the computation of confidence intervals for correlated signals, the present work makes use of a simple batch method. First, the whole time signal is divided into a number of batches n_b. For each batch, the mean is computed as

$$\widehat{\mathcal{X}}_i = \frac{1}{n_i}\sum_{j=1}^{n_i} \mathcal{X}_j, \quad \text{for} \quad i = 1, 2, 3, \ldots, n_b\,, \tag{3.9}$$

where n_i is number of samples in batch i. One can also define the time per batch as $t_b = n_i \Delta t$, where Δt is the signal time step. Next, the sample mean is given by the mean of the batches' mean and is computed as

$$\overline{\mathcal{X}} = \frac{1}{n_b}\sum_{i=1}^{n_b} \widehat{\mathcal{X}}_i. \tag{3.10}$$

Knowing the sample mean, one can proceed to compute the sample variance of the batches' mean as

$$S^2_{\widehat{\mathcal{X}}} = \frac{1}{n_b - 1}\sum_{i=1}^{n_b} \left(\widehat{\mathcal{X}}_i - \overline{\mathcal{X}}\right)^2 \tag{3.11}$$

and finally the variance of the sample mean (the square of the standard error) is approximated as the sample variance of the batch means divided by the number of

Table 3.3: Statistical parameters for $\mathcal{X}_1$ and $\mathcal{X}_2$.

t_b	$\overline{\mathcal{X}_1}$	$S_{\overline{\mathcal{X}_1}}$	$CI_{\overline{\mathcal{X}_1}}$	$\overline{\mathcal{X}_2}$	$S_{\overline{\mathcal{X}_2}}$	$CI_{\overline{\mathcal{X}_2}}$
2	-8.6×10^{-4}	9.84×10^{-4}	1.92×10^{-3}	4.7×10^{-3}	1.35×10^{-3}	2.65×10^{-3}
50	-8.6×10^{-4}	9.76×10^{-4}	1.92×10^{-3}	4.7×10^{-3}	3.74×10^{-3}	7.38×10^{-3}
80	-8.6×10^{-4}	9.98×10^{-4}	1.97×10^{-3}	4.7×10^{-3}	3.81×10^{-3}	7.54×10^{-3}

batches

$$S^2_{\overline{\mathcal{X}}} \approx \frac{S^2_{\widehat{\mathcal{X}}}}{n_b} = \frac{1}{n_b}\frac{1}{n_b - 1}\sum_{i=1}^{n_b}\left(\widehat{\mathcal{X}}_i - \overline{\mathcal{X}}\right)^2 , \tag{3.12}$$

in which it is assumed that the batch means are uncorrelated. With the standard error one can compute the $1-\alpha$ percent confidence interval of the sample mean as

$$CI_{\overline{\mathcal{X}}} = t_{\alpha/2,df} S_{\overline{\mathcal{X}}}. \tag{3.13}$$

where α is the significance level and $df = n_b - 1$ are the degrees of freedom. Figure 3.6 shows the estimated standard error as a function of the time per batch t_b for signals $\mathcal{X}_1$ and $\mathcal{X}_2$. One can see that for $\mathcal{X}_1$, the estimated standard error is approximately $S_{\overline{\mathcal{X}_1}} \approx 1.0 \times 10^{-3}$ for batch sizes between $2 < t_b < 300$. Therefore, the estimated confidence interval is also constant, with numbers listed in Table 3.3. On the other hand, for signal $\mathcal{X}_2$, the estimated standard error decreases as t_b decreases, because the batch means become increasingly correlated, which leads to an underprediction of the standard error. An example of this is that the estimated confidence intervals in Table 3.3 do not include the true mean for the smaller batch sizes. However, there is a region between $40 < t_b < 300$ where the estimated standard error is approximately constant at $S_{\overline{\mathcal{X}_2}} \approx 4.0 \times 10^{-3}$. In both signals, when t_b increases beyond the plateau region, the standard error increases and behaves erratically. The reason for this behavior is that, when t_b is sufficiently large, too few independent samples are used, leading to large random noise in the computed $S_{\overline{\mathcal{X}}}$.

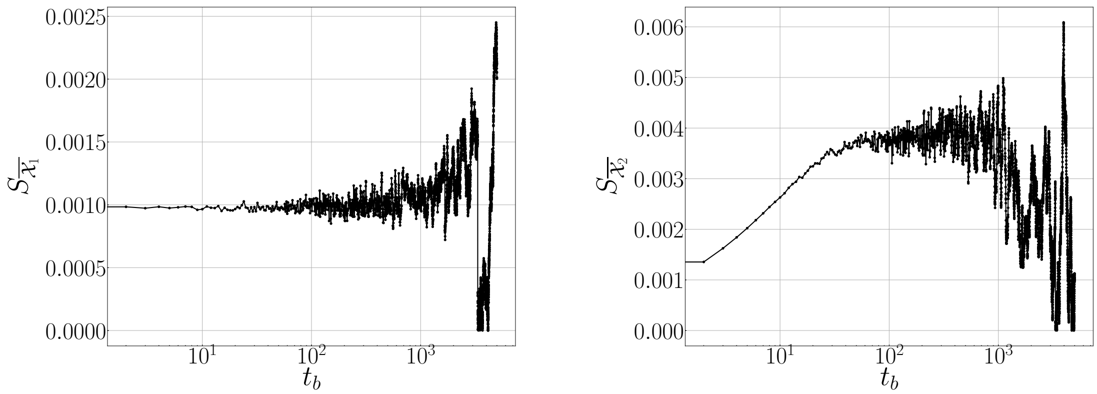

(a) (b)

Figure 3.6: Standard error for (a) signal $\mathcal{X}_1$ and (b) signal $\mathcal{X}_2$ as a function of the time per batch t_b.

3.5.2.2 Required sampling time to compute the difference of two random variables

In this work, there are two situations in which the difference of two random variables is needed. During the grid convergence study using hypothesis testing, and when estimating the true sensitivity using finite-differencing. Thus, this work makes use of statistical inference to draw conclusions about the difference between two independent random variables. Given two independent random variables, $\mathcal{X}_1$ and $\mathcal{X}_2$, with sample means $\overline{\mathcal{X}_1}$ and $\overline{\mathcal{X}_2}$, the standard error for the difference of these two sample means is

$$S_{\overline{\mathcal{X}_1}-\overline{\mathcal{X}_2}} = \sqrt{S^2_{\overline{\mathcal{X}_1}} + S^2_{\overline{\mathcal{X}_2}}} \quad , \tag{3.14}$$

and the $1-\alpha$ confidence interval is

$$CI_{\overline{\mathcal{X}_1}-\overline{\mathcal{X}_2}} = t_{\alpha/2,df} S_{\overline{\mathcal{X}_1}-\overline{\mathcal{X}_2}} \quad , \tag{3.15}$$

where df is the degrees of freedom given by

$$df = \frac{\left(S^2_{\overline{\mathcal{X}_1}} + S^2_{\overline{\mathcal{X}_2}}\right)^2}{\frac{S^4_{\overline{\mathcal{X}_1}}}{n_{b,1}-1} + \frac{S^4_{\overline{\mathcal{X}_2}}}{n_{b,2}-1}} \quad , \tag{3.16}$$

where $n_{b,1}$ and $n_{b,2}$ are the sizes of each sample.

A common question that arises when comparing two independent variables is for how long should the variables be sampled to obtain a desired value of the (1-α) CI for the sample mean? Or, from an engineering perpective, for how long should the experiments or the simulations be run to obtain a CI that is, at most, say 10 percent of the expected value?

To illustrate the problem at hand, consider three different and randomly created time series, $\mathcal{X}_1$, $\mathcal{X}_2$, and $\mathcal{X}_3$. The mean of each signal is $\mu_{\mathcal{X}_1} = 0.1$, $\mu_{\mathcal{X}_2} = 0.11$, and $\mu_{\mathcal{X}_3} = 0.101$, respectively. The standard deviation for the signal $\mathcal{X}_1$ is $S_{\mathcal{X}_1} = 0.01$. The standard deviations for signals $\mathcal{X}_2$, and $\mathcal{X}_3$ are the same at $S_{\mathcal{X}_2} = S_{\mathcal{X}_3} = 0.02$. All signals were created with an autocorrelation coefficient at time lag $\tau = 1$ of 0.3. Now, for how long should these signals be sampled to obtain a 95 percent CI that is at most 10 percent of the expected value?

Table 3.4 shows the confidence intervals for the difference of two random variables for different number of samples using equation (3.15). One can see that, first, the statistics for the difference between signals $\mathcal{X}_1$ and $\mathcal{X}_3$ using a sample length of 1000 show no convergence. In other words, the lower bound of the 95 percent CI for the difference is $\overline{\mathcal{X}_3} - \overline{\mathcal{X}_1} - CI_{\overline{\mathcal{X}_1}-\overline{\mathcal{X}_2}} = 1.06 \times 10^{-3}$, which is higher than the difference of the true means, $\mu_{\mathcal{X}_3} - \mu_{\mathcal{X}_1} = 1 \times 10^{-3}$. Second, a sample length of 5000 is long enough to give a 95 percent CI below 10 percent of the expected value for the difference between signals $\mathcal{X}_1$ and $\mathcal{X}_2$. However, to obtain a 95 percent CI lower than 10 percent of the expected value for the difference between signals $\mathcal{X}_1$ and $\mathcal{X}_3$, the sample length should be 1×10^5. Overall, one can see that the smaller the differ-

Table 3.4: Confidence intervals for the difference of two random variables for different number of samples. True means are $\mu_{\mathcal{X}_2} - \mu_{\mathcal{X}_1} = 0.01$, and $\mu_{\mathcal{X}_3} - \mu_{\mathcal{X}_1} = 1 \times 10^{-3}$.

No. of samples	$\overline{\mathcal{X}_2} - \overline{\mathcal{X}_1}$	$t_{\alpha/2,df}\sqrt{S^2_{\overline{\mathcal{X}_1}} + S^2_{\overline{\mathcal{X}_2}}}$	$\overline{\mathcal{X}_3} - \overline{\mathcal{X}_1}$	$t_{\alpha/2,df}\sqrt{S^2_{\overline{\mathcal{X}_1}} + S^2_{\overline{\mathcal{X}_3}}}$
1×10^3	0.0111	1.77×10^{-3}	2.75×10^{-3}	1.69×10^{-3}
5×10^3	0.0100	7.97×10^{-4}	8.00×10^{-4}	7.89×10^{-4}
1×10^5	—	–	1×10^{-3}	1.2×10^{-4}

ence of the means is, the larger the sample needed to obtain statistical convergence. This simple analysis will be useful when computing the sensitivity gradient using finite-differencing in chapter 5.

3.5.2.3 Statistical convergence of the aerodynamic drag and lift coefficients for the Base Cases

Figure 3.7 presents the statistics of the aerodynamic coefficients for Base Cases 1 and 2. The dotted curves are the instantaneous aerodynamic coefficients, $\mathcal{C}_D$ and $\mathcal{C}_L$, the black square dots represent the mean for each batch, $\widehat{\mathcal{C}_D}$ and $\widehat{\mathcal{C}_L}$, their running sample mean, $\overline{\mathcal{C}_D}$ and $\overline{\mathcal{C}_L}$ are shown as the solid thick lines, and their 95 percent confidence interval are plotted as the shadow region. These three quantities are shown as functions of the non-dimensionalized time $t^* = tU_\infty/c$. These results were obtained using the grid g_{LES3} which is the final grid obtained in the grid sufficiency study done in section 3.5.3.

Figure 3.8(a) shows the autocorrelation coefficient for Base Case 1. Note that for a delayed time of $\tau > 0.2$, both signals are weakly or not correlated, meaning that to guarantee independent samples, the signals must be divided into batches of length $t^*_b > 0.2$. This guarantees that each batch is an independent sample. Figure 3.8(b) shows the standard error for $\mathcal{C}_L$ as a function of the time per batch t^*_b. As shown in the figure, there is a range of batch sizes $0.15 < t^*_b < 0.8$ where the standard error is approximately constant, which means that batch sizes within this range produce a reasonable estimate of the confidence interval. Some sample

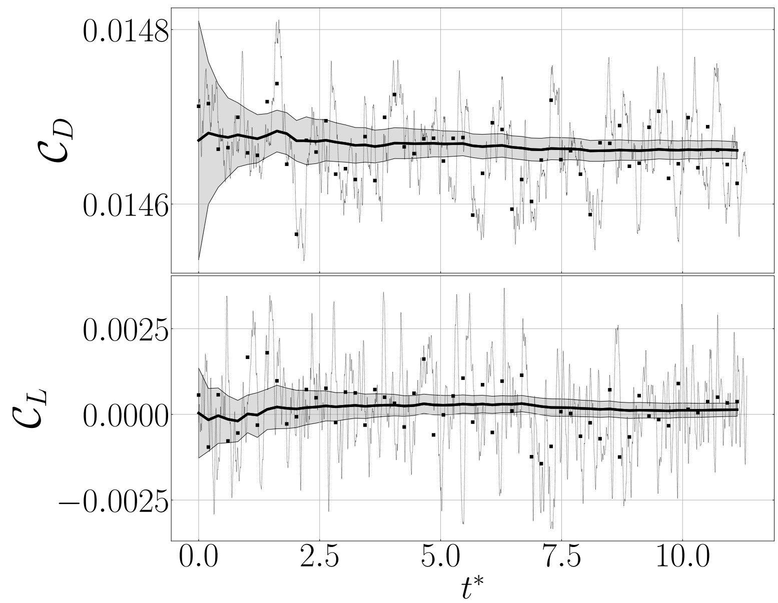

(a) Base Case 1

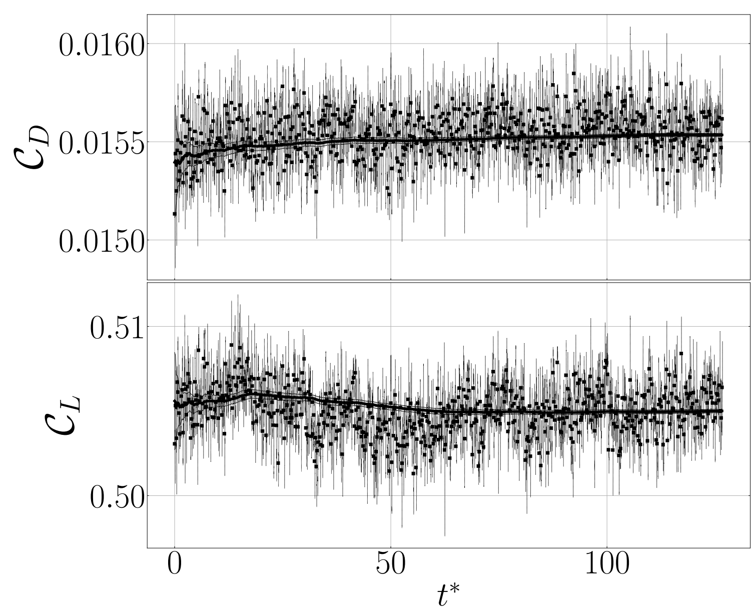

(b) Base Case 2

Figure 3.7: $\mathcal{C}_D$ and $\mathcal{C}_L$ vs t^* for (a) Base Case 1 and (b) Base Case 2. The dash-dotted lines are the instantaneous aerodynamic coefficients, $\mathcal{C}_D$ and $\mathcal{C}_L$, the black square dots represent the mean for each batch, $\widehat{\mathcal{C}_D}$ and $\widehat{\mathcal{C}_L}$, their running sample mean, $\overline{\mathcal{C}_D}$ and $\overline{\mathcal{C}_L}$ are shown as the solid thick lines, and their 95 percent confidence interval are plotted as the gray shadow region.

numbers are listed in Table 3.5. As a final sanity check for this particular case of zero incidence, the 95 percent confidence interval for the lift coefficient includes its

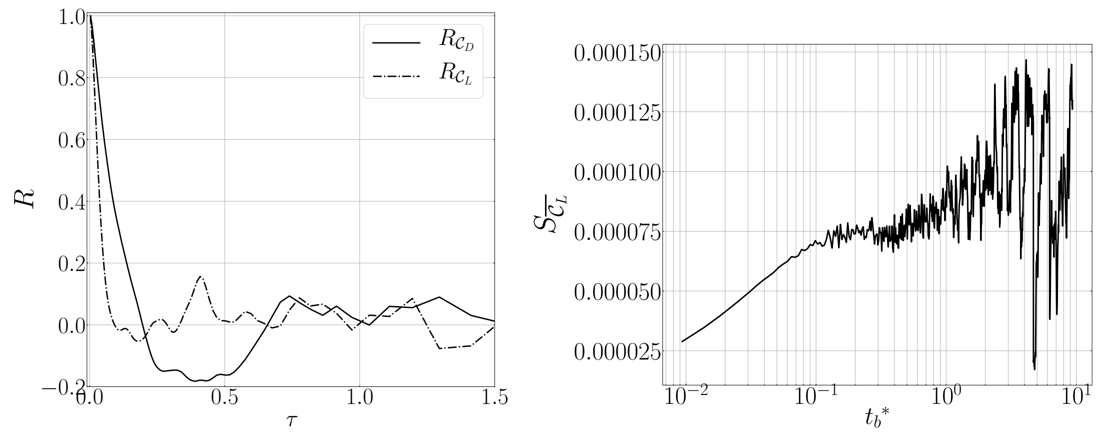

(a) (b)

Figure 3.8: (a) Autocorrelation coefficient for $\mathcal{C}_D$ and $\mathcal{C}_L$; (b) standard error for $\mathcal{C}_L$ vs t_b^* for Base Case 1.

Table 3.5: Statistics of $\mathcal{C}_D$ and $\mathcal{C}_L$ for Base Case 1.

n_b	t_b^*	$\overline{\mathcal{C}_D}$	$CI_{\overline{\mathcal{C}_D}}(95\%)$	$\overline{\mathcal{C}_L}$	$CI_{\overline{\mathcal{C}_L}}$ (95%)
60	0.20	0.0146615	9.60×10^{-6}	1.366×10^{-4}	1.88×10^{-4}
27	0.45	0.0146615	9.53×10^{-6}	1.366×10^{-4}	1.84×10^{-4}

Table 3.6: Statistics of $\mathcal{C}_D$ and $\mathcal{C}_L$ for Base Case 2.

n_b	$\overline{\mathcal{C}_D}$	$S_{\widehat{\mathcal{C}_D}}$	$S_{\overline{\mathcal{C}_D}}$	$CI_{\overline{\mathcal{C}_D}}$
655	0.0155349	1.06×10^{-4}	4.23×10^{-6}	8.31×10^{-6}
n_b	$\overline{\mathcal{C}_L}$	$S_{\widehat{\mathcal{C}_L}}$	$S_{\overline{\mathcal{C}_L}}$	$CI_{\overline{\mathcal{C}_L}}$
655	0.5050077	1.35×10^{-3}	5.35×10^{-5}	1.05×10^{-4}

true mean ($\mu_{\mathcal{C}_L} = 0.0$).

Figure 3.9(a) shows the autocorrelation coefficient for the signals from Base Case 2. For a time delay of $\tau > 0.15$ both signals are weakly- or un-correlated. Figure 3.9(b) shows the standard error for $\mathcal{C}_D$ as a function of the time per batch t_b^*. There is a region between $0.16 < t_b^* < 0.7$ where the standard error is approximately constant. Thus it follows that statistics can be drawn from within the region $0.2 < t_b^* < 0.7$ to guarantee uncorrelated and independent batch means for $\mathcal{C}_D$ (a similar region was found $\mathcal{C}_L$). Table 3.6 presents the statistics drawn within this region for $\mathcal{C}_D$ and $\mathcal{C}_L$.

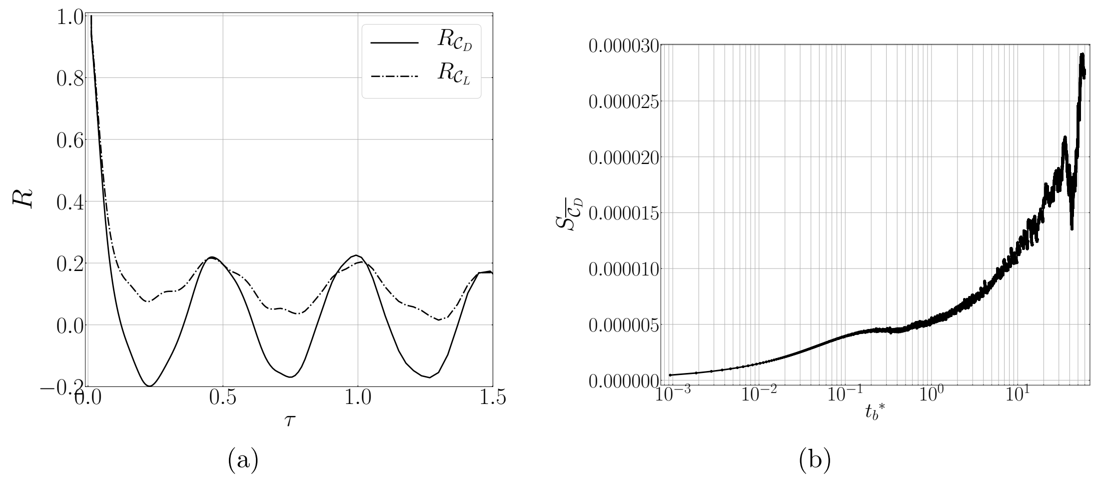

Figure 3.9: (a) Autocorrelation coefficient for $\mathcal{C}_D$ and $\mathcal{C}_L$; (b) standard error for $\mathcal{C}_D$ vs t_b for Base Case 2.

3.5.3 Grid sufficiency study using hypothesis testing

Large eddy simulation (LES) is clearly dependent on the computational grid and/or the imposed coarse-graining length scale ("filter width"). Thus meaningful predictions can only be made after having established that the quantities of interest (QoIs) are not affected too strongly by the chosen grid or length scale. Since an LES will never converge in the traditional numerical sense (i.e., to a specific solution at every point in space and time), it is necessary to discuss convergence and grid sufficiency only in the context of specific QoIs (cf. [37]). A meaningful definition of grid sufficiency is then that all QoIs must change by less than some acceptable tolerance between two sufficiently different grids. The chaotic nature of an LES then implies that the assessment of grid sufficiency necessarily centers around the difference between two imperfectly computed sample means or variances, and therefore should be viewed as a probabilistic hypothesis test.

The purpose of this section is three-fold: to make the case that the question of grid sufficiency in LES should be viewed as a hypothesis test, to demonstrate how this can be done using standard tools from statistics, and to illustrate how this leads to transparent ways to both judge grid sufficiency and to make decisions about how much averaging time is required to make that judgment.

3.5.3.1 Problem set-up and data processing

Consider again the time-varying quantity $\mathcal{X}(t)$, and that $\mathcal{X}_i(t)$ denotes this quantity computed from the solution on grid i. In this study, the quantity-of-interest is the expected value of $\mathcal{X}(t)$. In practice, we compute the sample mean $\overline{\mathcal{X}}_i$ over an averaging time of T_i after discarding an initial transient. Methods for determining the length of the initial transient were proposed and assessed in the context of turbulence simulations by, among others, [38] and [39]; the focus of the present analysis is entirely on the signal after the initial transient has been discarded. To facilitate a statistical treatment, the standard error of the mean, i.e., the standard deviation of the sample mean is computed. This could be done in multiple ways, for example with an auto-regressive model [40, 41] or the more recently proposed method by [42]. In the present work, we use the simple "non-overlapping batch method" which is described in section 3.5.2.1. The standard error for the difference of two sample means is computed following the process described in section 3.5.2.2.

3.5.3.2 Judging grid sufficiency

A practically meaningful definition of grid sufficiency is that a quantity-of-interest (QoI) must differ by less than some user-defined tolerance between two grids that are sufficiently different in terms of resolution. Mathematically, this is a hypothesis test that can be stated as (for a QoI that is an expected value)

$$H_0 : \; |\mu_{\mathcal{X}_A} - \mu_{\mathcal{X}_B}| \geq \Delta_0 \;\; (\text{insufficient grid}) \; ,$$

$$H_1 : \; |\mu_{\mathcal{X}_A} - \mu_{\mathcal{X}_B}| < \Delta_0 \;\; (\text{sufficient grid}) \, ,$$

where Δ_0 is a tolerance that is problem- and context-specific. The P-value of this test is

$$\text{P-value} = 2\,P\left(Z < \frac{|\overline{\mathcal{X}}_A - \overline{\mathcal{X}}_B| - \Delta_0}{S_{\overline{\mathcal{X}}_A - \overline{\mathcal{X}}_B}}\right), \tag{3.17}$$

with the probability evaluated from the normal distribution (or, alternatively, from a Student's T-distribution). The outcome of the test is either to declare the grid sufficient if the P-value is less than some specified α or, if not, to declare the test inconclusive and thus be forced to refine the grid. Note that α is the user-specified allowed probability of a Type I error. Also note that the criterion of a P-value $< \alpha$ can equivalently be stated as $|\overline{\mathcal{X}}_A - \overline{\mathcal{X}}_B| < \Delta_0 - z_{\alpha/2}\,S_{\overline{\mathcal{X}}_A - \overline{\mathcal{X}}_B}$ with the critical $z_{\alpha/2}$ computed from the normal distribution. In other words, the difference between the QoIs computed on the different grids must be less than Δ_0 by some margin determined by the standard error and the required confidence level in order to declare the grid sufficient.

Given the high computational cost of running on a refined grid, it is important that the hypothesis test declares the grid sufficient when it really is. Failure to do so is a Type II error, the probability β of which can be quantified only for a specific statement about the grid sufficiency. For example, we may require a specific value of β when the true difference of means is some fraction γ of the allowed tolerance, i.e., when $|\mu_{\mathcal{X}_A} - \mu_{\mathcal{X}_B}| = \gamma\Delta_0$. The probability of a Type II error is then

$$\beta = P\left(Z < z_{\alpha/2} - \frac{(1-\gamma)\Delta_0}{S_{\overline{\mathcal{X}}_A - \overline{\mathcal{X}}_B}}\right). \tag{3.18}$$

The main value in stating the question of grid sufficiency as a hypothesis test is the explicit acknowledgment that the conclusion carries uncertainty. Aside from being valuable in itself, this also allows us to transparently encode our tolerance for different types of errors. A Type I error implies erroneously trusting the results on an insufficient grid; this is a serious error, and one should thus assign a low value to α. A Type II error implies unnecessarily running on a further refined grid; this costs additional computational resources but is otherwise harmless, and one should thus assign a relatively larger value to the required β_{req}.

The process is then to first choose values for the parameters Δ_0, α, β_{req}, and γ, and then to run simulations on grids A and B until either the P-value $< \alpha$ (and thus stop and declare the grid sufficient) or $\beta < \beta_{\text{req}}$ (and thus stop, declare the grid insufficient, and proceed to create a finer grid).

3.5.3.3 Optimal averaging times

An interesting benefit of the hypothesis testing formulation outlined above is that it allows for the optimal averaging times on the different grids to be estimated. While many LES studies in practice have used the same averaging time on the different grids, this is in fact not the optimal use of resources.

Suppose that we average the two simulations for $\hat{T}_A$ and $\hat{T}_B$ units of time. The standard error of the mean in a single simulation is then $\sim \tau_{\text{decorr}}\, \sigma^2/\hat{T}$, where τ_{decorr} is a decorrelation time and σ^2 is the variance of the signal. We could then estimate the product $\tau_{\text{decorr}}\, \sigma^2$ (which we label ψ^2 here) as $T_i\, S^2_{\overline{\mathcal{X}}_i}$ from prior results. Assuming that the decorrelation time and the signal variance are approximately equal on the different grids (which should be true as we approach grid sufficiency), we can then estimate the standard error of the difference in sample means after arbitrary averaging times as

$$\widehat{S}_{\overline{\mathcal{X}}_A - \overline{\mathcal{X}}_B}\left(\hat{T}_A, \hat{T}_B\right) = \psi\sqrt{\frac{1}{\hat{T}_A} + \frac{1}{\hat{T}_B}}\,,$$

with $\psi^2 = \tau_{\text{decorr}}\, \sigma^2$ estimated from existing data as either

$$\psi^2 \approx T_A S^2_{\overline{\mathcal{X}}_A}$$

(using data from a single previous grid; e.g., if one has not yet started running on grid B) or

$$\psi^2 \approx \frac{T_A S^2_{\overline{\mathcal{X}}_A} + T_B S^2_{\overline{\mathcal{X}}_B}}{T_A + T_B}$$

(using data from two previous grids; e.g., if one has started running on both grids). Note that the second estimate is an approximate pooled variance; this estimate is biased without assumed knowledge of the decorrelation times, but presumably better than not using data from both grids if available.

The optimal averaging times are those that produce a given standard error at a minimal computational cost. Modeling this cost as $c_A \hat{T}_A + c_B \hat{T}_B$ and imposing the constraint

$$\widehat{S}_{\overline{f}_A - \overline{f}_B}\left(\hat{T}_A, \hat{T}_B\right) = \widehat{SE}_{\text{req}}\,,$$

the optimal solution can be found using a Lagrange multiplier as

$$\begin{aligned} \hat{T}_{A,\text{opt}} &= \frac{\psi^2}{\widehat{SE}^2_{\text{req}}}\left(1 + \sqrt{\frac{c_B}{c_A}}\right)\,, \\ \hat{T}_{B,\text{opt}} &= \frac{\psi^2}{\widehat{SE}^2_{\text{req}}}\left(1 + \sqrt{\frac{c_A}{c_B}}\right)\,. \end{aligned} \tag{3.19}$$

The ratio of the optimal averaging times is thus simply $\hat{T}_{A,\text{opt}}/\hat{T}_{B,\text{opt}} = \sqrt{c_B/c_A}$. For example, if one grid is four times more expensive than the other, the optimal choice is to average for twice as long on the cheaper grid.

If an estimate of ψ^2 is available, then the results of Eqn. (3.19) can also be used to estimate the maximum averaging time required, by using

$$\widehat{SE}_{\text{req}} = \frac{(1-\gamma)\Delta_0}{z_{\alpha/2} + z_{\beta_{\text{req}}}}\,, \tag{3.20}$$

which is required to produce the desired power (or Type II error) in the hypothesis test. Clearly, the optimal averaging times increase for greater desired confidence

Algorithm 2 Grid Convergence as a hypothesis testing.

```
1: procedure ( α, β, Δ_0, γ, Grids)
2:     compute \widehat{SE}_req from Eq. (3.20)
3:     Run Grid A (cheapest), compute first estimations and optimal times using
       equation 3.19.
4:     Run Grid B and test the null hypothesis.
5:     while 2P(Z < z_0) ≥ α & T ≤ \hat{T}_opt do
6:         Keep running Grids A and B till the optimal times, \hat{T}_Aopt and \hat{T}_Bopt, and
7:         Test the null hypothesis.
8:     if 2P(Z < z_0) ≤ α then
9:         Conclusive test. Grid B is sufficient with a (1-α) confidence.
10:    else
11:        Inconclusive results. Test declared Grid B insufficient with a power of (1
       -β).
12:        Restart from step 2.
13:        Grid B ⇒ Grid A
14:        Finer Grid ⇒ Grid B
```

(lower α and β_{req}, thus larger $z_{\alpha/2}$ and $z_{\beta_{\text{req}}}$), smaller allowed tolerance Δ_0, and for larger values of γ. The latter should be interpreted as follows: it is more expensive (larger averaging times) to require a specific probability β_{req} of a Type II error (unnecessarily running on a refined grid) when the true difference is close to the tolerance Δ_0 than when the true difference is closer to zero.

A computationally affordable path to find grid sufficiency is detailed in algorithm 2. Before executing it the experiment has to be designed, in which the parameters α, β, γ, and Δ_0 have to be set.

3.5.3.4 Application to Base Case 2

The proposed methodology for assessing grid sufficiency is demonstrated on Base Case 2 with the aerodynamic drag coefficient as the QoI. The flow is computed on three different grids, with key details listed in Table 3.7. The table shows the maximum values of the mesh resolution at the wall in viscous units and grid spacing along the edge of the boundary layer, where ξ, η, and z denote the airfoil azimuthal,

	cells	max $\Delta\xi_i^+$ at wall	max $\Delta\xi_i/\delta$ at BL edge	cost
Grid 1	$1960\times664\times156$	$(36.1, 1.35, 17.3)$	$(0.12, 0.019, 0.056)$	1.0
Grid 2	$2400\times800\times192$	$(29.3, 1.05, 14.1)$	$(0.10, 0.016, 0.048)$	2.8
Grid 3	$3072\times800\times256$	$(23.0, 1.05, 11.0)$	$(0.07, 0.015, 0.028)$	7.8

Table 3.7: Details on the computational grids, with directions ordered as streamwise, wall-normal, and spanwise. The normalized computational cost includes the number of time steps.

wall normal, and spanwise directions, δ is the boundary layer thickness, and all values are after the tripping. Each grid has 1.25 times more points in each direction than its predecessor, with the exception of the wall normal direction from g_{LES2} to g_{LES3}, since the mesh resolution at the wall, $\Delta\eta^+$, and at the boundary layer, $\Delta\eta/\delta$, achieved values representative of good quality LES simulations.

Note that the initial transients have been discarded and $t = 0$ is considered the start of the averaging, and that all times have been non-dimensionalized with the chord and the freestream velocity.

To illustrate the method, we choose the tolerance $\Delta_0 = 4\times^{-4}$ and $\alpha = 0.01$ ($z_{\alpha/2} = 2.58$); in other words, we want to be 99% confident that the drag coefficient differs by at most this tolerance between the final two grids. We also choose $\gamma = 0.85$ and $\beta_{\text{req}} = 0.05$ ($z_{\beta_{\text{req}}} = 1.65$); if the true difference between the grids is 85% of the allowable tolerance, then we want to be 95% certain of declaring the grid sufficient and thus avoiding having to create a finer grid.

The first step in determining grid sufficiency is to run on grids 1 and 2. The cost ratio of these grids is 2.8, and thus the optimal ratio of averaging times is $\sqrt{2.8} \approx 1.67$; we should therefore advance the simulations on grids 1 and 2 with that particular ratio until either the P-value (Eqn. 3.17) drops below α or the Type II error probability β (Eqn. 3.18) drops below β_{req}. The results are shown in Fig. 3.10 (left column). The P-value never reaches low levels (i.e., we are never close to declaring the grid sufficient), and we thus stop the simulations when the probability of a Type II error β reaches its required level after about $T_1 \approx 18$ and $T_2 \approx 11$.

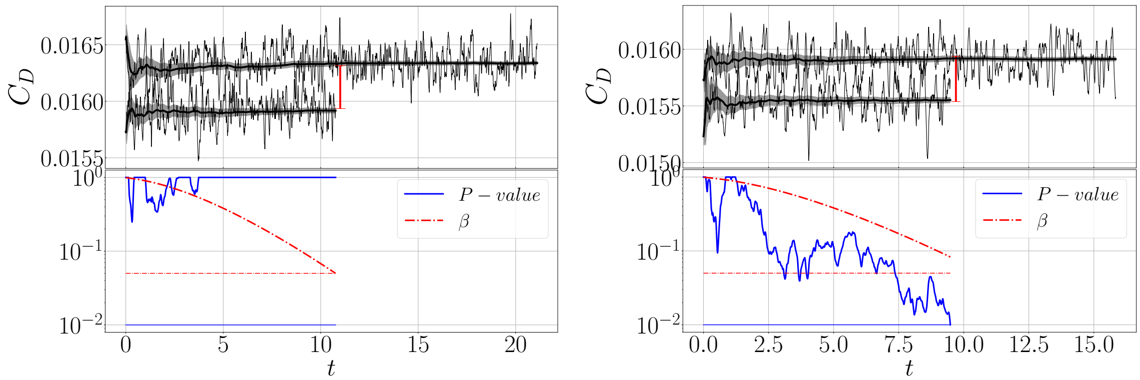

Figure 3.10: Results of the grid sufficiency tests, comparing grids 1-2 (left) and grids 2-3 (right). Top Figures: drag coefficient vs time for both grids (the coarser grid has the longer signal), with the tolerance Δ_0 visualized as the red errorbar. Bottom Figures: the P-value and Type II error β computed on-the-fly and compared to their required levels, plotted versus the time on the finer of the grids.

At this point, we are 95% confident that the grid is insufficient and we, therefore, create grid 3.

The cost ratio between grids 2 and 3 is the same and thus the optimal ratio of the averaging times is $\sqrt{2.8} \approx 1.67$. Since we already have results with $T_2 \approx 11$ we would start by running grid 3 up to $T_3 \approx 11/1.67 \approx 6.6$. If neither condition has been met, we would then advance the simulations on grids 2 and 3 at the optimal ratio $T_2/T_3 \approx 1.67$ until either condition is met. In the present case, this occurs after $T_2 \approx 16$ and $T_3 \approx 9.5$, when the P-value drops below the required $\alpha = 0.01$ level as can be seen in Fig. 3.10. At this point we are 99% confident that the grid is sufficient.

In practice, a user would presumably first run on grid 1 to get an initial view of the results before asking whether the grid is sufficient. In that spirit, imagine that grid 1 had been averaged for a time of $T_1 = 25$. The logical process is then to run on grid 2 until either $T_2 = 25/1.67 \approx 15$ or one of the conditions before are met. In this case, the β requirement would be met after $T_2 \approx 9$; earlier than the optimal scenario due to the additional averaging of signal 1.

3.5.4 Verification: Comparison against the *Nek5000* code

To verify the *Tortuga* LES solver when computing a flow past a wing, a large eddy simulation of an incompressible flow over a NACA 0012 profile with zero incidence and Reynolds number $Re = 4 \times 10^5$ is performed. *Tortuga* results are compared against results of the *Nek5000* code, developed by Fischer, Lottes and Kerkemeier [43], and obtained by Tanarro *et al.* [2]. The boundary layers are tripped at the same location in both studies. The grid used in this case is g_{LES3}, which for zero incidence has the following maximun spatial resolution at the wall in viscous units of $\max(\Delta\xi^+ = 22.0, \Delta\eta^+ = 1.0, \Delta z^+ = 10.0)$, where ξ, η, and z denote the airfoil azimuthal, wall normal, and spanwise directions, respectively. A maximum grid spacing along the edge of the boundary layer of $\max(\Delta\xi/\delta = 0.08, \Delta\eta/\delta = 0.016, \Delta z/\delta = 0.036)$, where δ is the boundary layer thickness. All values are after the tripping. It is worth noting that, although the tripping point is the same, the tripping mechanism between this study and the study of Tanarro *et al.* [2] is different, causing some differences near the tripping region. Important boundary layer quantities obtained by the current study are shown and compared with the benchmark in Fig. 3.11. Here, the black solid lines are the results of the present study, and the blue dash-dotted lines are the results of Tanarro *et al.* [2]. Figures 3.11(a) -(c) show the friction velocity at the wall, u_τ, the momentum thickness, θ_{BL}, and the streamwise velocity at the edge of the boundary layer, U_e, along the airfoil chord, respectively. Figure 3.11(d) exhibits the skin friction coefficient, $c_f = \tau_w/(1/2\rho U_e^2)$, as a function of $Re_\theta = U_\infty\theta/\nu$. Figures 3.11$(e)$ and (f) display the boundary layer thickness, δ, and displacement thickness, δ^*, along the airfoil chord, respectively.

Despite the difference in the tripping mechanism, good agreement is observed for the relevant boundary layer quantities. The main difference between the two studies is seen in the values of the boundary layer thickness, which is shown in Fig. 3.11(d). The reason for this could be that the *Tortuga* tripping mechanism increases the projected frontal area of the airfoil, causing a thicker boundary layer.

Figures 3.12(a) - (c) show a comparison of the velocity profiles, $\overline{u_t}^+$, at $x/c = 0.3$, $x/c = 0.65$ and $x/c = 0.9$, respectively. Here, black solid lines are the results of the present study and the blue dash-dotted are the results of Tanarro *et al.* [2]. Figures 3.12(d) - (f) display the profiles of the Reynolds stress tensor components, $\overline{u_t^2}^+$, $\overline{u_n^2}^+$, $\overline{w^2}^+$, and $\overline{u_t u_n}^+$, at $x/c = 0.3$, $x/c = 0.65$ and $x/c = 0.9$, respectively. Good agreement is seen in these profiles.

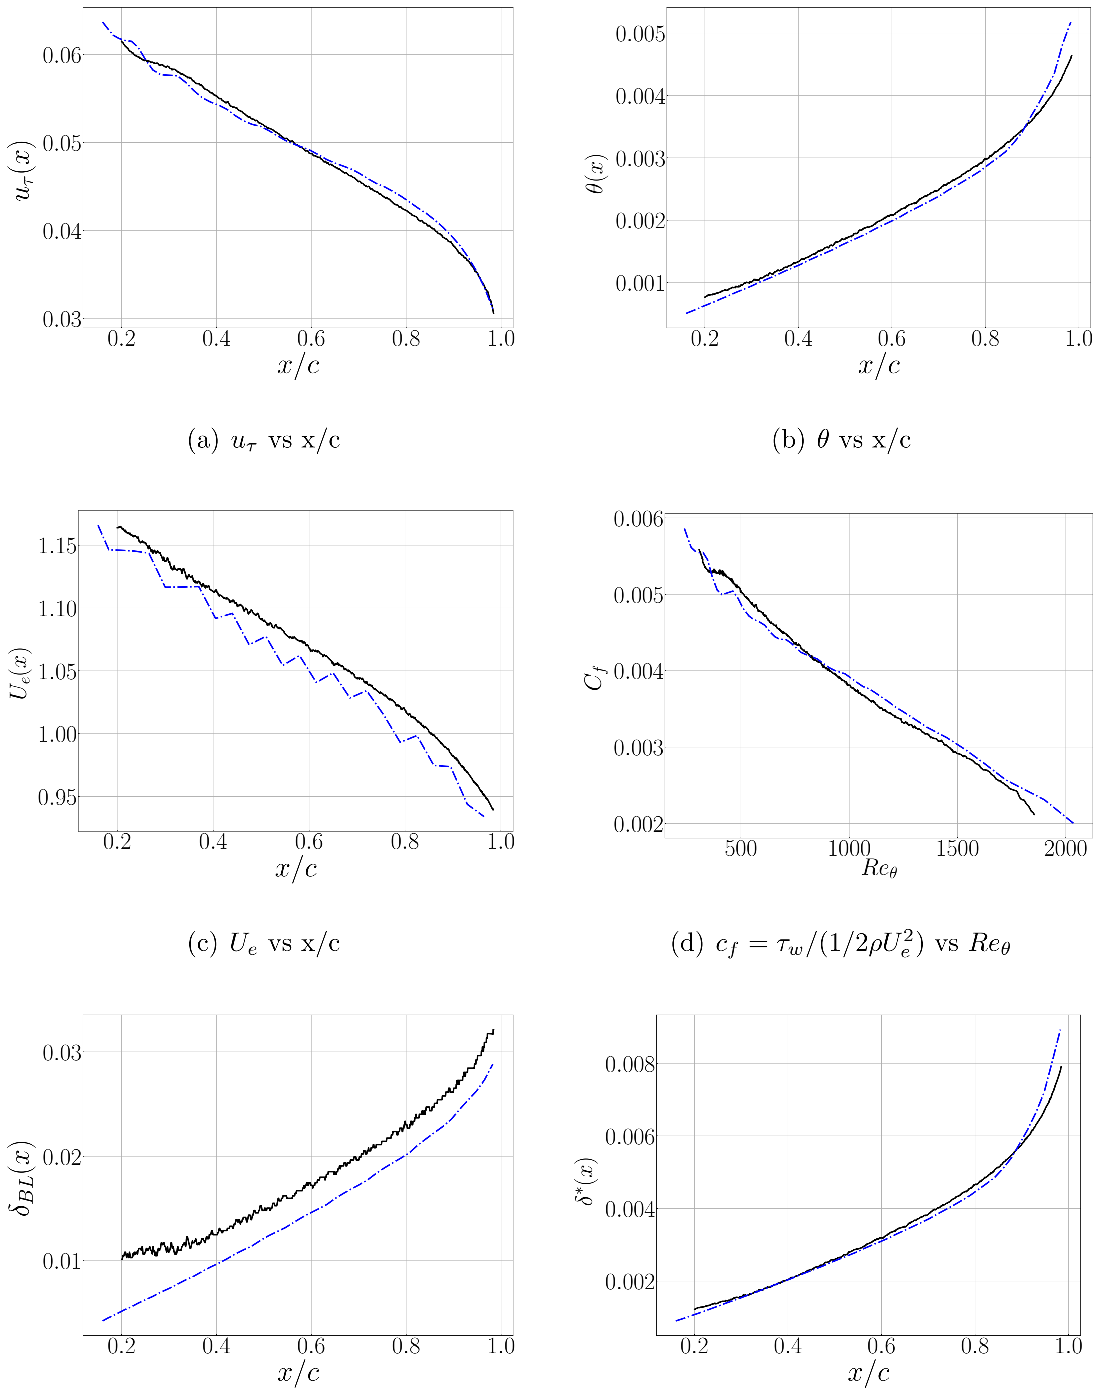

(a) u_τ vs x/c

(b) θ vs x/c

(c) U_e vs x/c

(d) $c_f = \tau_w/(1/2\rho U_e^2)$ vs Re_θ

(e) δ vs x/c

(f) δ^* vs x/c

Figure 3.11: Boundary layer quantities along the airfoil chord for Base Case 1; (*a*) friction velocity at the wall; (*b*) momentum thickness; (*c*) streamwise velocity at the edge of the boundary layer; (*d*) skin friction coefficient as a function of Re_θ; (*e*) boundary layer thickness; (*f*) displacement thickness. Black solid lines are the results of the present study, and the blue dash-dotted line are results of Tanarro *et al.* [2]

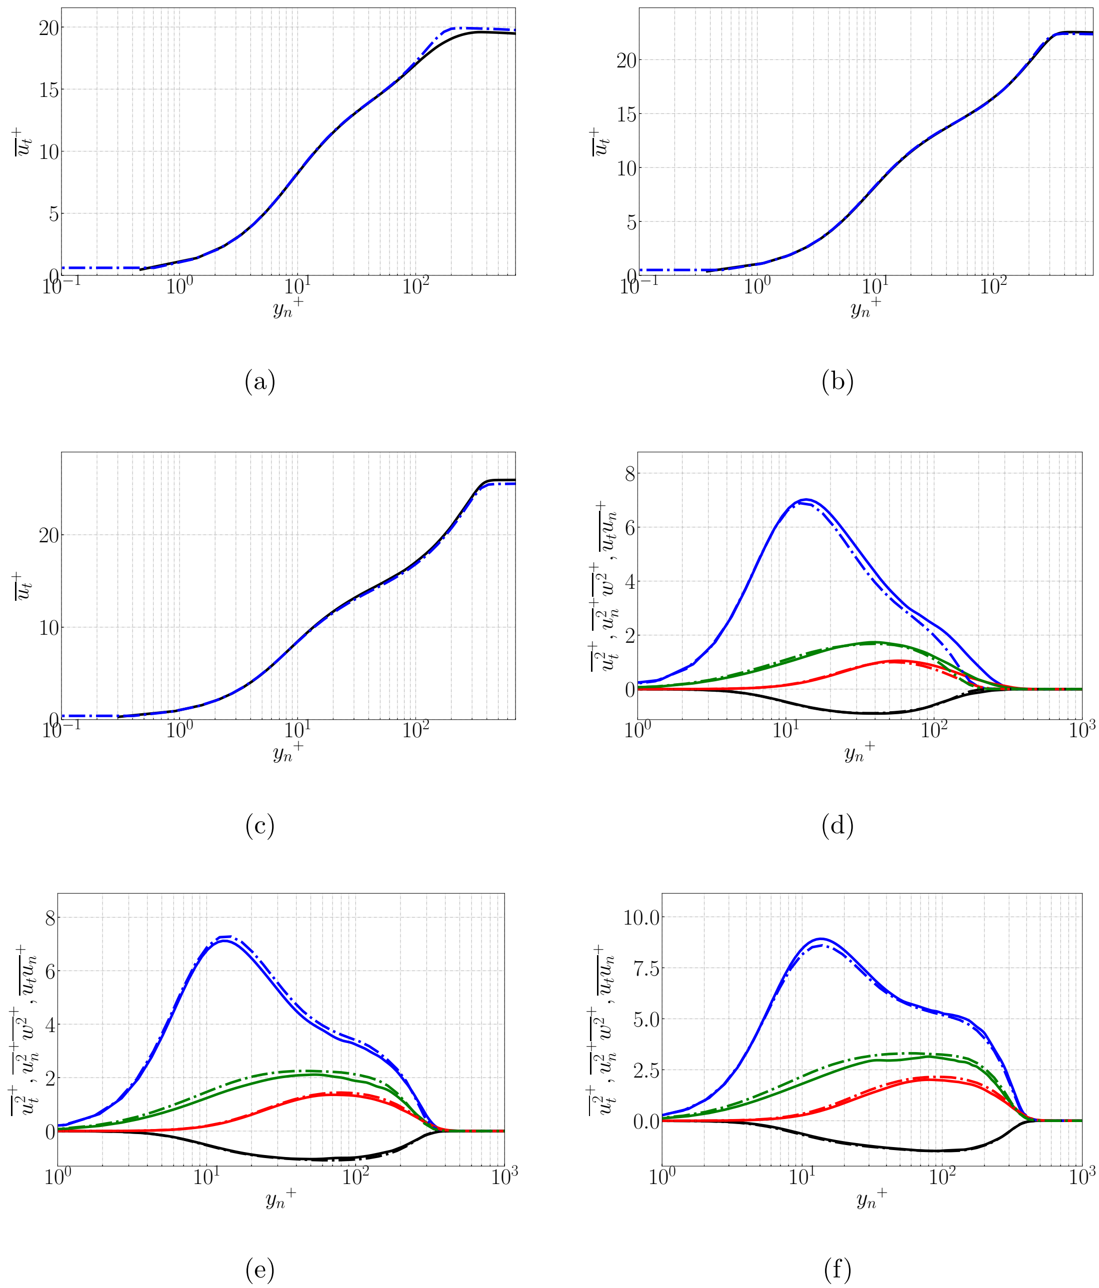

Figure 3.12: (a) - (c) Mean-velocity profiles in inner-scaled units at $x/c = 0.3$, $x/c = 0.65$ and $x/c = 0.9$, respectively; (d) - (f) Reynolds stress tensor components in inner-scaled units at $x/c = 0.3$, $x/c = 0.65$ and $x/c = 0.9$, respectively. Solid lines are the results of the present study, and the dash-dotted lines are from Tanarro *et al.* [2]

Chapter 4: Reynolds-averaged Navier-Stokes

This study required the development and use of a RANS solver for several reasons. First, the current state-of-the-art method to compute the sensitivity in design and engineering is based on RANS models. Since the assessment methodology considers a direct comparison against the current state-of-the-art method in industry, the RANS solver is used to compute the sensitivity of the QoI as the finite-difference in the parametric space. Second, the linearized problem is based on the non-linear RANS model, therefore, the formulation and implementation of the RANS solver facilitated the implementation of the linearized RANS. Third, one of the most common assumptions in standard sensitivity analysis using RANS is to assume the eddy viscosity is constant when perturbing the flow (known as the "frozen eddy viscosity assumption" or "frozen turbulence" [44–47]). In order to assess the impact of this assumption on the sensitivity predictions, a non-linear RANS frozen eddy viscosity solver was derived from the standard RANS solver.

The RANS solver developed by the author follows a standard approach but its formulation and verification are shown in this document for completeness and credibility.

4.1 Mathematical formulation

The derivation of the governing equations for RANS follows that of LES, except for a replacement of the filter operation $\widetilde{\langle\cdot\rangle}$ with an averaging operation $\overline{\langle\cdot\rangle}$. The Reynolds-averaged Navier-Stokes (RANS) equations for an incompressible and density constant flow of a Newtonian fluid in Cartesian tensor notation are

$$\frac{\partial \overline{u}_i}{\partial x_i} = 0 \quad , \tag{4.1}$$

$$\frac{\partial \overline{u}_i}{\partial t} + \overline{u}_j \frac{\partial \overline{u}_i}{\partial x_j} = -\frac{\partial \overline{p}}{\partial x_i} + \frac{\partial \left(2\nu \overline{S}_{ij} - \overline{u_i' u_j'}\right)}{\partial x_j} \quad , \tag{4.2}$$

where the mean velocity and pressure used in RANS are represented as $\overline{u}_i$ and $\overline{p}$, for simplicity.

4.1.1 Turbulence modeling

The current study only considers the constitutive relation (or Boussinesq assumption) for modeling the Reynolds stress tensor. The Reynolds stress is modeled as

$$-\overline{u_i' u_j'} = \tau_{ij} = 2\nu_t \overline{S}_{ij} - \frac{2}{3} k \delta_{ij} \quad , \tag{4.3}$$

and equation 4.2 can be re-written as

$$\frac{\partial \overline{u}_i}{\partial t} + \overline{u}_j \frac{\partial \overline{u}_i}{\partial x_j} = -\frac{\partial \overline{p}}{\partial x_i} + \frac{\partial \left[2(\nu + \nu_t) \overline{S}_{ij}\right]}{\partial x_j} \quad , \tag{4.4}$$

where the isotropic part of the Reynolds stress tensor has been absorbed into the pressure.

The current study uses two models to predict the eddy viscosity, the two-equation k-ω version 2006 developed by Wilcox [3] and a zero-equation algebraic model based on the Prandtl's mixing length. Their formulation are shown here for purposes of clarity and completeness.

The numerical methods used in this solver are the same as the ones used for LES in section 3.2.

4.1.1.1 The k-ω 2006 model

The governing equations for this model are [3]

$$\nu_t = \frac{k}{\tilde{\omega}}, \quad \tilde{\omega} = \max\left[\omega, C_{lim}\sqrt{\frac{2\overline{S}_{ij}\overline{S}_{ij}}{\beta^*}}\right], \quad C_{lim} = 7/8 \quad , \tag{4.5}$$

$$\frac{\partial k}{\partial t} + \overline{u}_j \frac{\partial k}{\partial x_j} = \tau_{ij}\frac{\partial \overline{u}_i}{\partial x_j} - \beta^* k\omega + \frac{\partial\left[(\nu + \sigma^* \frac{k}{\omega})\frac{\partial k}{\partial x_j}\right]}{\partial x_j} \quad , \tag{4.6}$$

$$\frac{\partial \omega}{\partial t} + \overline{u}_j \frac{\partial \omega}{\partial x_j} = \alpha\frac{\omega}{k}\tau_{ij}\frac{\partial \overline{u}_i}{\partial x_j} - \beta\omega^2 + \frac{\sigma_d}{\omega}\frac{\partial k}{\partial x_j}\frac{\partial \omega}{\partial x_j} + \frac{\partial\left[(\nu + \sigma \frac{k}{\omega})\frac{\partial \omega}{\partial x_j}\right]}{\partial x_j} \quad , \tag{4.7}$$

where k is the turbulent kinetic energy, ω is the specific dissipation rate and σ^*, β^*, α, σ_d, and σ are closure coefficients.

The k-ω RANS equations can be denoted by

$$\mathcal{R}(\overline{\mathbf{v}}^{k\omega}, \nu_t^{k\omega}, k, \omega; \mathbf{a}) \quad , \tag{4.8}$$

here, the superscript $\langle\cdot\rangle^{k\omega}$ is used to differentiate the solution given by the the k-ω RANS model from the one inferred using the LES model.

4.1.1.2 Boundary conditions for the k-ω model

In this work, three different configurations - a turbulent incompressible flow in a plane channel, over a flat plate, and past an airfoil - are used to verify the solver. Each of the three different configurations makes use of the following boundary conditions for the turbulent kinetic energy: for inflow, a Dirichlet boundary condition; for no-slip walls, a Dirichlet boundary condition equal to zero; and for outflow, a homogeneous Neumann boundary condition. For the specific dissipation rate: the inflow and outflow boundary conditions are similar to $\overline{u}_i$ and k. However, the boundary condition for a no-slip wall is set to be a function of the friction velocity, as suggested by Saffman for a "slightly-rough-surface" boundary condition

$$\omega = \frac{{u_\tau}^2}{\nu} S_R \quad , \tag{4.9}$$

with S_R being the surface roughness factor, defined as

$$S_R = \left(\frac{200}{{k_s}^+} \right)^2 , \quad k_s^+ \leq 5 \quad . \tag{4.10}$$

4.1.1.3 Prandtl's algebraic eddy viscosity model

The eddy viscosity model using Prandtl's mixing length hypothesis is

$$\nu_t = \kappa^2 d^2 \left| S \right| D \quad , \tag{4.11}$$

where $|S|$ is the magnitude of the strain rate tensor

$$|S| = \sqrt{2 \overline{S}_{ij} \overline{S}_{ij}}, \tag{4.12}$$

D is the VanDriest damping function given by

$$D = \left[1 - \exp \left[\frac{-d^+}{A^+} \right] \right]^2 , \tag{4.13}$$

κ and A^+ are model constants, and d is the wall-normal distance.

4.2 Verification: RANS of plane channel flow at $Re_\tau = 395$

A turbulent incompressible flow in a plane channel at Reynolds number $Re_\tau = 395$ is used as a first verification case for the RANS solver. The channel has a periodic boundary condition in the streamwise and spanwise directions and a no-

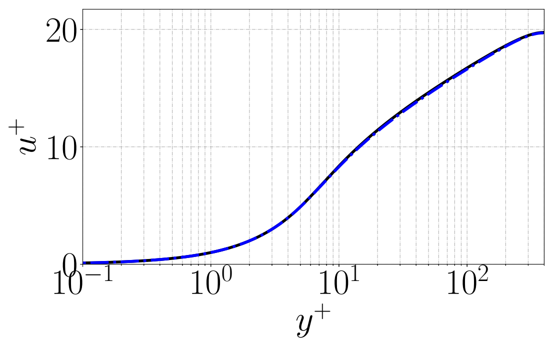

(a)

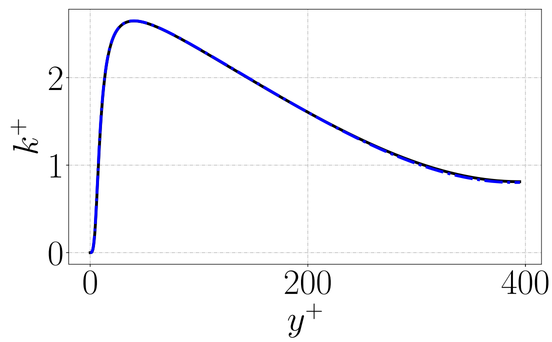

(b)

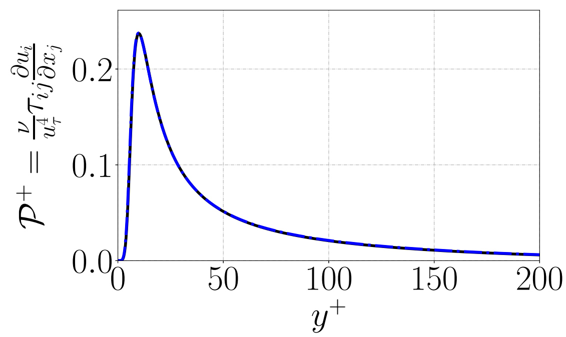

(c)

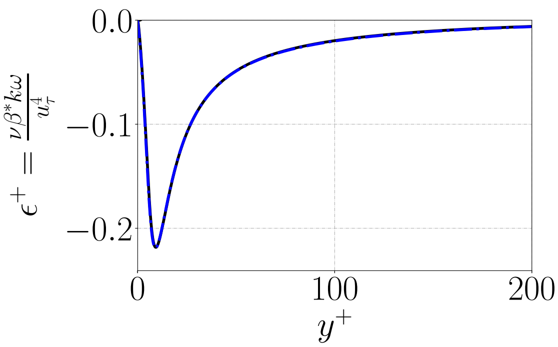

(d)

Figure 4.1: Turbulent channel flow at $Re_\tau = 395$; (a) velocity profile; (b) turbulent kinetic energy; (c) production of turbulent kinetic energy; (d) dissipation of turbulent kinetic energy. All profiles are in viscous units. The black solid lines are the results obtained from the in-house RANS solver and the blue dash-dotted lines are the results from Wilcox [3].

slip boundary condition at the walls. The spatial resolution at the wall in viscous units is $\Delta\eta^+ = 0.1$, where η denotes the airfoil wall normal direction. The k-ω model is used and results are compared with results from Wilcox [3]. Figures 4.1 (a) to (d) show profiles in viscous units for velocity $u^+ = u/u_\tau$, turbulent kinetic energy $k^+ = k/u_\tau^2$, production of turbulent kinetic energy $\mathcal{P}^+ = \nu\tau_{ij}\partial u_i/\partial x_j/u_\tau^4$, and dissipation of turbulent kinetic energy $\epsilon^+ = \nu\beta^* k\omega/u_\tau^4$, respectively. The black solid lines are the results obtained from the in-house RANS solver and the blue dash-dotted lines are the results from Wilcox [3]. One can see that the in-house solver matches the results of the benchmark, and it captures the model's intended behavior.

4.3 Verification: RANS of the flow over a flat plate at $Re_L = 5 \times 10^6$

A turbulent flow over a flatplate at Reynolds number $Re_L = 5 \times 10^6$ is used as a second verification case for the in-house solver. The k-ω model is used. As a benchmark, the results obtained from the CFL3D code published at the Turbulence Modeling Resource (TMR) [4] are used. The boundary conditions are as follows. The inflow uses a Dirichlet boundary condition in which a uniform flow is set, the outflow uses a homogeneous Neumann boundary condition, the top uses a freestream boundary condition, and the solid wall has two conditions: a symmetry, or slip, boundary condition from the inflow (x=-0.5) to x=0, and then a no-slip boundary condition from x=0 to the end. The spatial resolution at the wall in viscous units is $\Delta\eta^+ = 0.12$, where η denotes the airfoil wall normal direction. The results are shown and compared with the results from CFL3D in Fig. 4.2. The black solid lines are the results obtained from the in-house RANS solver and the blue dash-dotted lines are the results from CFL3D. Figure 4.2 (a) shows the skin friction coefficient as a function of x. Figures 4.2 (b) to (d) show profiles in viscous units for velocity, turbulent kinetic energy, and eddy viscosity, respectively. One can see that the in-house solver matches the results of the benchmark. Furthermore, Figs. 4.2 (e) and (f) show the behavior of turbulent kinetic energy and eddy viscosity in the viscous sublayer. It can be seen that these turbulent quantities behave as

$$k^+ = \frac{k}{u_\tau^2} \sim (y^+)^{3.3}, \quad \omega^+ = \frac{\omega\nu}{u_\tau^2} \sim \frac{1}{(y^+)^2}, \quad \nu_t^+ = \frac{\nu_t}{\nu} \sim (y^+)^{5.3} \quad . \tag{4.14}$$

The fact that these quantities goes as these particular power laws means two things. First, it verifies the results obtained by the in-house solver, particularly the turbulent quantities in this case, follow the originally designed behavior in the viscous sublayer [48]. Second, it explains why the k-ω model requires a finer grid spacing near the wall, where usually a value of $\eta^+ \leq 0.1$ is suggested.

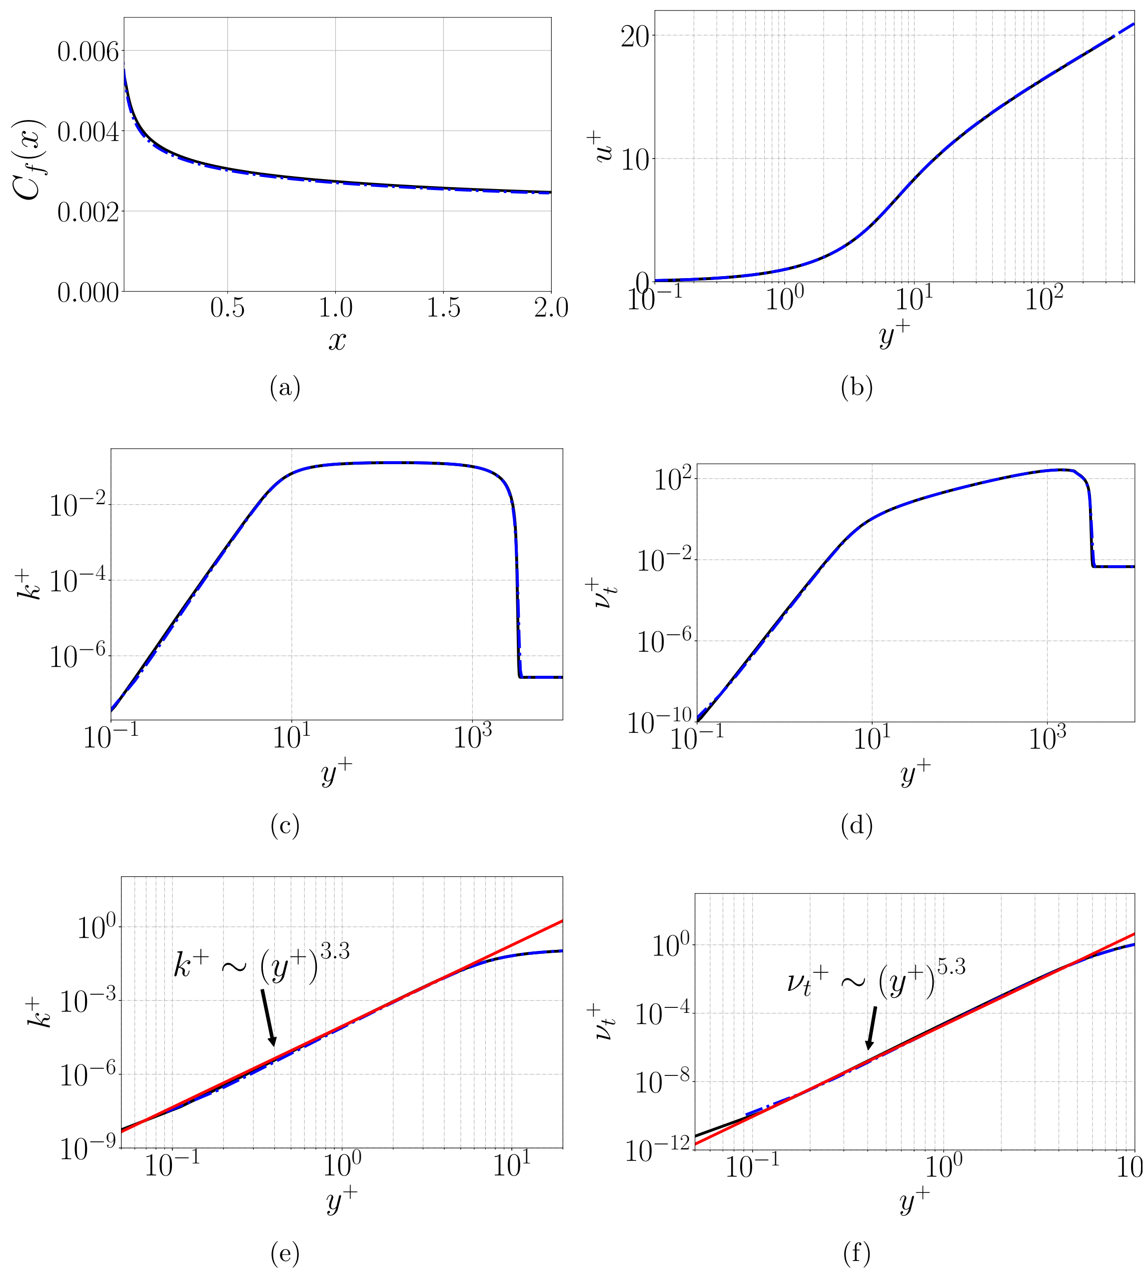

Figure 4.2: Turbulent flatplate flow at $Re_{x=1} = 5 \times 10^6$; (a) skin friction coefficient profile as a function of x; (b) velocity profile at $x = 1.0$; (c) turbulent kinetic energy as function of y at $x = 1.0$; (d) eddy viscosity as function of y at $x = 1.0$; (e) behavior of the turbulent kinetic energy in the viscous sublayer; (f) behavior of the eddy viscosity in the viscous sublayer. All profiles are in viscous units. The black solid lines are the results obtained from the in-house RANS solver and the blue dash-dotted lines are the results from CFL3D. The red solid line represents the respective slope in figures (e) and (f).

4.4 RANS of a turbulent flow past a NACA 0012 profile

4.4.1 Verification: RANS of the flow past a NACA 0012 profile at $Re = 6 \times 10^6$ and two angles of attack, $\theta = 0°$ and $\theta = 10°$

Two cases are chosen to verify the RANS solver on a wing. In both cases, we use a flow past an airfoil NACA 0012 at a Reynolds number $Re = 6 \times 10^6$. In the first case, we use an angle of attack $\theta = 0°$, while the second case uses an angle of attack $\theta = 10°$. As a benchmark, we use results obtained from the CFL3D code mentioned in the previous section. A symmetric body-fitted O-grid block with a computational domain of 500 chords is used. For these two particular cases, the NACA0012 profile used differs from the one used throughout this study. It has a sharp trailing edge, and it uses the revised NACA 0012 formula given in [4] (`https://turbmodels.larc.nasa.gov/naca0012_val.html`). The sharp trailing edge is used only in these two cases for purposes of comparison against benchmark results. The boundary conditions are as follows. The outer farfield circumference uses a freestrem boundary condition, and the solid wall is set to use a no-slip boundary condition. Results are show and compared with results from CFL3D in Fig. 4.3. The black solid lines are the results obtained from the in-house RANS solver and the blue dash-dotted lines are the results from CFL3D. Figures 4.2 (a) and (b) show the skin friction coefficient and pressure coefficient profiles as a function of x/c for $\theta = 0°$. Figures 4.2 (c) and (d) show the skin friction coefficient and pressure coefficient profiles as a function of x/c for $\theta = 0°$.

Table 4.1: Grid spacing information for three different grids used in the RANS convergence study.

Grid	$_{max}(\Delta\eta^{+})$	$\mathcal{C}_D$	$\mathcal{C}_L$
g_{RANS1}: 1920×658	(0.15)	0.0153238	0.534848
g_{RANS2}: 2400×750	(0.12)	0.0150800	0.535400
g_{RANS3}: 3072×940	(0.095)	0.0147847	0.536801

4.4.1.1 Freestream boundary conditions for k and ω

In this work, when solving the RANS model of a turbulent flow past a NACA 0012, the farfield values of k and ω are set as

$$\omega_{farfield} = \frac{5U_\infty}{L} \quad , \quad k_{farfield} = 1 \times 10^{-6} U_\infty^2 \tag{4.15}$$

As suggested in the TMR website, The farfield boundary condition for k in equation (4.15) corresponds to a freestream turbulence level of 0.08 [4]

4.4.2 RANS of Base Case 2: $\theta = 5^\circ$ and $Re = 4 \times 10^5$

The in-house RANS solver is used to predict the aerodynamic coefficients for Base Case 2 from Table 1.1. A two-dimensional and symmetric body-fitted O-grid block is used to simulate a turbulent flow past an airfoil. The grid domain and the airfoil trailing edge are the same as the ones shown in Figs. 3.3 (a) and (b), respectively. A traditional grid convergence study in which the relative change in the aerodynamic coefficients between grids is used to decide wether these QoIs converge. The relevant information for the three finest grids used in this study is shown in Table 4.1. The relative change seen $\mathcal{C}_D$ and $\mathcal{C}_L$ between g_{RANS1} and g_{RANS2} is 2.05 percent and 0.26 percent, respectively. The relative change in $\mathcal{C}_D$ and $\mathcal{C}_L$ between g_{RANS1} and g_{RANS2} is 1.5 percent and 0.09 percent. If the threshold is set at $\Delta = 2$ percent, one can declare the aerodynamic coefficients grid converged and the values of grid g_{RANS3} are used.

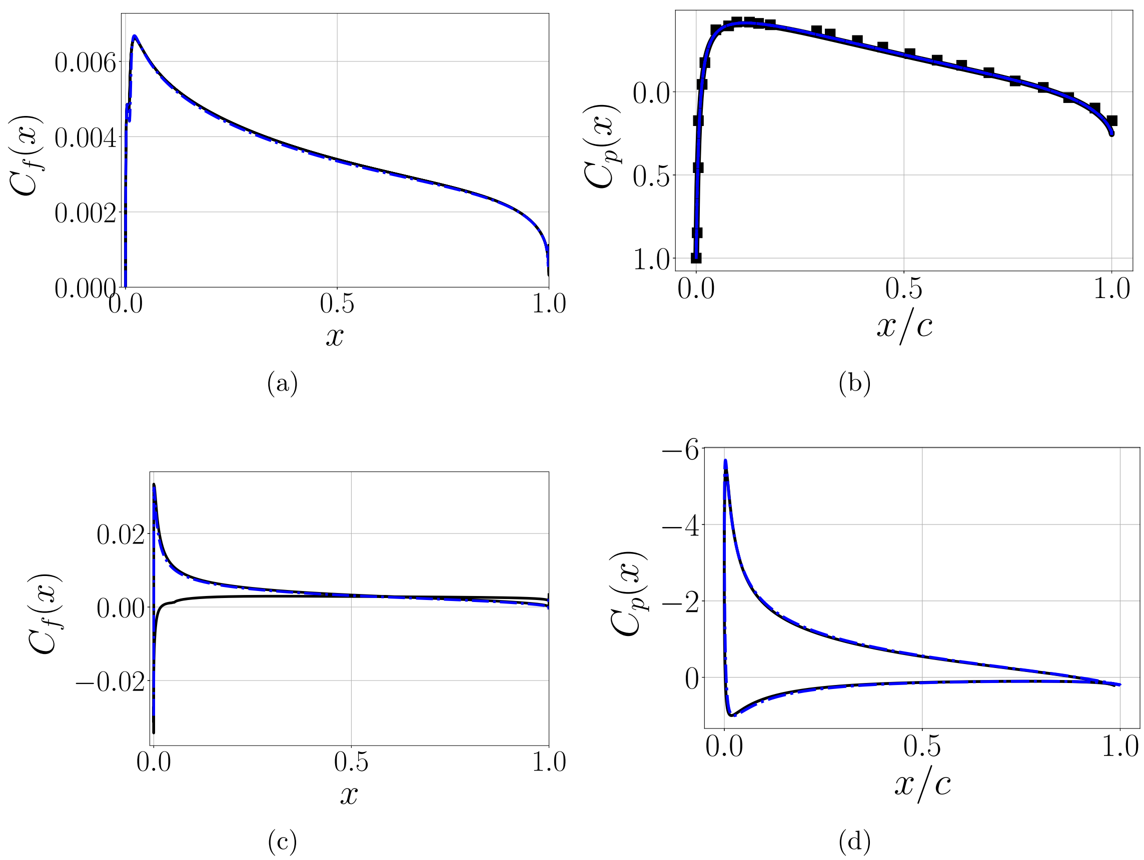

Figure 4.3: Turbulent flow past an airfoil NACA 0012 at a Reynolds number $Re = 6 \times 10^6$ and two angles of attack, $\theta = 0°$ and $\theta = 10°$; (a) skin friction coefficient profile as a function of x/c for $\theta = 0°$; (b) pressure coefficient profile as a function of x/c for $\theta = 0°$; (c) skin friction coefficient profile as a function of x/c for $\theta = 10°$; (d) pressure coefficient profile as a function of x/c for $\theta = 10°$. The black solid lines are the results obtained from the in-house RANS solver and the blue dash-dotted lines are the results from CFL3D. The black square dots in figure (b) are experimental results obtained by Gregory and O'Reilly and obtained from [4].

Chapter 5: Finite-difference sensitivity

Finite differences (FD) in the parametric space is the method that obtains the closest approximation to the true sensitivity of J_j in the context of LES. The sensitivity $\partial J_j/\partial a_i$ can be approximated by computing the difference of J_j from $n+1$ different LES runs, one at $a_{i,b}$ and n at $a_{i,p} = a_{i,b} + \Delta a_i$, where n is the order of the finite-difference stencil, the subscripts b and p represent base and perturbed conditions, respectively, and Δa_i will vary accordingly. In this research, the sensitivity $\partial J_j/\partial a_i$ computed using a first-order forward finite-difference stencil, is considered to be the true value. This FD-sensitivity is used as the benchmark against which the results obtained by the current state-of-the-art (fully RANS-based) and the proposed MFSA methods will be compared. However, when computing the sensitivity of a time-averaged quantity using finite-differences, two types of errors are inherently incurred: "TR error", or the truncation error due to finite-differencing, and "AV error", or the uncertainty due to insufficient average of the QoI.

The FD-sensitivity of $\mathcal{C}_D$ with respect to Re using a first-order forward finite-difference stencil from LES is given by

$$\left.\frac{d\overline{\mathcal{C}_D}}{dRe}\right|_{FD} \approx \frac{\overline{\mathcal{C}_D}(\bar{\mathbf{v}}(Re_p)) - \overline{\mathcal{C}_D}(\bar{\mathbf{v}}(Re_b))}{\Delta Re} - \underbrace{\frac{\Delta Re}{2}\frac{d^2\overline{\mathcal{C}_D}}{dRe^2}}_{\text{TR error}} \pm \underbrace{\frac{CI_{\overline{\mathcal{C}_D}(\bar{\mathbf{v}}(Re_p))-\overline{\mathcal{C}_D}(\bar{\mathbf{v}}(Re_b))}}{\Delta Re}}_{\text{AV error}}, \quad (5.1)$$

The same equation applies for all other QoIs. In this chapter, the in-house LES solver (already verified) is used to approximate the true sensitivity for Base Case 2. The RANS solver (also verified) is going to be leveraged to approximate the truncation error. The statistical procedure presented in section 3.5 is used to assess the behavior of the uncertainty due to insufficient averaging of the QoI.

5.1 TR error: Truncation error due to finite-differencing

Truncating the finite-difference series expansion to the first-order results in an error, which is a function of ΔRe. To compute the truncation error, one should ideally run a suite of large eddy simulations at different Reynolds numbers, followed by the statistical convergence and the computation of the sample mean of J_j for each simulation. However, the computational cost of such an approach is prohibitively high in this study. Instead, because the flow around the airfoil remains attached for Base Case 2, the RANS solver is leveraged to approximate the truncation error.

A suite of RANS at $\theta = 5°$ with different Reynolds numbers are solved. Figures 5.1(a) and (b) show the curves of $\Delta\mathcal{C}_D = \mathcal{C}_D(\overline{\mathbf{v}}^{k\omega}(Re_p)) - \mathcal{C}_D(\overline{\mathbf{v}}^{k\omega}(Re_b))$ and $\Delta\mathcal{C}_L = \mathcal{C}_L(\overline{\mathbf{v}}^{k\omega}(Re_p)) - \mathcal{C}_L(\overline{\mathbf{v}}^{k\omega}(Re_b))$ plotted against $\Delta Re = Re_p - Re_b$, respectively. The blue square dots represent the discrete values obtained by solving the RANS, and the black solid line is a least-squares fit quadratic polynomial interpolated from these square dots. This least-squares fit is leveraged to approximate the analytical sensitivity at Re_b. The approximated analytical sensitivities computed from the least-square fit are denoted from here on as $d\mathcal{C}_D/dRe|_{lsq}$ and $d\mathcal{C}_L/dRe|_{lsq}$ and are reported in Table 5.4.

Figures 5.2(a) and (b) show the FD-sensitivities in blue solid lines for $d\mathcal{C}_D/dRe|_{FD}$ and $d\mathcal{C}_L/dRe|_{FD}$, respectively. The FD-sensitivities are shown as functions of ΔRe, and compared against the approximate analytical sensitivity (black solid line) computed at Re_b. One can see that, as ΔRe increases, the blue lines diverge from the black lines. The difference between these two lines is an approximation of the truncation error. Also, it is observed that truncation error behaves linearly within the span of ΔRe shown in the plot.

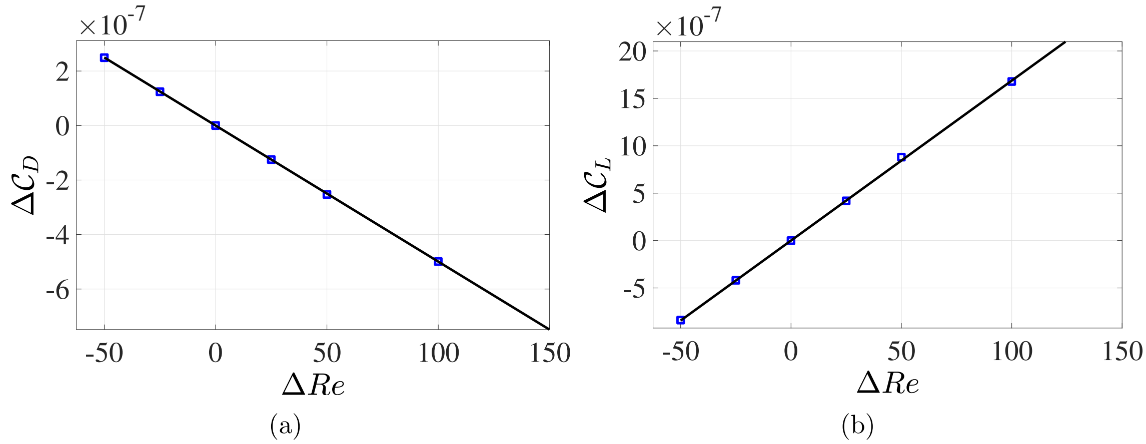

Figure 5.1: Curves of (a) $\Delta\mathcal{C}_D = \mathcal{C}_D(\overline{\mathbf{v}}^{k\omega}(Re_p)) - \mathcal{C}_D(\overline{\mathbf{v}}^{k\omega}(Re_b))$ and (b) $\Delta\mathcal{C}_L = \mathcal{C}_L(\overline{\mathbf{v}}^{k\omega}(Re_p)) - \mathcal{C}_L(\overline{\mathbf{v}}^{k\omega}(Re_b))$, plotted against $\Delta Re = Re_p - Re_b$.

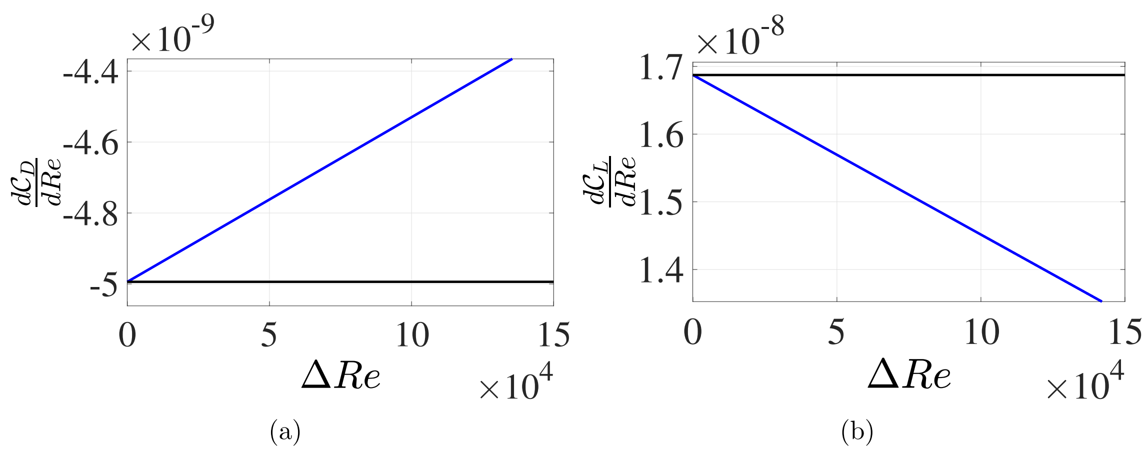

Figure 5.2: Behavior of the sensitivity using finite-difference (blue solid line) and analytical derivative (black solid line) computed at Re_b for (a) $\mathcal{C}_D$ and (b) $\mathcal{C}_L$.

The absolute relative difference between the approximate analytical and the finite-difference sensitivity is plotted in Figs. 5.3(a) and (b). One can see, for example, that using a $\Delta Re = 1 \times 10^5$ would create a TR, or truncation, error of approximately 10 percent and 15 percent in the sensitivity values of $\mathcal{C}_D$ and $\mathcal{C}_L$, respectively. Having an approximation of the behavior of the TR error is going to be useful in section 5.3, in which it is discussed how to obtain an "optimal" ΔRe.

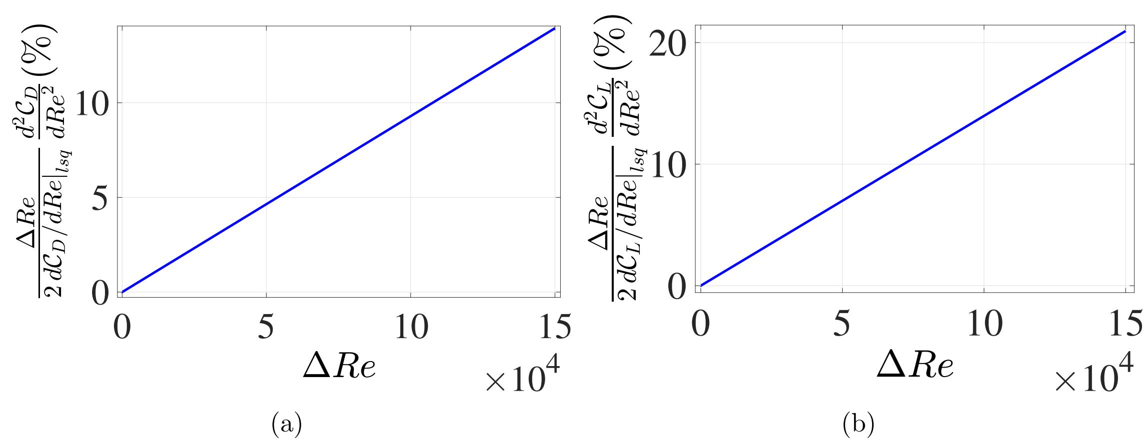

Figure 5.3: Absolute relative difference between the analytical and finite-difference sensitivities for (a) $\mathcal{C}_D$ and (b) $\mathcal{C}_L$, respectively.

A similar analysis is done to assess the effects of the truncation error in the sensitivity of $\mathcal{C}_f$ and $\mathcal{C}_p$ with respect to Re, and shown in Fig. (5.4). In this figure, the sensitivities of the skin friction and pressure coefficients are computed using two different values of ΔRe. The blue dashed lines represent the sensitivities using a $\Delta Re = 50$ and the red dotted lines represents the sensitivities using a $\Delta Re = 1 \times 10^5$. The most perceptible effect is observed between $0.1 < x/c < 0.25$. In this region, one can see that the larger the ΔRe, the difference in the transition point shifts upstream with a lower value for both sensitivities. Similarly, the larger the ΔRe, the smaller the area inside the curves of $d\mathcal{C}_f/dRe$ and $d\mathcal{C}_p/dRe$.

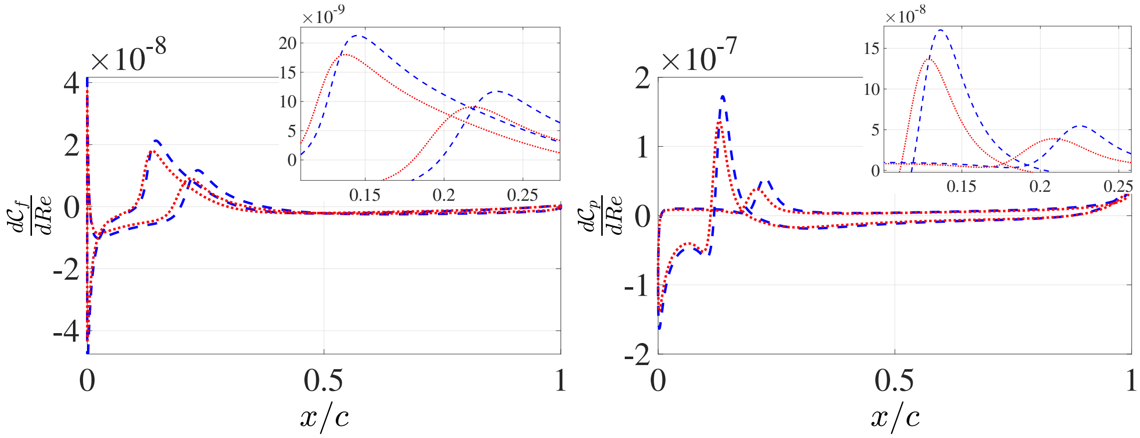

Figure 5.4: Effects of the truncation error in the sensitivity of the skin friction coefficient (left), and the pressure coefficient (right), with respect to the Reynolds number.

5.2 AV error: Uncertainty due to insufficient averaging of the QoI

Computing the FD-sensitivity of two LES requires computing the difference between two imperfectly averaged quantities. Following the procedure explained in section 3.5.2.2, the confidence interval of the FD-sensitivity for $\mathcal{C}_D$ is computed as (the same equation applies for $\mathcal{C}_L$)

$$\frac{CI_{\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_p))-\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_b))}}{\Delta Re} = \frac{t_{\alpha/2,df}\sqrt{S^2_{\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_p))} + S^2_{\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_b))}}}{\Delta Re} \quad . \tag{5.2}$$

In order to quantify the confidence interval of the FD-sensitivity, one would need to solve two LES. However, before running a second LES, a value of ΔRe has to be chosen. Currently, the only term known in equation (5.2) is the standard error for the Base Case, $S_{\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_b))}$. To be able to do a prospective analysis, the sample variance for the perturbed case, $S_{\widehat{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_p))}$, is approximated to be the same as the one found for the Base Case 2. This approximation is made taking into account that the perturbed case is going to be solved using the same parameters used in Base Case 2, with the only difference being the Reynolds number. With these approximations, equation (5.2) can be written as

$$\frac{CI_{\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_p))-\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_b))}}{\Delta Re} = \frac{t_{\alpha/2,df}S_{\widehat{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_b))}}{\Delta Re}\sqrt{\frac{1}{n_{b,p}}+\frac{1}{n_{b,b}}} \quad , \tag{5.3}$$

where $n_{b,p}$ and $n_{b,b}$ are the number of batches used to compute statistics for the perturbed and base cases, respectively. In the particular case in which the number of batches is the same in both simulations, equation (5.3) reduces to

$$\frac{CI_{\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_p))-\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_b))}}{\Delta Re} = \frac{t_{\alpha/2,df}}{\Delta Re}\sqrt{2S^2_{\overline{\mathcal{C}_D}(\overline{\mathbf{v}}(Re_b))}} \quad , \tag{5.4}$$

which computes the lowest confidence interval in the isosurface $n_{b,p} + n_{b,p} = C$.

In the following analysis, the equation (5.4) is used to estimate the confidence interval for the FD-sensitivity of $\mathcal{C}_D$ using the time series shown in Fig. 3.7(b). Figure 5.5 shows the behavior of the confidence interval as a function of the total simulation time T_s^* and ΔRe for the Base Case 2. The total simulation time T_s^* is the sum of both simulation times (base and perturbed cases). In this figure, the confidence interval estimations are divided by $m_{\mathcal{C}_D}$ and $m_{\mathcal{C}_L}$, the approximate analytical sensitivity using the surrogate RANS curves. Since all curves use the same t_b^*, doubling the total simulation time T_s^* doubles the number of batches n_b. For example, for a $\Delta Re = 4 \times 10^4$ and $\Delta Re = 1 \times 10^5$, the uncertainties for different T_s^* are given in Table 5.1. If one chooses a $\Delta Re = 1 \times 10^5$, one would need an estimated $T_s^* = 20$ to obtain an uncertainty of 9.8 percent for $\mathcal{C}_D$ and 33.5 percent for $\mathcal{C}_L$. Under the current approximations, the total simulation time T_s^* is evenly split, meaning that both simulations are required to run for 10 units of time (the simulation time referenced in this study does not include the initial transient simulation time).

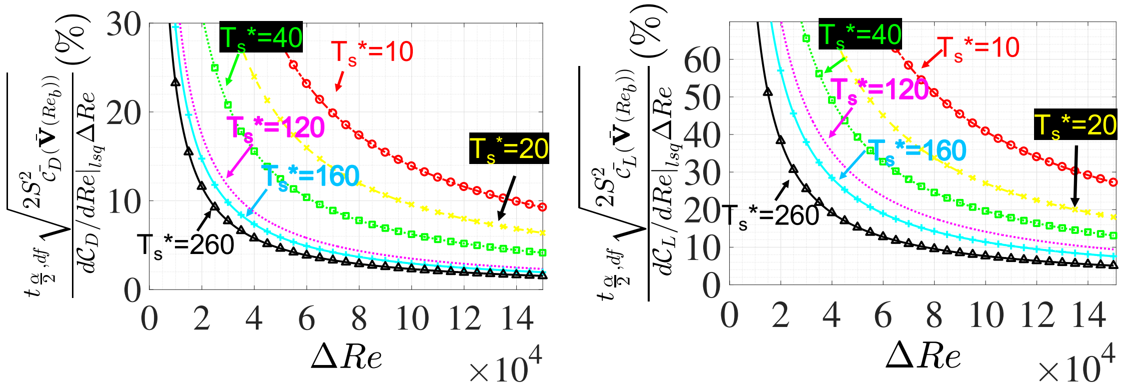

(a) (b)

Figure 5.5: Confidence interval (equation 5.4) for the FD-sensitivity of (a) drag and (b) lift coefficients with respect to Reynolds number for Base Case 2. The number of batches is assumed to be the same in both simulations. The red dash-dotted lines with circle symbols, the yellow dashed lines with cross symbols, the green dotted lines with square symbols, the magenta dotted lines, the cyan solid lines with plus sign symbols, the orange solid lines with asteric symbols, and the black solid lines with triangle symbols represent the estimation of the uncertainty due to insufficient averaging for $T_s^* = 10$, $T_s^* = 20$, $T_s^* = 40$, $T_s^* = 120$, $T_s^* = 160$, and $T_s^* = 260$, respectively.

Table 5.1: Uncertainty due to insufficient averaging for the sensitivity of drag and lift coefficients with respect to Reynolds number for Base Case 2.

T_s^*	$\Delta Re = 1 \times 10^5$		$\Delta Re = 4 \times 10^4$	
	$CI_{d\overline{C_D}/dRe}$ (%)	$CI_{d\overline{C_L}/dRe}$(%)	$CI_{d\overline{C_D}/dRe}$ (%)	$CI_{d\overline{C_L}/dRe}$(%)
10	13.90	40.85	34.70	102.20
20	9.57	27.02	23.87	67.64
40	6.23	19.65	15.57	48.96
80	4.30	15.58	10.83	38.94
120	3.49	14.13	8.73	35.33
160	2.94	11.38	7.33	28.44
200	2.63	9.87	6.55	24.72
260	2.35	7.64	5.88	19.00

5.3 Finding the optimal ΔRe to compute the finite-difference sensitivity from LES

To compute the finite-difference sensitivity, one needs to decide what value of ΔRe to use. This decision affects the finite-difference sensitivity predictions, as shown in the last two sections. Furthermore, the two types of error (truncation and insufficient averaging) behave inversely and at different rates when ΔRe changes. Since this is the closest approximation to the true sensitivity values, it is necessary to understand the implications of choosing a specific ΔRe value. Figures 5.6(a) and (b) present the estimations of the finite-difference error (blue dash-dotted line) and the uncertainty due to imperfect averaging for $\mathcal{C}_D$ and $\mathcal{C}_L$, respectively. Naively, one could choose a $\Delta Re = 1 \times 10^4$ which would have a truncation error of 1.5 percent for $\mathcal{C}_L$. However, to achieve a confidence interval of only 20 percent at the same ΔRe, one would need an estimated total simulation time of 4230 units, which is infeasible and computationally unaffordable. For that reason, it is pertinent to analyze the behavior of the equation (5.1) in order to make an educated decision about what should be the optimal value of ΔRe.

In this section, the word optimal is used to mean the most computationally affordable approach to obtain a similar error percentage in both types of error for a QoI. As it is going to be shown, the optimal value of ΔRe depends on which QoI, the desired precision, and the total simulation time (or computational cost).

Initially, one needs estimations of both types of error and their behavior as functions of the total simulation time T_s^* and ΔRe, which were presented in the last two sections. Next, one needs to decide which is the deciding parameter: precision or computational cost. As shown in Fig. 5.6 there is a trade-off between the uncertainty due to insufficient averaging and T_s^*. For example, if the decision is to set a tolerable error to be 10 percent of the true sensitivity, this would set the total estimation time required to achieve it. For $\mathcal{C}_D$ the optimal value would be around $\Delta Re = 1 \times 10^5$ and would require an estimated total simulation time of 20 units (or running both simulations for at least 10 units of time). However, for $\mathcal{C}_L$, the optimal value is approximately $\Delta Re = 7.5 \times 10^4$ and it will require an estimated $T_s^* = 260$. A similar procedure could be done if the computational cost is the limiting factor. Therefore, the optimal value of ΔRe should be chosen based on the QoI, and a trade-off between the desired precision and computational cost.

On the other hand, in this study, there is interest in understanding the impact of each type of error on the FD-sensitivity. As it is going to be shown next, the two types of error have different effects on the sensitivity profiles of the skin friction and pressure coefficients. For that reason, two values of ΔRe are chosen and two Perturbed Cases are solved. The tolerable levels of error are set to be no more than 10.0 and 20.0 percent for $\mathcal{C}_D$ and $\mathcal{C}_L$, respectively. These levels were chosen based on the computational resources available. Based on the tolerable levels of error, the first value chosen in this study is $\Delta Re = 4 \times 10^4$. This requires a total simulation time of 260 units, evenly split. Under this condition, the second type of error is dominant with an estimated confidence interval of 5.8 percent for $\mathcal{C}_D$ and 19.0 percent for $\mathcal{C}_L$ (the truncation error is 3.4 percent and 5.6 percent, respectively). This case is named Perturbed Case 1 and the parameters are shown in Table 5.2. The second value chosen is $\Delta Re = 1 \times 10^5$. This case is named Perturbed Case 2 and the conditions are listed in Table 5.2. In this case, the truncation error is 9.3 percent and 14.0 percent for $\mathcal{C}_D$ and $\mathcal{C}_L$, respectively. Since the intention for this case is to have a truncation error dominant, one would need a $T_s^* > 160$. However, the simulation time for the Base Case 2 has 130 units, which is longer than needed. In this scenario, equation (5.3) is used to find how much simulation time would be needed from the Perturbed Case 2 to obtain a confidence interval lower than 14.0 percent for this case. Setting the confidence interval to be less than or equal to 10 percent, one would need to run Perturbed Case 2 for 45 units of time.

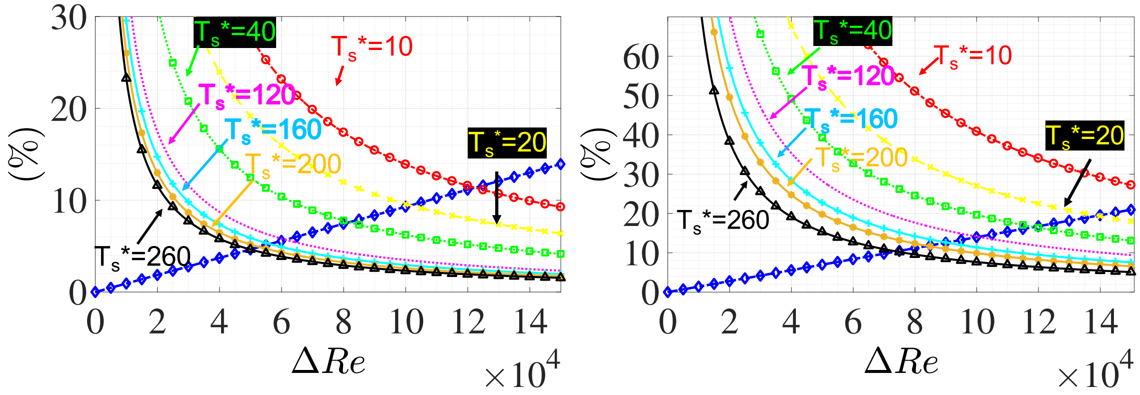

(a) (b)

Figure 5.6: Comparison between the two types of error for (a) $\mathcal{C}_D$ and (b) $\mathcal{C}_L$ of the Base Case 2. The blue dash-dotted lines with diamond symbols represent the finite-difference error. The red dash-dotted lines with circle symbols, the yellow dashed lines with cross symbols, the green dotted lines with square symbols, the magenta dotted lines, the cyan solid lines with plus sign symbols, the orange solid lines with asteric symbols, and the black solid lines with triangle symbols represent the estimation of the uncertainty due to insufficient averaging for $T_s^* = 10$, $T_s^* = 20$, $T_s^* = 40$, $T_s^* = 120$, $T_s^* = 160$, $T_s^* = 200$, and $T_s^* = 260$, respectively.

Table 5.2: Perturbed Cases

Perturbed Case	θ	Re
1	$5°$	4.4×10^5
2	$5°$	5×10^5

5.4 Statistical convergence of the aerodynamic drag and lift coefficients for the Perturbed Cases

Figures 5.7(a) and (b) show the statistics of the aerodynamic coefficients for the Perturbed Cases 1 and 2, respectivey. The dotted curves are the instantaneous aerodynamic coefficients, the black square dots represent the mean for each batch, $\widehat{\mathcal{C}_D}$ and $\widehat{\mathcal{C}_L}$, their running sample mean are shown as the solid thick lines, and their 95 percent confidence interval are plotted as the gray shadow region, as functions of the non-dimensionalized time $t^* = tU_\infty/c$. By looking at them one can see that, as expected, when the Reynolds number increases the aerodynamic drag decreases and the aerodynamic lift decreases. This is confirmed when computing the expected values of $\mathcal{C}_D$ and $\mathcal{C}_L$ in Table 5.3.

Figure 5.8(a) shows the autocorrelation coefficients for the signals shown for the Perturbed Case 1. For a delayed time of $\tau > 0.17$ both signals are weakly- or un-correlated. Figure 5.8(b) shows the standard error for $\mathcal{C}_D$ as a function of the time per batch t_b^* for the same case. There is a region between $0.2 < t_b^* < 0.9$ where the standard error is approximately constant. This means that statistics can be drawn from within the region $0.2 < t_b^* < 0.9$ to guarantee uncorrelated and independent samples. Table 5.3 presents the statistics drawn within this region for $\mathcal{C}_D$ and $\mathcal{C}_L$.

Figures 5.9(a) and (b) show the autocorrelation coefficients and the standard error of $\mathcal{C}_D$ for the Perturbed Case 2, respectively. Similarly, for a delayed time of $\tau > 0.2$ both signals are weakly- or un-correlated. Figure 5.9(b) shows the standard error for $\mathcal{C}_D$ as a function of the time per batch t_b^*. One interesting observation is that the autocorrelation coefficient shows a weak correlation for delayed times $\tau > 0.2$ in all three cases. This might be linked to the aerodynamic profile itself. Table 5.3 presents the statistics of $\mathcal{C}_D$ and $\mathcal{C}_L$ for this case.

Table 5.3: Statistics of $\mathcal{C}_D$ and $\mathcal{C}_L$ for Perturbed Cases 1 and 2.

Case	n_b	$\overline{\mathcal{C}_D}$	$S_{\widehat{\mathcal{C}_D}}$	$S_{\overline{\mathcal{C}_D}}$	$CI_{\overline{\mathcal{C}_D}}$
1	660	0.0154204	9.57×10^{-5}	3.72×10^{-6}	7.31×10^{-6}
2	242	0.0150952	1.08×10^{-4}	6.98×10^{-6}	1.37×10^{-5}
Case	n_b	$\overline{\mathcal{C}_L}$	$S_{\widehat{\mathcal{C}_L}}$	$S_{\overline{\mathcal{C}_L}}$	$CI_{\overline{\mathcal{C}_L}}$
1	660	0.50565421	1.45×10^{-3}	5.67×10^{-5}	1.11×10^{-4}
2	242	0.50581463	1.33×10^{-3}	8.59×10^{-5}	1.69×10^{-4}

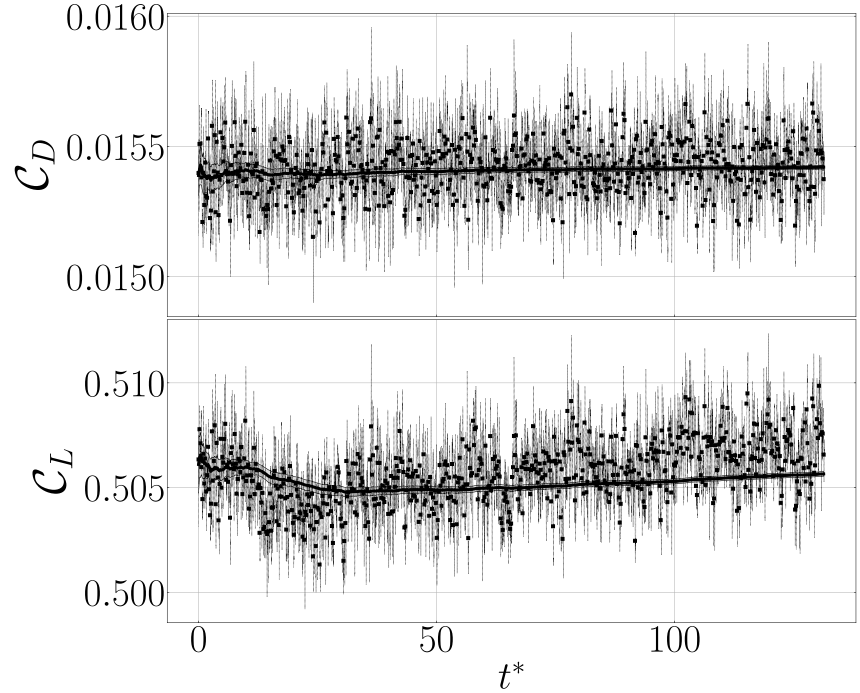

(a) Perturbed Case 1

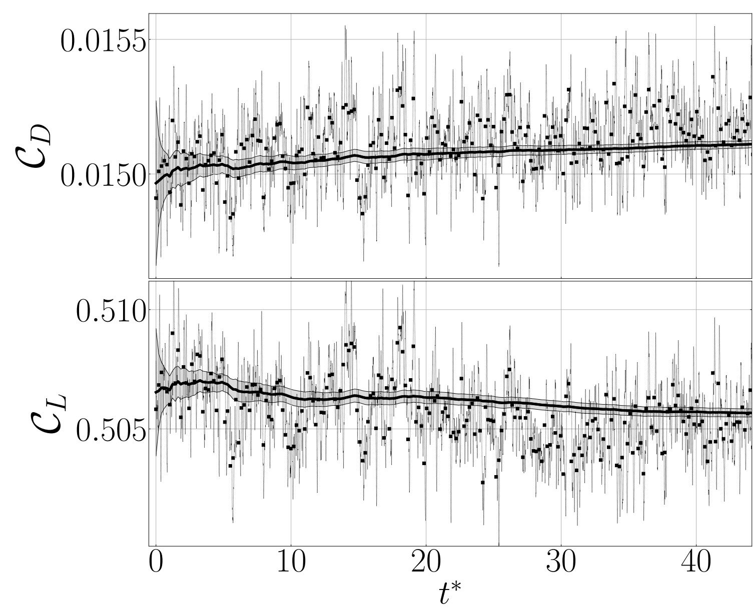

(b) Perturbed Case 2

Figure 5.7: $\mathcal{C}_D$ and $\mathcal{C}_L$ vs t^* for (a) Perturbed Case 1 and (b) Perturbed Case 2. The dash-dotted lines are the instantaneous aerodynamic coefficients, $\mathcal{C}_D$ and $\mathcal{C}_L$, the black square dots represent the mean for each batch, $\widehat{\mathcal{C}_D}$ and $\widehat{\mathcal{C}_L}$, their running sample mean, $\overline{\mathcal{C}_D}$ and $\overline{\mathcal{C}_L}$ are shown as the solid thick lines, and their 95 percent confidence interval are plotted as the gray shadow region.

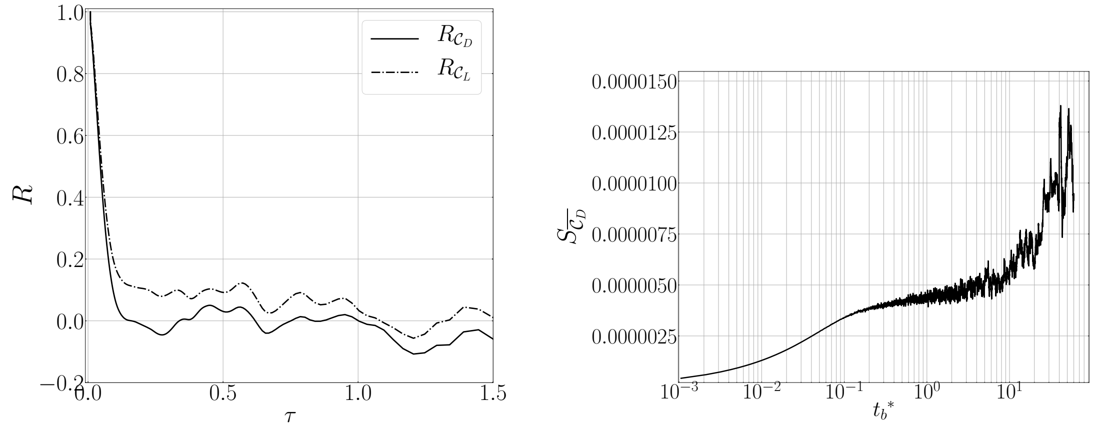

(a) (b)

Figure 5.8: (a) Autocorrelation coefficient for $\mathcal{C}_D$ and $\mathcal{C}_L$; (b) Standard error for $\mathcal{C}_D$ vs t_b. Results for Perturbed Case 1.

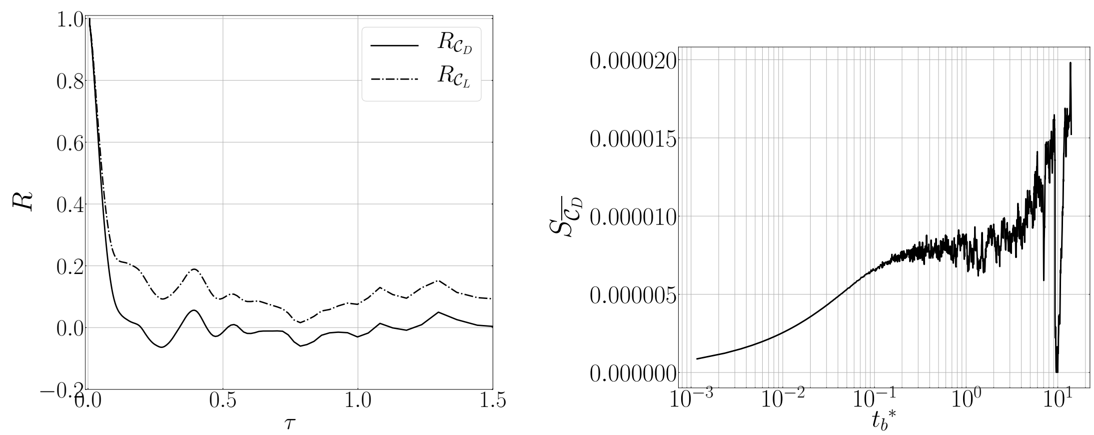

(a) (b)

Figure 5.9: (a) Autocorrelation coefficient for $\mathcal{C}_D$ and $\mathcal{C}_L$; (b) Standard error for $\mathcal{C}_D$ vs t_b. Results for Perturbed Case 2.

5.5 Comparison of the mean quantities between the Base and Perturbed Cases

Before computing the finite-difference sensitivities, it is insightful to compare the skin friction and pressure coefficient profiles obtained for the different cases. Figure 5.10 shows the skin friction (upper) and pressure (lower) coefficients for the Base and Perturbed Cases. The black dash-dotted curves are the results of the Base Case 2, the blue dashed and red dotted curves are the results for the Perturbed Cases 1 and 2, respectively. The curves with triangle symbols represent the upper surface of the airfoil. Several observations are important to highlight. First, one can see the violent effects of the tripping mechanism between $0.05 \leq x/c \leq 0.1$ illustrated by the identical spikes in the skin friction coefficient on the suction region, between $0.05 \leq x/c \leq 0.075$, followed by the sudden decrease in the blowing region. Downstream near the trailing edge, where the flow is the most turbulent, the skin friction follows the expected behavior of decreasing as the Reynolds number increases in both surfaces. However, the most interesting difference is observed in the region $0.1 \leq x/c \leq 0.4$, or the transition region. In this region, the upper surface experiences a sharp transition between $0.1 \leq x/c \leq 0.14$. This transition begins earliest in the Perturbed Case 2, followed by the Perturbed Case 1, and then Base Case 2. The transition observed on the lower surface starts a bit later and occurs more slowly than on the upper surface. The transition on the lower surface occurs between $0.155 \leq x/c \leq 0.22$ for the Perturbed Case 2 and between $0.18 \leq x/c \leq 0.26$ for the Perturbed Case 1. For the Base Case 2, the transition starts at $x/c = 0.22$, and smoothly extends until $x/c \approx 0.4$. A similar transition behavior is observed on the pressure coefficient. There is a sharp transition on the upper surface between $0.1 \leq x/c \leq 0.125$, and a slower and delayed transition on the lower surface between $0.17 \leq x/c \leq 0.23$.

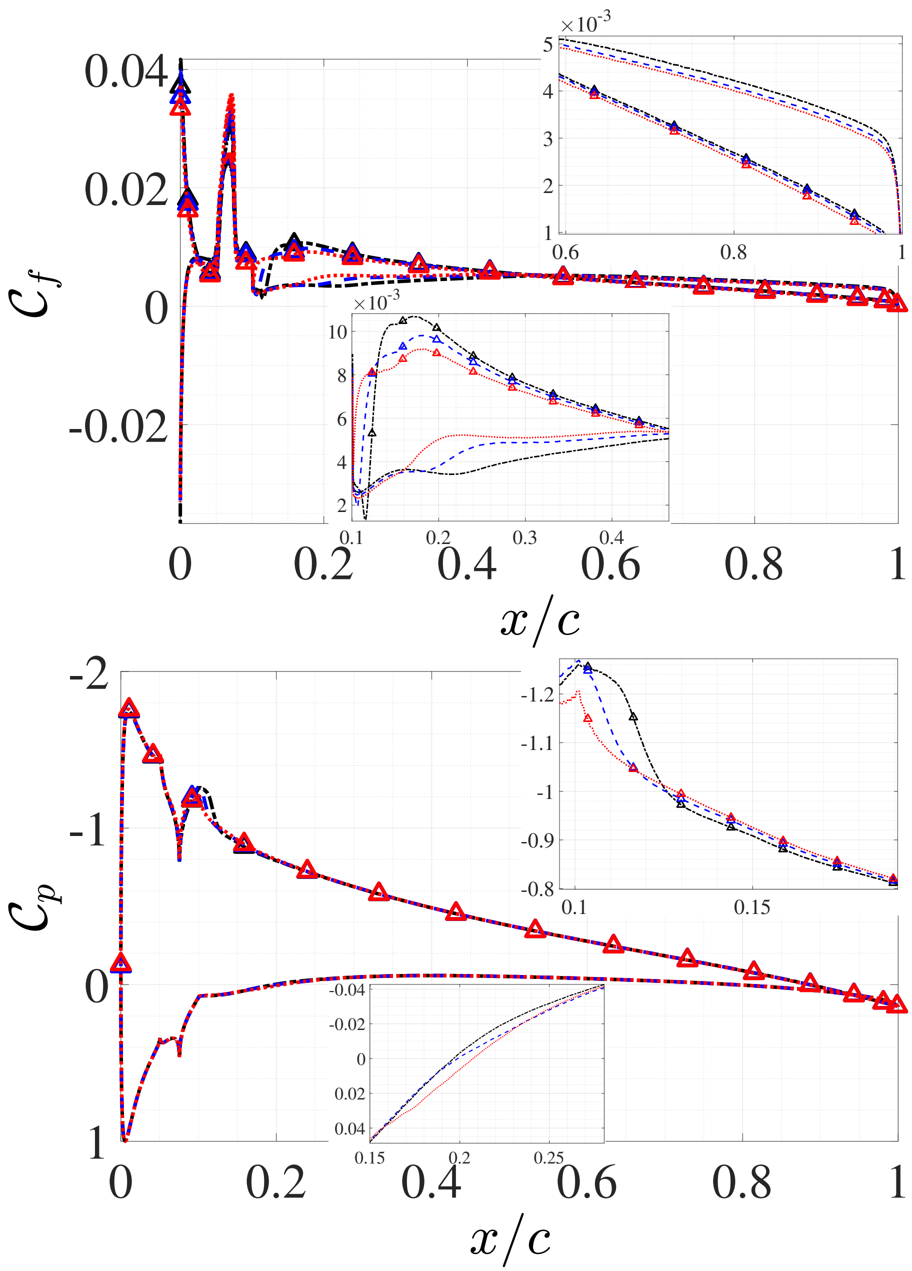

Figure 5.10: Skin friction (upper) and pressure (lower) coefficient along the airfoil surface. The black dash-dotted curve is for the Base Case 2, the blue dashed curve and red dotted curve are the results for the Perturbed Cases 1 and 2, respectively. The curves with triangle symbols represent the upper surface of the airfoil.

Figure 5.11 shows the profiles of the mean velocity and inferred eddy viscosity at different locations of the airfoil with respect to the wall-normal distance. In this figure, the black dash-dotted lines are Base Case 2 results, the blue dashed lines and red dotted lines are the results for the Perturbed Cases 1 and 2, respectively. The important take-aways that can be drawn from these profiles are the expected deceleration of the flow in both surfaces and the increment of the eddy viscosity as it goes downstream. One can observe that the major differences in the velocity profiles are at $x = 0.1$ on the upper surface and $x = 0.3$ on the lower surface, following the same behavior found in the skin-friction coefficient. The big sharp changes shown in the inferred eddy viscosity, especially at $x/c = 0.3$ on the lower surface, are due to the cutoff condition on the denominator of equation (2.16), in which the value is set to be zero when $\overline{S}_{kl}\overline{S}_{kl} \leq$ threshold, to make sure the inferred values are well-defined.

5.6 Finite-difference (FD) sensitivity predictions at Base Case 2

Finally, after deciding what ΔRe to use, running the Perturbed Cases, and computing the QoI's statistics, it is possible to compute the true sensitivity for Base Case 2. This section first presents the sensitivities of the skin friction and pressure coefficients in order to visualize how the truncation and imperfect averaging errors affect the predictions. Next, the sensitivities obtained by LES and RANS are compared for the purpose of detailing the major differences between the models. And last, the sensitivities of the drag and lift coefficients are reported.

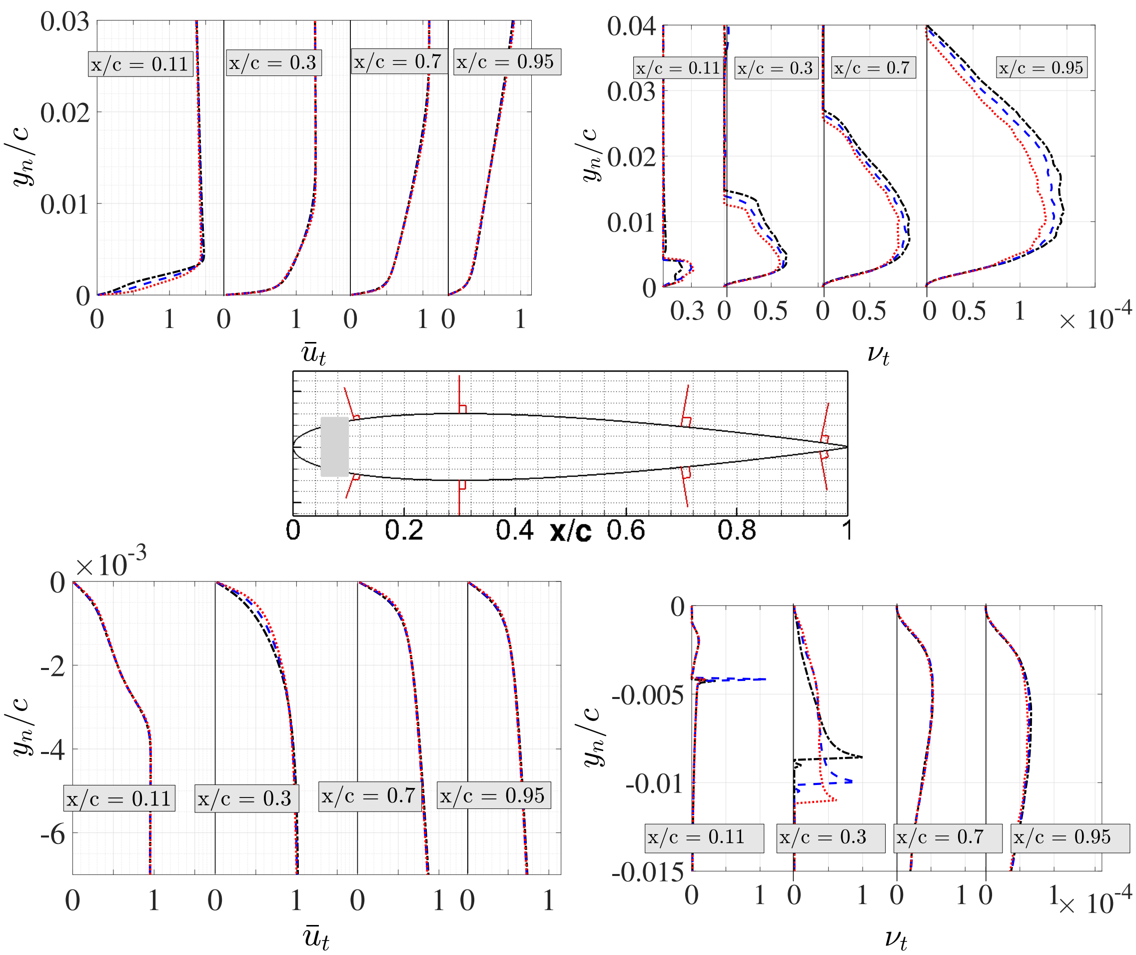

Figure 5.11: Mean velocity and inferred eddy viscosity profiles with respect to the wall-normal distance, y_n/c, at different airfoil locations. The black dash-dotted lines are Base Case 2 results, the blue dashed lines and red dotted lines are the results for the Perturbed Cases 1 and 2, respectively. The gray square on the airfoil profile represents the suction and blowing region.

5.6.1 Effects of the truncation and imperfect averaging errors on the FD-sensitivity predictions

The skin friction and pressure coefficient sensitivities with respect to the Reynolds number are shown in Figs. 5.12 and 5.13, respectively. They are shown as functions of the airfoil chord and split in lower and upper surfaces. The airfoil surface is symmetrically divided by the leading edge location. In here, the cyan dashed curves represent the FD sensitivity using a $\Delta Re = 1 \times 10^5$ and a $T_s^* = 150$, the blue dotted, the orange dash-dotted, and the black solid curves are the FD sensitivity using a $\Delta Re = 4 \times 10^4$ with a $T_s^* = 40$, $T_s^* = 80$, and a $T_s^* = 215$, respectively. The shadow regions around every curve represent the 95 percent confidence interval. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

Several differences are noticeable when the sensitivity is computed using different values of ΔRe. First, the truncation error has a clear effect in the region $0.1 < x/c < 0.5$ on both sides of the airfoil. The sensitivity of the transition point is clearly affected by the ΔRe, having a lower magnitude and shifting upstream on both sides of the airfoil as the ΔRe increases (the shifting being more noticeable on the lower surface). The cyan dashed and black solid $d\mathcal{C}_f/dRe$ curves have their peaks on the upper surface located at $(x/c, d\mathcal{C}_f/dRe) = (0.1151, 6.62 \pm 0.16 \times 10^{-8})$ and $(x/c, d\mathcal{C}_f/dRe) = (0.1164, 12.61 \pm 0.37 \times 10^{-8})$. The difference in the peak's magnitude on the upper surface between the black solid and the cyan dotted curves is of 1.9 times for $d\mathcal{C}_f/dRe$ and 2.0 times for $d\mathcal{C}_p/dRe$, respectively. On the lower surface, the curves' peak shifts upstream 9.5 percent for $d\mathcal{C}_f/dRe$ and 11.0 percent for $d\mathcal{C}_p/dRe$. One could argue that the peak of the transition point sensitivity for the actual $d\mathcal{C}_f/dRe$ and $d\mathcal{C}_p/dRe$ could be located further downstream and have a larger amplitude. If one assumes that the peak's magnitude and location changes

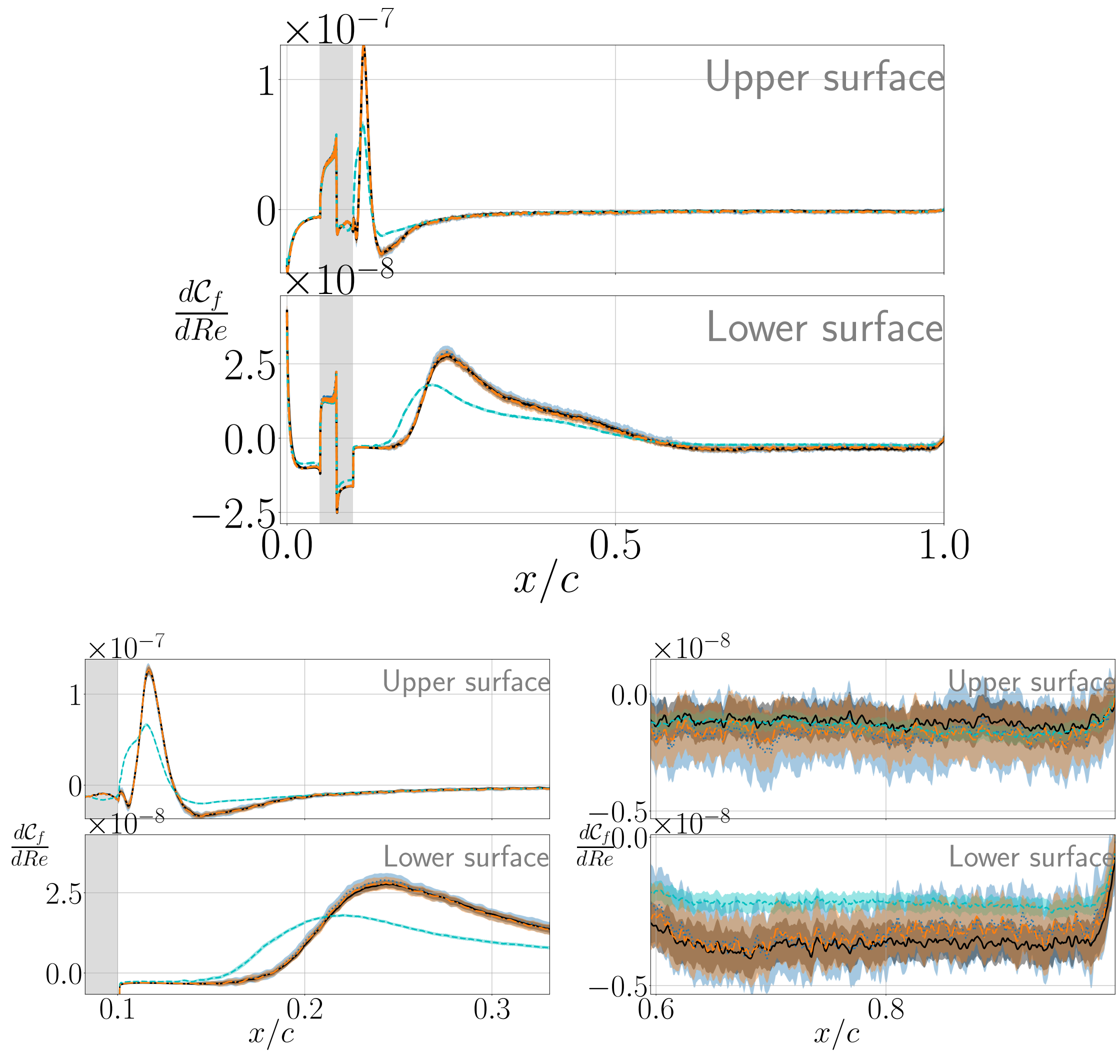

Figure 5.12: Sensitivity of the skin friction coefficient with respect to the Reynolds number using LES. The cyan dashed curves represent the sensitivity using a $\Delta Re = 1 \times 10^5$ and a $T_s^* = 150$, the blue dotted, the orange dash-dotted, and the black solid curves are the sensitivities using a $\Delta Re = 4 \times 10^4$ with a $T_s^* = 40$, $T_s^* = 80$, and a $T_s^* = 215$, respectively. The shadow regions around every curve represent the 95 percent confidence interval. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

linearly with ΔRe (which might not be the case, but useful as a rough estimate), the peak for $d\mathcal{C}_f/dRe$ would be located around $(x/c, d\mathcal{C}_f/dRe) = (0.1172, 16.59\times10^{-8})$.

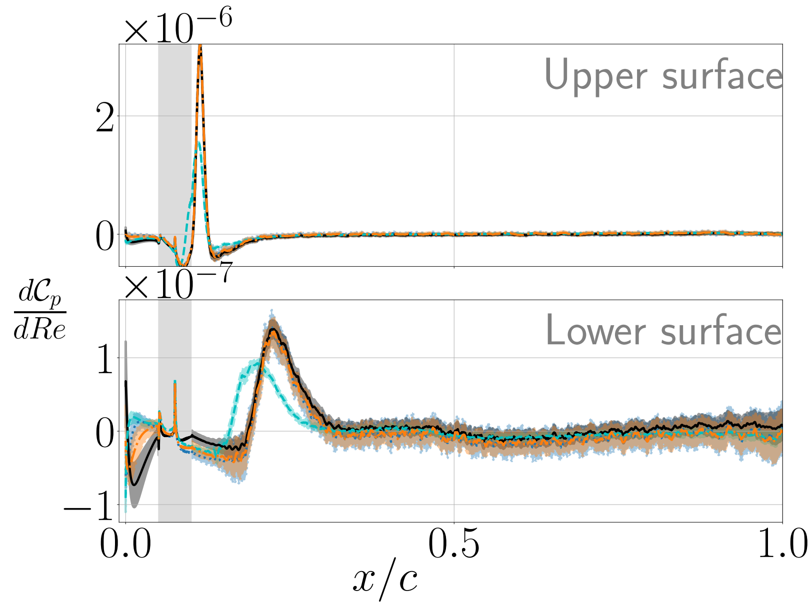

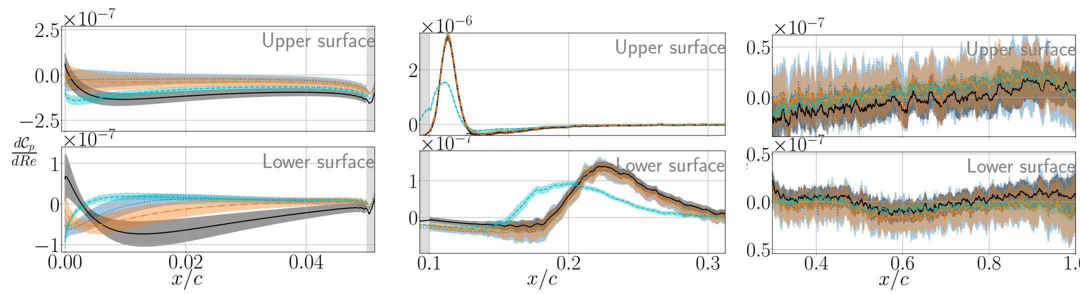

Figure 5.13: Sensitivity of the pressure coefficient with respect to the Reynolds number using LES. The cyan dashed curves represent the sensitivity using a $\Delta Re = 1\times10^5$ and a $T_s^* = 150$, the blue dotted, the orange dash-dotted, and the black solid curves are the sensitivities using a $\Delta Re = 4\times10^4$ with a $T_s^* = 40$, $T_s^* = 80$, and a $T_s^* = 215$, respectively. The shadow regions around every curve represent the 95 percent confidence interval. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

When using a $\Delta Re = 1\times10^5$ and different T_s^* values, there were no major differences on the expected value of $d\mathcal{C}_p/dRe$. However, the insufficient averaging has a noticeable effect on the expected value of $d\mathcal{C}_p/dRe$ at the leading edge (see Fig. 5.13 lower-left) when computing the sensitivity using a $\Delta Re = 4\times10^4$. A less noticeable effect is observed at the trailing edge (see Fig. 5.13 right-left). Interestingly, the lack of averaging did not have major effects on the transition region between $0.1 < x/c < 0.5$.

The black solid results are the definitive FD-sensitivities regarded as the true sensitivities for the skin friction and pressure coefficients which will be used for the remainder of this work. The cyan dashed results are going to be considered to qualitatively estimate the truncation error.

5.6.2 Finite difference sensitivity predictions at Base Case 2 using LES and RANS

Comparing sensitivities between the fully RANS-based method, which is the state-of-the-art in industry, against the finite-difference LES should not be understood as a direct comparison. As explained in chapter 1, one of the purposes of the proposed method is to complement, rather than replace, the state-of-the-art method in situations in which sensitivities with higher fidelity are required within the design process. In other words, one of the objectives of the MFSA method is to predict the features shown by the finite-difference LES sensitivity (which would not be predicted by RANS) at the cost of one LES and RANS simulations. With that clarification, Figs. 5.14 and 5.15 compare the sensitivities of the skin friction and pressure coefficients obtained by LES and RANS. The black solid and cyan dashed curves are the results of sensitivities using LES with a $\Delta Re = 4 \times 10^4$ and a $\Delta Re = 1 \times 10^5$, respectively. The red dash-dotted lines are the finite difference sensitivity obtained using RANS with a $\Delta Re = 50$.

There are several differences worth mentioning between these predictions. First, as expected, the sensitivity of the transition point is predicted differently between the models. In particular, one can see that the sensitivity of the transition point obtained by the k-ω RANS model is delayed on the suction side and predicted at $(x/c, d\mathcal{C}_f/dRe) = (0.1454, 2.127 \times 10^{-8})$ which is located approximately

0.03 chords downstream when compared to LES. Overall, the largest difference between the models is seen between $0.1 \leq x/c \leq 0.2$ on the upper surface and between $0.2 \leq x/c \leq 0.4$ on the lower surface. Second, the LES sensitivities are approximately constant for $d\mathcal{C}_f/dRe$ between $0.6 \leq x/c \leq 0.98$ on both surfaces; however, RANS predictions show a slope that becomes larger approaching the trailing edge.

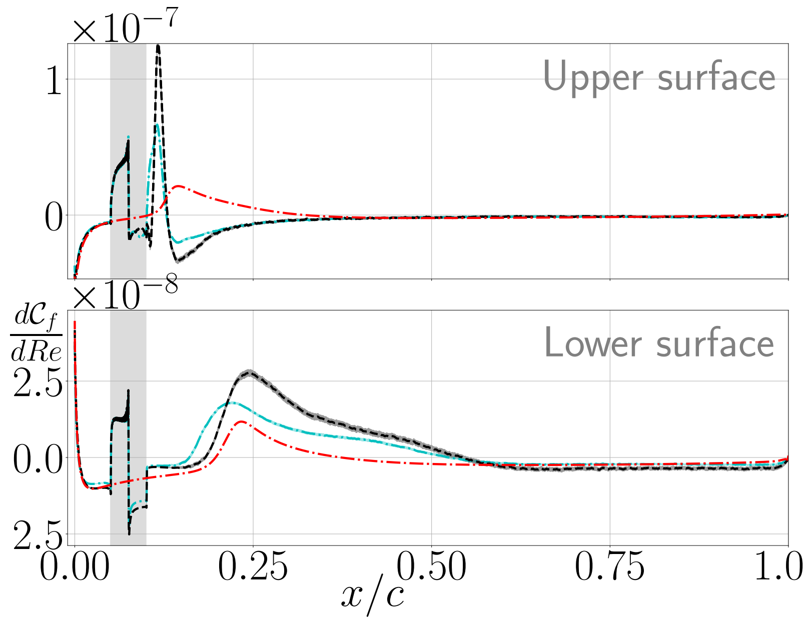

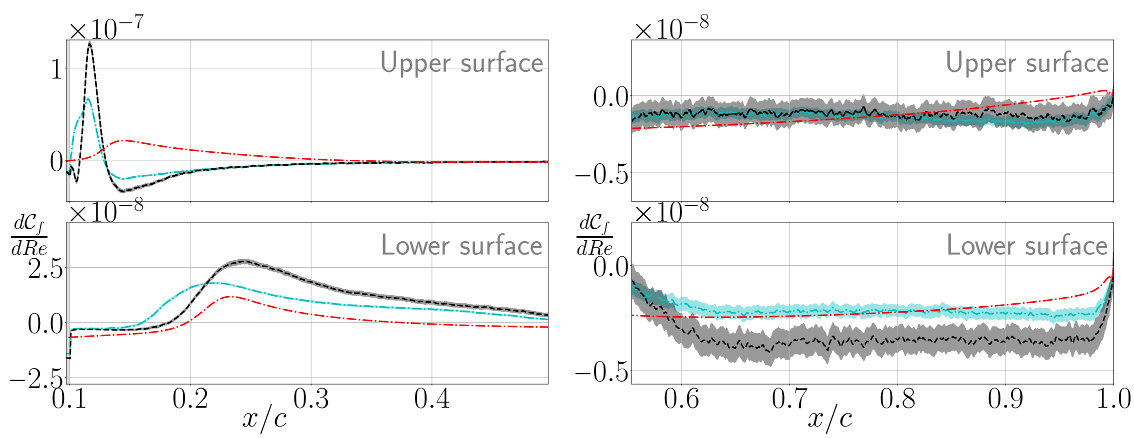

Figure 5.14: Sensitivity of the skin friction coefficient with respect to the Reynolds number using LES and RANS. The black solid and cyan dashed curves are the results of sensitivities using LES with a $\Delta Re = 4 \times 10^4$ and a $\Delta Re = 1 \times 10^5$, respectively. The red dash-dotted curves are the finite difference sensitivity obtained using RANS with a $\Delta Re = 50$. The gray rectangle area between $0.05 < x/c < 0.1$ is where the suction and blowing is applied.

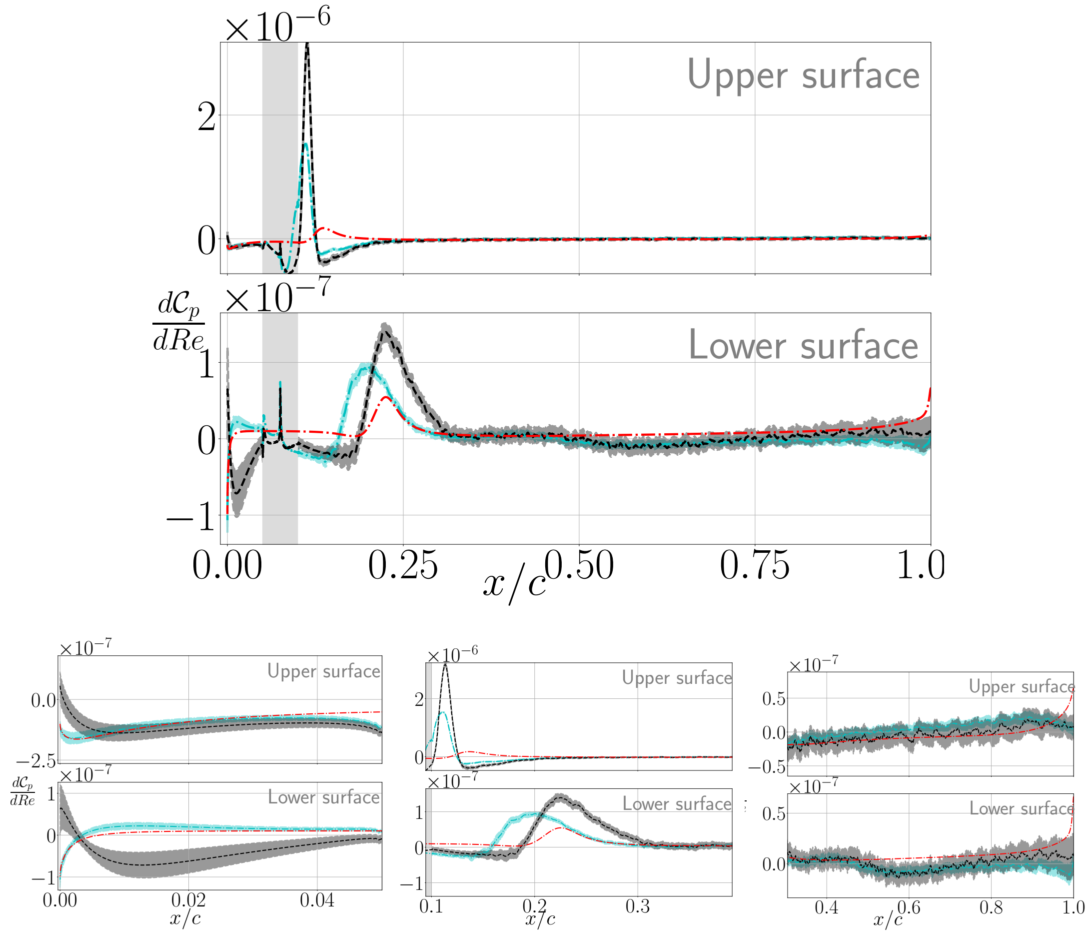

Figure 5.15: FD-sensitivities for the pressure coefficient with respect to the Reynolds number using LES and RANS. The cyan dashed and the black solid lines are the results of the FD-sensitivity using a $\Delta Re = 1 \times 10^5$ and $\Delta Re = 4 \times 10^4$ using LES, respectively. The red dash-dotted curves are the RANS FD-sensitivity and uses a $\Delta Re = 50$. The gray rectangle area between $0.05 < x/c < 0.1$ is where suction and blowing is applied.

Figure 5.16 shows the profiles for the variation in the mean velocity and the inferred eddy viscosity with respect to the Reynolds number, $d\overline{u}_{t,FD}/dRe$ and $d\nu_{t,FD}/dRe$. The profiles are computed as the finite difference between the Perturbed Cases and the Base Case. The black solid and cyan dashed curves represent the difference between the Perturbed Case 1 and Base Case 2 ($\Delta Re = 4 \times 10^4$) and between the Perturbed Case 2 and Base Case 2 ($\Delta Re = 1 \times 10^5$), respectively. The red dash-dotted curves are the RANS finite-difference profiles using a $\Delta Re = 50$.The differences in the profiles of $d\overline{u}_{t,FD}/dRe$ obtained by LES and RANS are aligned with the differences found in $d\mathcal{C}_f/dRe$ and $d\mathcal{C}_p/dRe$. However, the inferred eddy viscosity profiles obtained by LES and RANS show some clear differences. First, near the wall, LES predicts a larger change in the turbulence than RANS in all the profiles. At $x/c = 0.3$ the LES and RANS models predict profiles of $d\nu_{t,FD}/dRe$ with different signs on the upper surface. The big sharp changes shown in the profiles of $d\nu_{t,FD}/dRe$, especially at $x/c = 0.3$ on the lower surface, are due to the cutoff condition given to the denominator in equation 2.16, as explained in the previous section.

Finally, the Table 5.4 presents the sensitivity predictions of the aerodynamic coefficients for Base Case 2 using LES and RANS. Several trends can be observed from this table. First, the “best available estimates” of the sensitivities of the aerodynamic coefficients (first two rows) have a similar predictions for $d\mathcal{C}_L/dRe$, with only a relative difference of 4 percent. However, they report different $d\mathcal{C}_D/dRe$, where the RANS predicts a sensitivity twice of the one obtained by LES. Second, one can see the effects of the truncation error on these sensitivities. The larger the ΔRe the $d\mathcal{C}_D/dRe$ decreases for RANS, however, for LES increases; and the larger the ΔRe the $d\mathcal{C}_L/dRe$ decreases for both models. Third, the truncation error may have a larger effect on the LES predictions than the ones obtained by RANS.

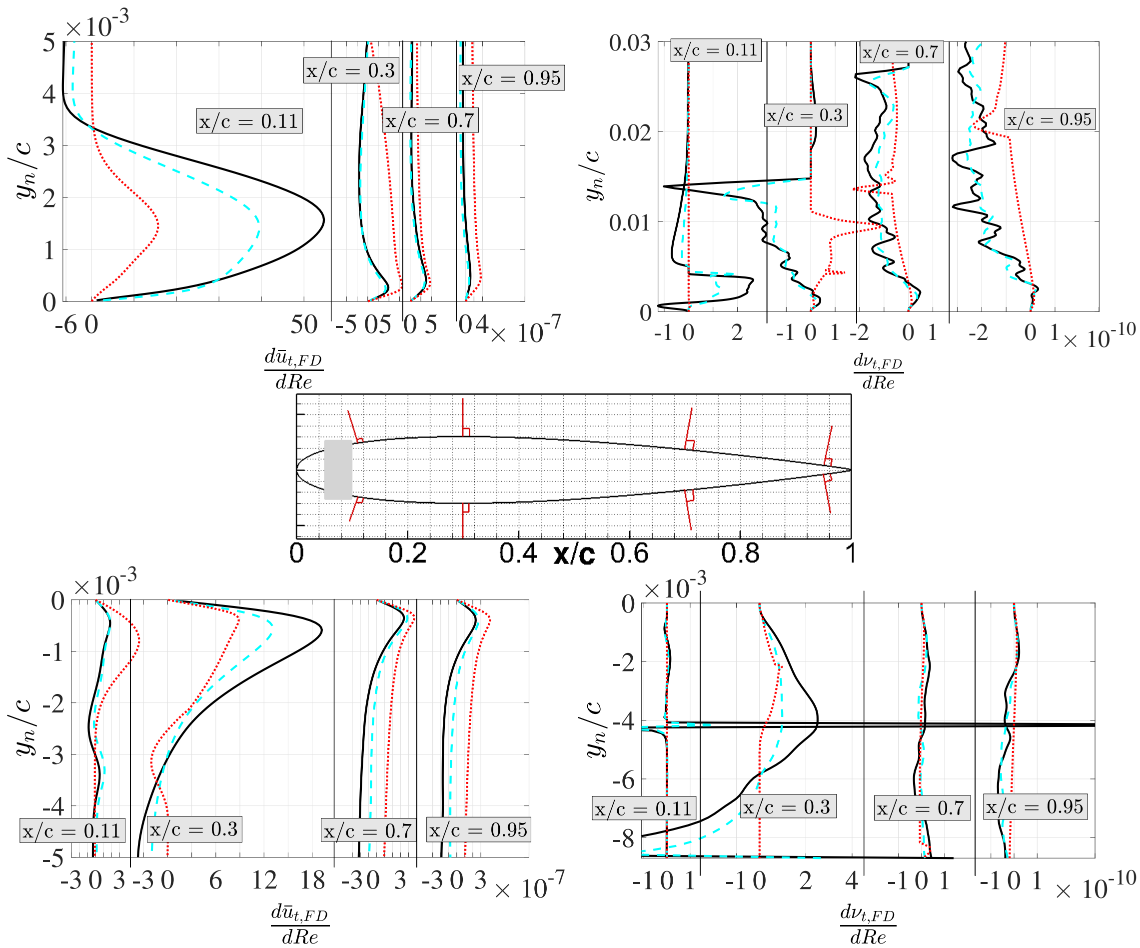

Figure 5.16: Profiles for the variation in the mean velocity and the inferred eddy viscosity with respect to the Reynolds number, $d\bar{u}_{t,FD}/dRe$ and $d\nu_{t,FD}/dRe$, plotted vs wall-normal distance, y_n/c, at different airfoil locations. All profiles are computed by taking the finite-difference between the Perturbed Cases and the Base Case. The black solid and the cyan dashed lines represent the difference between the Perturbed Case 1 and Base Case 2 ($\Delta Re = 4 \times 10^4$) and the difference between the Perturbed Case 2 and Base Case 2 ($\Delta Re = 1 \times 10^5$) using LES, respectively. The red dash-dotted lines correspond to the difference between the Perturbed and Base cases using RANS with a $\Delta Re = 50$.

Table 5.4: Comparison of sensitivity predictions of the aerodynamic coefficients for Base Case 2 using LES and RANS.

Method	$\frac{d\mathcal{C}_D}{dRe} \times 10^9$	$\frac{d\mathcal{C}_L}{dRe} \times 10^8$
FD LES $\Delta Re = 4 \times 10^4$	-2.86±0.27	1.61±0.38
$d\mathcal{C}_D/dRe\|_{lsq}$ and $d\mathcal{C}_L/dRe\|_{lsq}$	-4.994	1.687
FD LES $\Delta Re = 1 \times 10^5$	-4.39±0.15	0.80±0.19
FD RANS $\Delta Re = 50$	-4.990	1.678
FD RANS $\Delta Re = 1 \times 10^5$	-4.53	1.45

Chapter 6: Linearized Reynolds-averaged Navier-Stokes solver

This study makes use of the linearized Reynolds-averaged Navier-Stokes (LRANS) to estimate the sensitivity $\partial J_j/\partial a_i$. As previously said, the tangent equation method is more suited for situations in which there are fewer random parameters a_i than QoIs J_j, $(N_a < N_J)$. This approach could estimate the sensitivities of each of these QoI's with respect to Reynolds numbers by solving N_a simulations of the linear model plus one simulation of the governing model. Furthermore, since the adjoint method is derived from the linearized RANS, the proposed method represents a feasibility study for a multi-fidelity sensitivity analysis in which the adjoint could be used to compute the gradient $\partial J_j/\partial a_i$.

In this section, the formulation, implementation and verification of the LRANS solver is shown for clarity purposes.

6.1 Mathematical formulation

Starting from the mean LES equations (2.12), one can perturb the mean solution about the base flow, i.e., $\overline{u}_i \to \overline{u}_i + \delta\overline{u}_i$, $\overline{p} \to \overline{p} + \delta\overline{p}$. Additionally, since the Reynolds numbers is used as the random parameter, the kinematic viscosity is perturbed as $\nu \to \nu + \delta\nu$. The perturbed mean LES equations can be linearized to

$$\frac{\partial \delta\overline{u}_j}{\partial x_j} = 0\,,$$

$$\frac{\partial \delta\overline{u}_i}{\partial t} + \overline{u}_j\frac{\partial \delta\overline{u}_i}{\partial x_j} + \delta\overline{u}_j\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \delta\overline{p}}{\partial x_i} - \frac{\partial}{\partial x_j}\left[2\nu\delta\overline{S}_{ij} + 2\delta\nu\overline{S}_{ij} - \delta\overline{\tau}_{ij} - \delta\overline{u_i'u_j'}\right] = 0\,, \tag{6.1}$$

these equations are still exact, but suffer from a closure problem just like standard RANS equations. Specifically, we need to model the perturbed mean Reynolds plus subgrid stress, which is accomplished in this work using an eddy viscosity approach.

The eddy viscosity hypothesis for the full Reynolds stress tensor is

$$\overline{u_i' u_j'} = -2\nu_t \overline{S}_{ij} + \frac{2k}{3}\delta_{ij}\,, \tag{6.2}$$

and $k = \overline{u_k' u_k'}/2$ is the turbulence kinetic energy. Linearization yields

$$\delta\overline{u_i' u_j'} = -2\nu_t \delta\overline{S}_{ij} - 2\overline{S}_{ij}\delta\nu_t + \frac{2\delta k}{3}\delta_{ij}\,. \tag{6.3}$$

We absorb the perturbed mean subgrid stress into the model of the perturbed Reynolds stress (or, equivalently, we neglect it), and also neglect the perturbed turbulence kinetic energy. This then yields equation (6.1) for infinitesimal perturbations of the mean equations as

$$\begin{aligned} \frac{\partial \delta\overline{u}_j}{\partial x_j} &= 0\,, \\ \frac{\partial \delta\overline{u}_i}{\partial t} + \overline{u}_j \frac{\partial \delta\overline{u}_i}{\partial x_j} + \delta\overline{u}_j \frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \delta\overline{p}}{\partial x_i} - \frac{\partial}{\partial x_j}\Big[2(\nu+\nu_t)\delta\overline{S}_{ij} + 2(\delta\nu + \delta\nu_t)\overline{S}_{ij}\Big] &= 0\,. \end{aligned} \tag{6.4}$$

This requires the specification of the eddy viscosity ν_t and its perturbation $\delta\nu_t$. Using an existing RANS turbulence model for ν_t is not appealing since that would produce an eddy viscosity that is inconsistent with the LES mean solution. We therefore use the procedure described in [31] to compute the ν_t field inferred from the LES mean solution $\overline{\mathbf{v}}$ as

$$\nu_t = -\frac{\overline{S}_{ij}\overline{u_i' u_j'}}{2\overline{S}_{ij}\overline{S}_{ij}}. \tag{6.5}$$

The only term left to predict for closing the equation system (6.4) is the perturbartion of the eddy viscosity, which models the change in turbulence. In this work, this term is either assumed to be zero (the "frozen turbulence" assumption), computed as the first-order finite-difference of the eddy viscosity fields at two different points in the parametric space as

$$\frac{\delta \nu_{t,FD}}{\delta \nu} = \frac{\nu_t(\overline{\mathbf{v}}(Re_p)) - \nu_t(\overline{\mathbf{v}}(Re_b))}{\nu_p - \nu_b} \tag{6.6}$$

and given as input when solving the equation system (6.4), or modeled using Prandtl's zero-equation model. In the particular case where the Reynolds number is the perturbed parameter, and κ and d are constant, the eddy viscosity perturbation is given by

$$\delta \nu_t = \nu_t \left[\frac{\overline{S}_{ij} \delta \overline{S}_{ij}}{\overline{S}_{kl} \overline{S}_{kl}} + \frac{d \,\mathrm{sgn}(\tau_w) \exp\left[\frac{-d^+}{A^+}\right] \delta \tau_w}{\nu A^+ \sqrt{D|\tau_w|}} - \frac{d^+ \exp\left[\frac{-d^+}{A^+}\right] \delta \nu}{\nu A^+ \sqrt{D}} \right] \quad , \tag{6.7}$$

where d is the wall-normal distance, $d^+ = u_\tau d/\nu$, D is the Van Driest damping function, A^+ is the damping function constant, and τ_w is wall shear stress. The eddy viscosity in equation (6.7) is computed as

$$\nu_t = \kappa^2 d^2 \left| \overline{S} \right| D \quad . \tag{6.8}$$

The finite-difference eddy viscosity perturbation, $\delta \nu_{t,FD}$, is used in two situations. First, to verify the LRANS solver. Second, to perform the proof-of-concept of the proposed MFSA method in chapter 7. The linearized equation (2.17) is denoted by the operator

$$\mathcal{L}(\delta \overline{\mathbf{v}}, \delta \nu_t; \overline{\mathbf{v}}, \nu_t, \mathbf{a}, \delta \mathbf{a}) = 0 \,.$$

The linearized RANS equation is implemented in the same software framework as used for the LES and RANS.

6.2 Verification: LRANS of Base Case 2

To verify the LRANS solver, its results are compared against results obtained by the RANS solver, already verified. Two RANS at $\theta = 5^o$ with different Reynolds numbers are solved, $\mathcal{R}(\overline{\mathbf{v}}, \nu_t; \theta, Re_b) = 0$ and $\mathcal{R}(\overline{\mathbf{v}}, \nu_t; \theta, Re_p) = 0$, and the finite-difference sensitivity for all QoIs is computed as

$$\left.\frac{\partial J}{\partial Re}\right|_{FD} = \frac{J(\overline{\mathbf{v}}(Re_p)) - J(\overline{\mathbf{v}}(Re_b))}{\Delta Re}, \tag{6.9}$$

and compared against the sensitivity obtained by the LRANS solver as

$$\left.\frac{\partial J}{\partial Re}\right|_{\overline{\mathbf{v}}(Re_b)} = \lim_{\delta Re \to 0} \left[\frac{\left\langle \left.\frac{\partial J}{\partial u}\right|_{\overline{\mathbf{v}}(Re_b)}, \delta u \right\rangle}{\delta Re} + \frac{\left\langle \left.\frac{\partial J}{\partial \nu}\right|_{\overline{\mathbf{v}}(Re_b)}, \delta \nu \right\rangle}{\delta Re} \right] \tag{6.10}$$

6.2.1 LRANS verification: k-ω eddy viscosity model as benchmark

The sensitivities of the skin friction, pressure, drag and lift coefficients are predicted using the k-ω eddy viscosity model. Two RANS simulations are solved using the grid g_{RANS3} from Table 4.1 at $\theta = 5^o$ with $Re_b = 4 \times 10^5$ and $Re_p = 400025$. Two LRANS simulations are solved using the same g_{RANS3} grid and different δRe, to verify its linearity. A comparison of QoIs predictions using RANS k-ω and LRANS are shown in Table 6.1 and Figure 6.1. The aerodynamic coefficients reported in this table from the two different approaches have less than a 0.1 percent relative difference. In Figure 6.1, the black solid lines are the FD-sensitivity from RANS k-ω,which in this particular case represents the true values, the red dashed blue dotted lines represent the LRANS solution using $\delta Re = 50$ and $\delta Re = 25$, respectively. Two important observations can be noted. First, the LRANS results agree with the finite-difference results. Second, the LRANS results (properly non-dimensionalized) are independent of the value of the perturbation, δRe, indicating linearity.

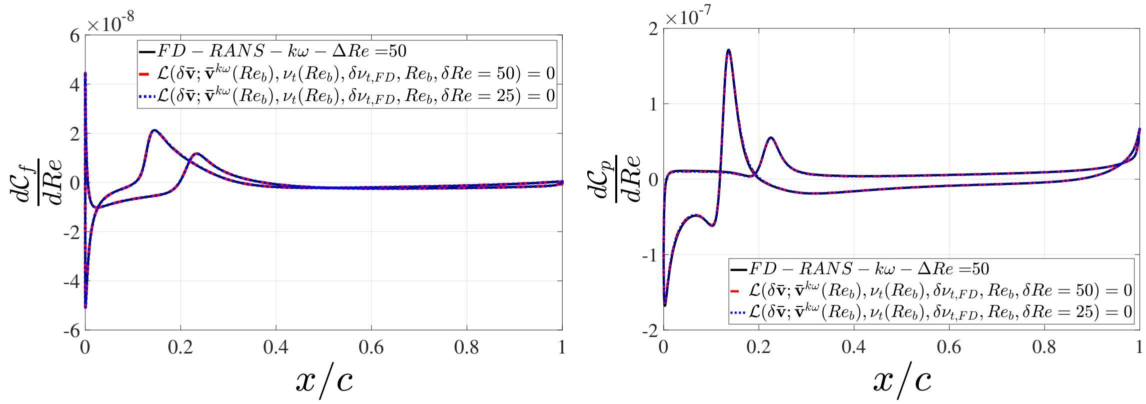

(a) (b)

Figure 6.1: LRANS verification - QoI values along the airfoil chord using the Wilcox $k\omega$ eddy viscosity model. (a) sensitivity of the skin friction coefficient. (b) sensitivity of the pressure coefficient.

Table 6.1: LRANS verification using k-ω eddy viscosity as benchmark. Sensitivity of $d\mathcal{C}_D/dRe$ and $d\mathcal{C}_L/dRe$ for Base Case 2.

Method	$\frac{d\mathcal{C}_D}{dRe}\Big\vert_{Re_b}$	$\frac{d\mathcal{C}_L}{dRe}\Big\vert_{Re_b}$
$\mathcal{R}(\overline{\mathbf{v}}^{k\omega}, \nu_t; \theta, Re_p) - \mathcal{R}(\overline{\mathbf{v}}^{k\omega}, \nu_t; \theta, Re_b) = 0$ $\Delta Re = 50$	-4.990×10^{-9}	1.678×10^{-8}
Surrogate RANS (m_{C_D} and m_{C_L})	-4.994×10^{-9}	1.687×10^{-8}
$\mathcal{L}(\delta\overline{\mathbf{v}}; \overline{\mathbf{v}}^{k\omega}(Re_b), \nu_t(Re_b), \delta\nu_{t,FD}, Re_b, \delta Re) = 0$	-4.993×10^{-9}	1.6766×10^{-8}

Previous results show that the LRANS solver can retrieve the correct sensitivity for all QoIs. As a clarification, in this section, the FD eddy viscosity perturbation, $\delta\nu_{t,FD}$, is used. Although it was helful, it is necessary a model to predict it.

6.2.2 LRANS Verification: Prandtl's eddy viscosity model as benchmark

This thesis represents a first attempt towards the prediction of a computationally affordable sensitivity in LES. In that sense, the criteria to select a first model to predict the change in turbulence due to variations in the design parameters are that is computationaly cheap solving it and that is "simple-enough" to implementing it. In this attempt, the eddy viscosity model based on Prandtl's mixing length theory is chosen since meets the requirements. This model does not require the solution of an extra transport equation and its implementation in LRANS, already verified, is not instrusive.

This section presents the process to verify the eddy viscosity perturbation model based on Prandtl's mixing length theory. This algebraic model is going to be used in the proposed MFSA method. A similar verification process as the one shown in the previous section is presented here. FD-sensitivity predictions from the RANS solver using Prandtl's eddy viscosity model are compared against the LRANS solver using the eddy viscosity perturbation model based on Prandtl's. The development, as previouly said, is for a density constant, and incompressible flow.

The eddy viscosity model using Prandtl's mixing length hypothesis is

$$\nu_t = \kappa^2 d^2 \, |S| \, D \quad , \tag{6.11}$$

In the particular case where the Reynolds number is the perturbed parameter, and κ and d are constant, the eddy viscosity perturbation is given by:

$$\delta\nu_{t,Pr} = \nu_t \left[\frac{\overline{S}_{ij}\delta\overline{S}_{ij}}{\overline{S}_{kl}\overline{S}_{kl}} + \frac{d\,\mathrm{sgn}(\tau_w)\exp\left[\frac{-d^+}{A^+}\right]\delta\tau_w}{\nu A^+\sqrt{D|\tau_w|}} - \frac{d^+\exp\left[\frac{-d^+}{A^+}\right]\delta\nu}{\nu A^+\sqrt{D}} \right] \quad , \tag{6.12}$$

A comparison of the skin friction and pressure coefficients predictions using RANS and LRANS is presented in Figure 6.1. Similarly, profiles of the eddy viscosity change in the wall-normal direction are shown in Figure 6.3. In this figures, the blue solid lines are the FD-sensitivity from RANS, and the red dashed lines are the results obtained by the LRANS solver using the equation (6.7) to predict $\delta\nu_t$. From both figures one can note to observations. First, one can note that the LRANS results agree with the finite-difference results. Second, and most importantly, the LRANS solver correctly predicts the change of the eddy viscosity.

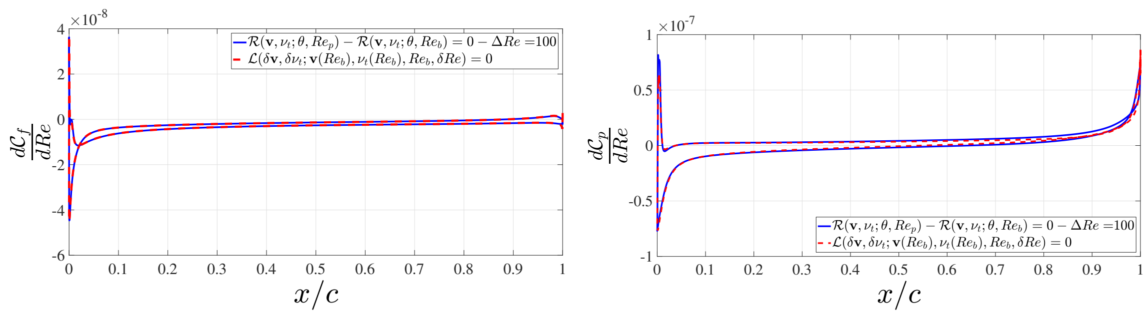

(a) (b)

Figure 6.2: LRANS verification - QoI values along the airfoil chord using the Prandtl's mixing length eddy viscosity model. (a) sensitivity of the skin friction coefficient. (b) sensitivity of the pressure coefficient.

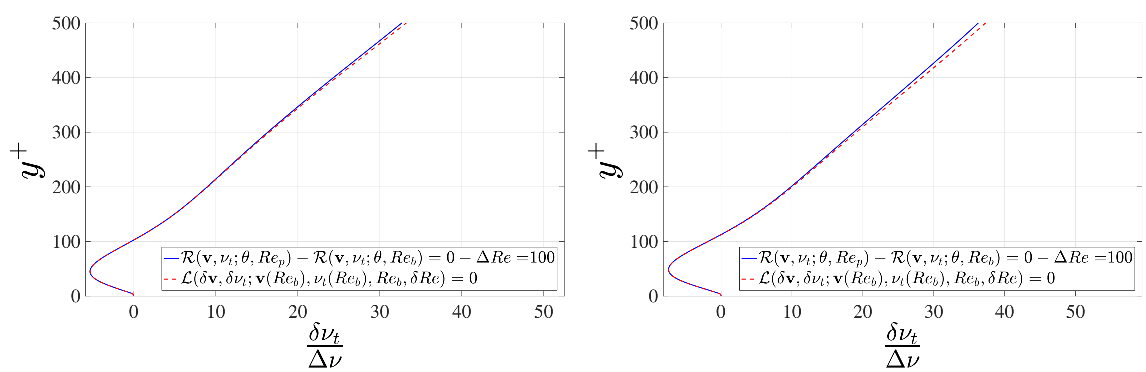

(a) (b)

Figure 6.3: LRANS verification - eddy viscosity perturbation model using the Prandtl's mixing length. (c) $\delta\nu_t/\Delta\nu$ profile at $x = 0.6$. $\delta\nu_t/\Delta\nu$ profile at $x = 0.9$

Chapter 7: Results

This thesis represents the first attempt towards the prediction of a computationally affordable sensitivity of an LES using the multi-fidelity sensitivity analysis (MFSA). For that reason, it was not fully known whether the coupling of the different fidelity models would be able to predict the main hydrodynamic features of the QoI's sensitivity. Similarly, it was assumed that the inferred eddy viscosity was the only term important in the system of equations (6.4), partly because the literature review showed that in standard sensitivity analysis it is common to assume a constant eddy viscosity when perturbing the flow, and partly because the importance of the perturbed eddy viscosity was not fully understood.

This chapter presents the systematic approach followed in this research to assess the proposed method. The sensitivity predictions obtained by the benchmark and the fully RANS-based method, which is the state-of-the-art in industry, are leveraged to: first, prove the feasibility of the MFSA; second, quantify the accuracy of the MFSA when different models for closure are used; and third, identify the possible sources of error.

In all plots, the black solid curves show the best available estimate of the true sensitivity, computed using a finite-difference approximation in parameter space from two different LES cases spaced $\Delta Re = 4 \times 10^4$ apart. The cyan dashed curves show the finite-difference LES sensitivity using a larger $\Delta Re = 1 \times 10^5$, shown here to provide a sense of the truncation error in the "true" sensitivity. The shaded regions around these curves represent the 95 percent confidence interval due to the imperfect averaging.

7.1 Proof-of-concept

Mathematically, the proposed method should be able to predict the exact sensitivities provided the exact infinitesimal perturbed mean Reynolds stress, $\delta\overline{u_i'u_j'}$. In this first attempt, $\delta\overline{u_i'u_j'}$ is modeled using

$$\delta\overline{u_i'u_j'} = -2\nu_t\delta\overline{S}_{ij} - 2\overline{S}_{ij}\delta\nu_t + \frac{2\delta k}{3}\delta_{ij}\,, \tag{7.1}$$

where ν_t is approximated by minimizing the error between the deviatoric part of the Reynolds stress tensor given by high- and low-fidelity models using

$$\nu_t = -\frac{\overline{S}_{ij}\overline{u_i'u_j'}}{2\overline{S}_{kl}\overline{S}_{kl}}, \tag{7.2}$$

and δk is assumed to be close to zero and therefore in this analysis is neglected. However, there is no direct solution for $\delta\nu_t$. This section checks the feasibility of the MFSA and asks how close its predictions are to the true values if the "best available estimate" for the perturbed eddy viscosity is given as input. Within the context of the eddy viscosity constitutive model, the "best available estimate" to the infinitesimal perturbed eddy viscosity is established as the finite-difference of the inferred eddy viscosity from the LES of Perturbed Case 1 and Base Case 2 ($\delta\nu_t \approx \delta\nu_{t,FD}$), and is computed as

$$\frac{\delta\nu_{t,FD}}{\delta\nu} = \frac{\nu_t(\overline{\mathbf{v}}(Re_p)) - \nu_t(\overline{\mathbf{v}}(Re_b))}{\nu_p - \nu_b}. \tag{7.3}$$

This field is then given as an input to the LRANS model to solve for the solution perturbation $\delta\overline{\mathbf{v}}$. It is important to notice that this approach also carries a truncation error. However, it is not possible to quantify its impact because there is not access to the infinitesimal perturbed eddy viscosity.

Figures 7.1 and 7.2 compare the sensitivities of the skin friction and pressure coefficients with respect to the Reynolds number using different methods. The green solid curves are the MFSA prediction using $\delta\nu_{t,FD}$ for modeling closure. The most important observation that can be drawn from these figures is that the MFSA method is able to predict the general features shown on the true predictions in both sensitivities. In other words, provided an accurate model for the perturbed eddy viscosity, the proposed method predicts correctly the different regions on the airfoil like the laminar region before the tripping mechanism, the sensitivity of the transition point after the tripping mechanism, the approximately constant behavior downstream between $0.4 < x/c < 0.9$, and even the sharp changes due to the suction and blowing mechanism. However, despite the use of $\delta\nu_{t,FD}$, there are some discrepancies between the MFSA and the true sensitivities. The first, and most noticeable, difference is the delayed start in the sensitivity of the transition point (difference between the green and black lines in the region $0.1 < x/c < 0.2$ in the upper surface, and in the region $0.2 < x/c < 0.4$ in the lower surface in Figs. 7.1 and 7.2). The MFSA predictions have a sharp decrement in the sensitivity of the skin friction before transitioning, a behavior that the black solid line shows much less (see Fig. 7.1 lower left). Second, there is a magnitude difference in the sensitivity of the transition point between the black and green solid lines. The MFSA under predicts the magnitudes in three of the four points (the exception is on the lower surface for the skin friction sensitivity in Fig. 7.1). Third, the MFSA using $\delta\nu_{t,FD}$ captures the behavior and magnitude of the skin friction sensitivity at the leading edge and trailing edge in Fig. 7.1. Fourth, although the proposed method captures the behavior and magnitude of the pressure coefficient sensitivity on the upper surface leading and trailing edges, the predictions on the lower surface have differences at the leading edge and towards the end of the trailing edge in Fig. 7.2.

However, these discrepancies between the green and black solid lines could be due to either the limitations of the true sensitivity predictions (truncation error) or the use of the eddy viscosity hypothesis. In order to separate both types of errors in the differences between the green and black solid lines, the following section presents an additional analysis. The inherent truncation error in $\delta\nu_{t,FD}$ might have an effect on the predictions. However, it is not possible to quantify it and therefore is going to be added to the imperfect modeling error.

7.1.1 Error analysis: imperfect modeling of unclosed terms

To quantify the incurred error when using the eddy viscosity model, one would solve the MFSA provided the exact infinitesimal perturbed mean Reynolds stress, and quantify the difference of the results obtained with the ones shown in Figs. 7.1 and 7.2. However, there is no access to the exact perturbed mean Reynolds stress. For that reason, this term is approximated as the finite-difference from the LES of the Perturbed Case 1 and Base Case 2 as

$$\frac{\delta\overline{u_i'u_j'}_{FD}}{\delta\nu} = \frac{\overline{u_i'u_j'}(\bar{\mathbf{v}}(Re_p)) - \overline{u_i'u_j'}(\bar{\mathbf{v}}(Re_b))}{\nu_p - \nu_b}, \tag{7.4}$$

and given as input to the LRANS solver. Although $\delta\overline{u_i'u_j'}_{FD}$ is not exact, it is the "best available estimate" to the infinitesimal value in this study without using any constitutive model.

Figures 7.3 and 7.4 show similar plots of the sensitivities of skin friction and pressure coefficients with respect to the Reynolds number. By comparing these curves, one might be able to differentiate the errors associated to the imperfect modeling and the inherent errors on the true sensitivities (truncation and imperfect averaging errors). The first and most noticeable observation is that, when comparing the green and brown curves at $0.1 < x/c < 0.2$ on the upper surface, and at

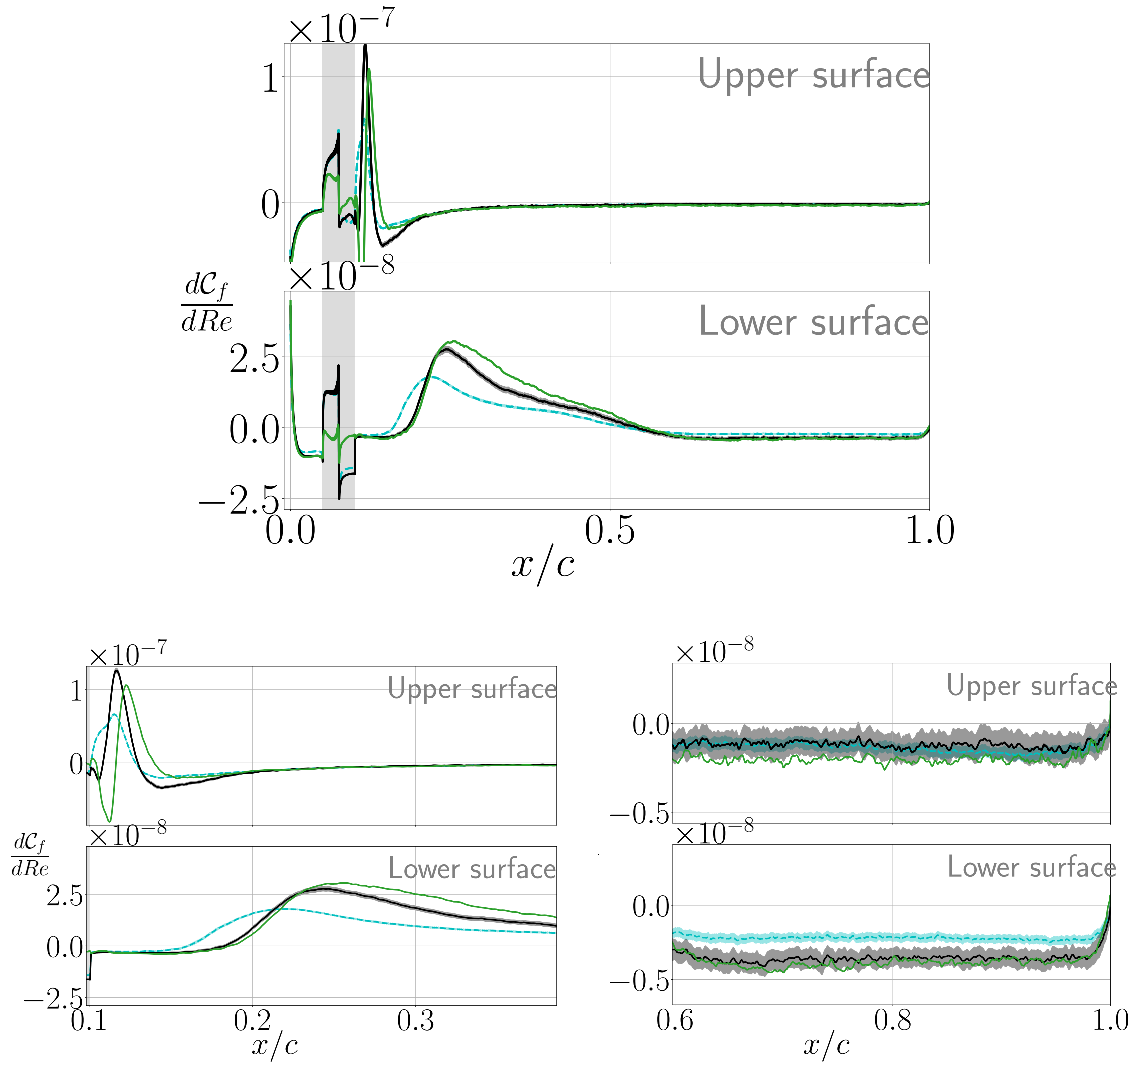

Figure 7.1: Sensitivity of the skin friction coefficient with respect to the Reynolds number. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The cyan dashed curve is the same sensitivity but with a larger $\Delta Re = 1 \times 10^5$. The shaded region around each of these curves represent the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

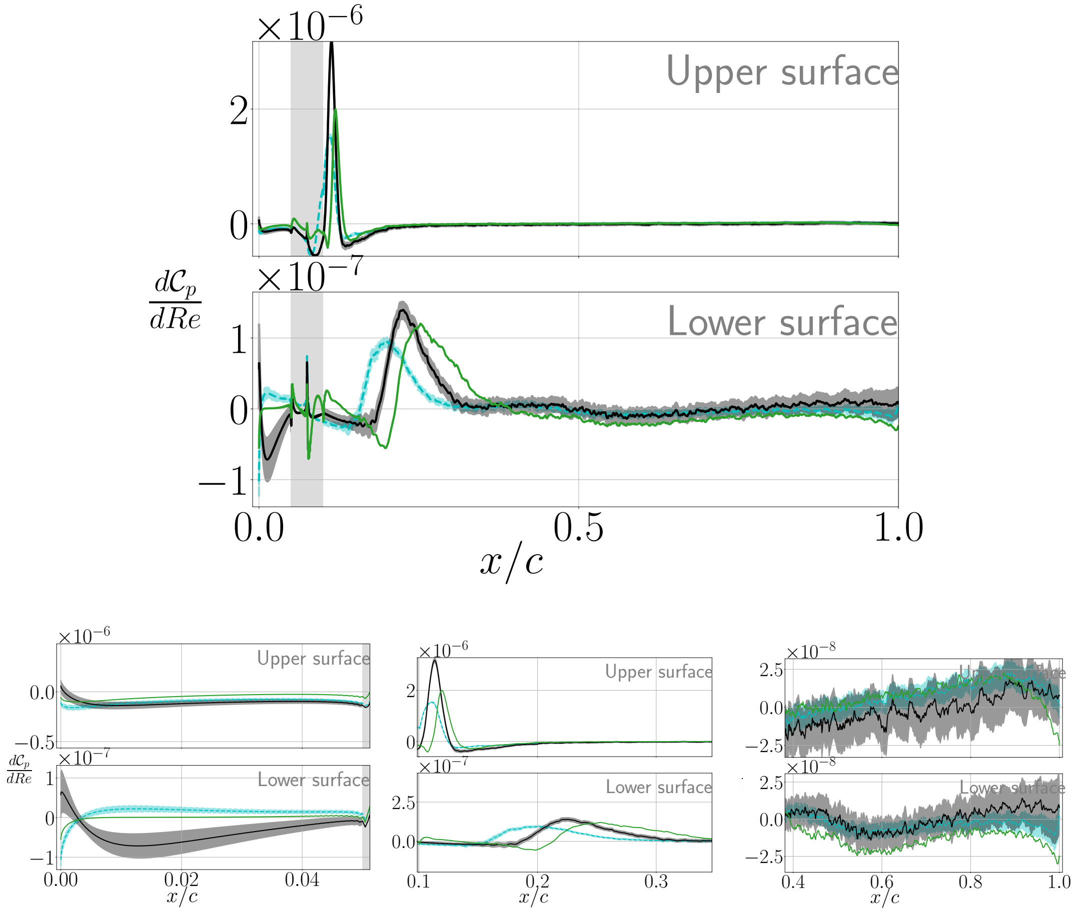

Figure 7.2: Sensitivities for the pressure coefficient with respect to the Reynolds. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The cyan dashed curve is the same sensitivity but with a larger $\Delta Re = 1 \times 10^5$. The shaded region around each of these curves represent the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

$0.2 < x/c < 0.4$ on the lower surface for both sensitivities, one can see that the delayed start in the sensitivity of the transition point may be due to the eddy viscosity model approximation. Second, the sharp decrement in the sensitivity of the skin friction before transitioning (green line in between $0.1 < x/c < 0.15$ on the upper surface in Fig. 7.3) seems to be due to the imperfect modeling.

Figure 7.5 shows the profiles for the variation in the mean velocity with respect to the Reynolds number, $d\overline{u}_{t,FD}/dRe$, at different airfoil locations. The color scheme is the same as the previous plot. First, by looking at the location $x/c = 0.1$ on the upper surface, the mean perturbed velocity obtained by the proposed method has the opposite sign to the one obtained by the rest of the predictions, including the true sensitivity. This aligns with the negative sensitivity in the skin friction on the same location shown in Fig. 7.3. Overall, Fig. 7.5 shows that, with the exception of the $x/c = 0.1$ location, the proposed MFSA method is capable of predicting the sensitivity of the mean perturbed velocity with respect to the Reynolds number, provided an accurate closure modeling.

7.1.2 Error analysis: true sensitivity limitations

Now, regarding the limitations of the true sensitivity, one can see make several observations. First, by looking at the defined trend between the cyan, black, and brown lines in Figs. 7.3 and 7.4, the magnitude of the peaks in the sensitivity of the transition point might be underpredicted due to the truncation error. Second, the discrepancies observed between the green and black lines at the lower surface leading edge of the pressure coefficient sensitivity may be attributed to the imperfect averaging of the benchmark, this is based on the fact that the cyan curve shows a similar profile, and the expected value of this profile has converged. The predictions using $\delta\nu_{t,FD}$ and $\delta\overline{u_i'u_j'}_{FD}$ show a similar behavior. Third, the analysis of the difference between the green and black lines of the pressure coefficient sensitivity on the

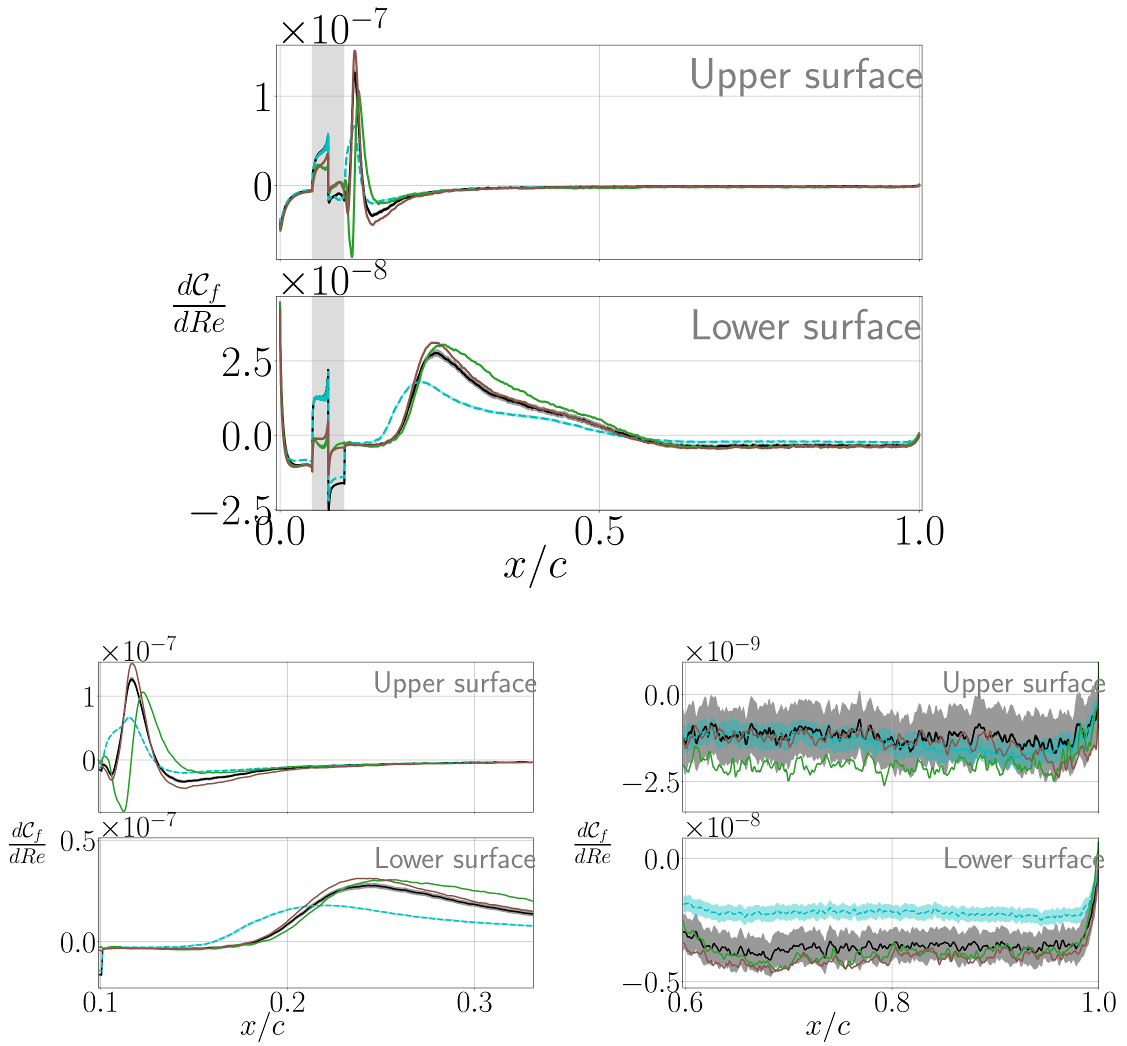

Figure 7.3: Sensitivity of the skin friction coefficient with respect to the Reynolds number. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The cyan dashed curve is the same sensitivity but with a larger $\Delta Re = 1 \times 10^5$. The shaded region around each of these curves represent the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The brown solid curve is the MFSA prediction using the LES-inferred $\delta\overline{u_i'u_j'}_{FD}$. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

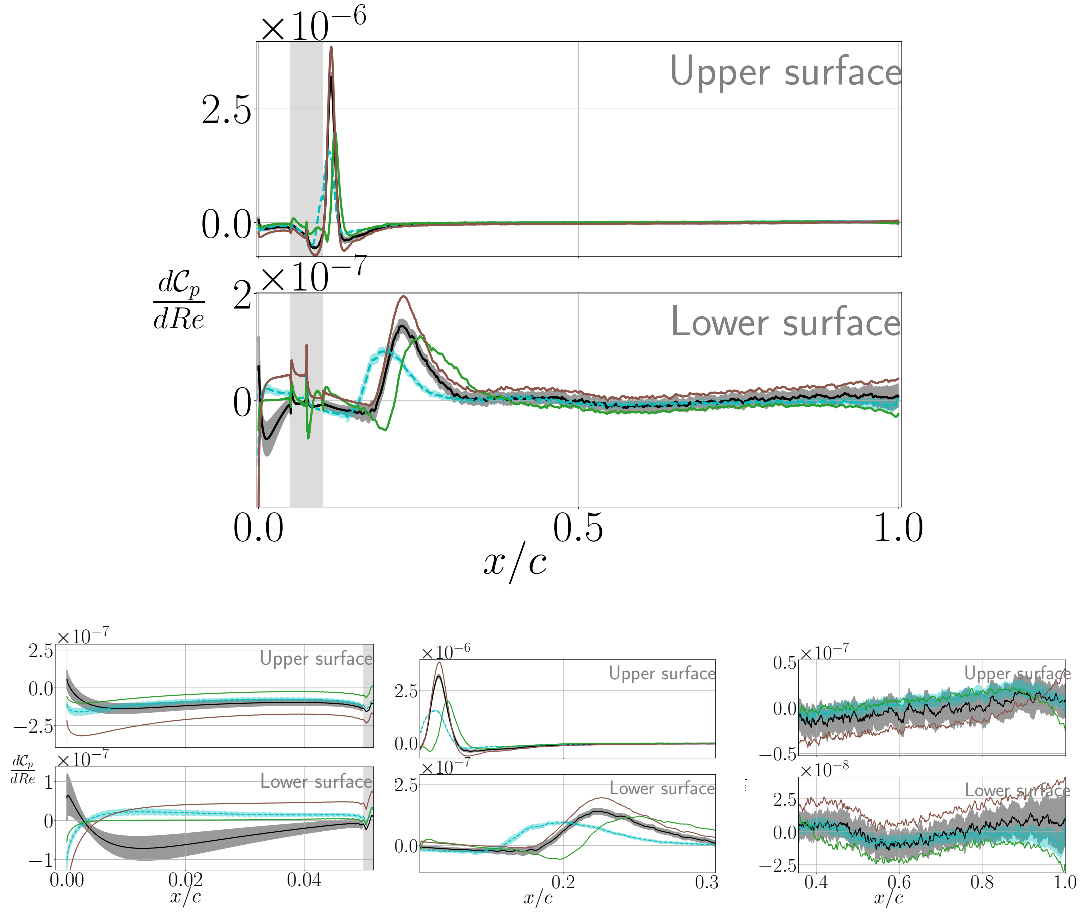

Figure 7.4: Sensitivity of the pressure coefficient with respect to the Reynolds number. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The cyan dashed curve is the same sensitivity but with a larger $\Delta Re = 1 \times 10^5$. The shaded region around each of these curves represent the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The brown solid curve is the MFSA prediction using the LES-inferred $\delta\overline{u_i'u_j'}_{FD}$. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

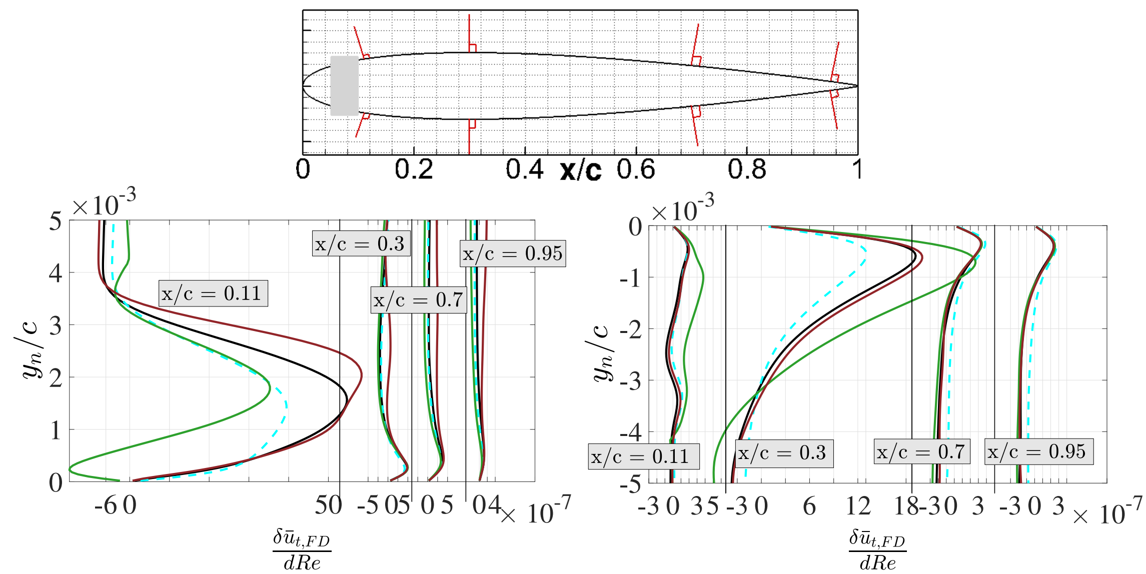

Figure 7.5: Profiles for the variation in the mean velocity with respect to the Reynolds number, $d\overline{u}_{t,FD}/dRe$ at different airfoil locations. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The cyan dashed curve is the same sensitivity but with a larger $\Delta Re = 1 \times 10^5$. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The brown solid curve is the MFSA prediction using the LES-inferred $\delta\overline{u_i'u_j'}_{FD}$.

lower surface trailing edge is inconclusive.

7.2 Importance of the perturbed eddy viscosity on the proposed MFSA method

The previous section showed that, provided an accurate closure model, the proposed multi-fidelity approach is able to predict the sensitivity of the flow. To create a meaningful method, however, we must define a model that does not require additional LES runs. The process of inferring the base eddy viscosity from the LES data at nominal conditions is difficult to improve upon, and the main challenge lies in the modeling of the perturbed eddy viscosity $\delta\nu_t$. For that reason, the following sections analyze the effects of the perturbed eddy viscosity on the MFSA predictions. The results of stand-alone RANS finite differencing will also be shown, since this represents the current state-of-the-art in engineering practice.

7.2.1 Frozen eddy viscosity assumption

Perhaps the easiest and one of the most common assumptions in standard sensitivity analysis is to assume that the eddy viscosity is constant when perturbing the flow (known as the "frozen eddy viscosity" or "frozen turbulence" assumption [44–47]). This would imply solving the system of equations (6.4) assuming $\delta\nu_t = 0$. Figure 7.6 shows the results obtained when this assumption is used (blue dotted curves). The frozen eddy viscosity assumption is insufficiently accurate. None of the important features on the transition region are captured, the predictions at the leading edge and towards the trailing edge on both surfaces for the pressure coefficient sensitivity have different magnitude and behavior when compared to the the true values. Figure 7.5 shows the profiles for the variation in the mean velocity with respect to the Reynolds number, $d\bar{u}_{t,FD}/dRe$ at different airfoil locations. One can notice that after the suction and blowing region, the velocity profiles are approximately constant throughtout the remainder of the surface.

These results show the significance of having an accurate perturbed eddy viscosity model on the sensitivity predictions. In addition, and more importantly, the results confirm that the perturbed eddy viscosity carries the information on how the turbulence changes when design parameters are changed.

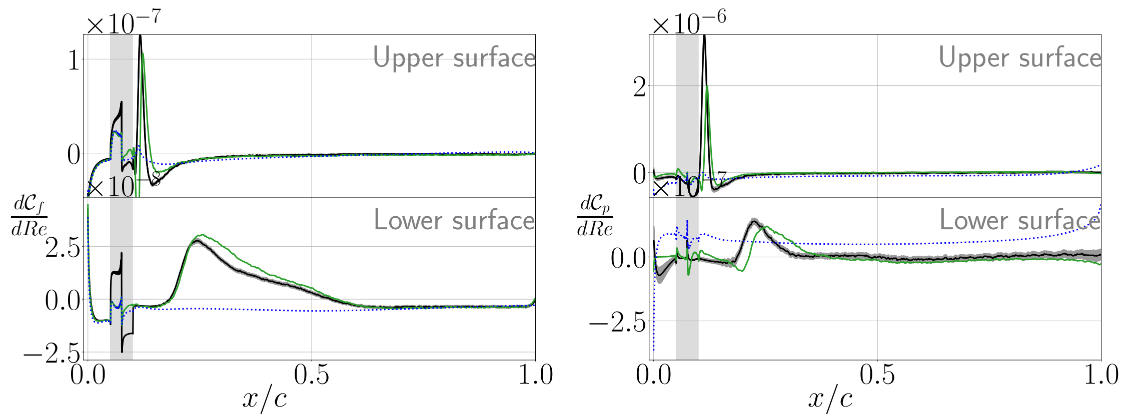

Figure 7.6: Sensitivities of the skin friction coefficient (left) and pressure coefficient (right) with respect to the Reynolds number. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The shaded region around this curve represents the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The blue dotted curve is the MFSA prediction using the frozen turbulence assumption, $\delta\nu_t = 0$. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

7.2.2 Modelled perturbed eddy viscosity

Having concluded that the perturbed eddy viscosity is crucial, the next natural step is trying to model it. In this study, as previously explained in chapter 2, the eddy viscosity perturbation is modeled using Prandtl's zero-equation model given

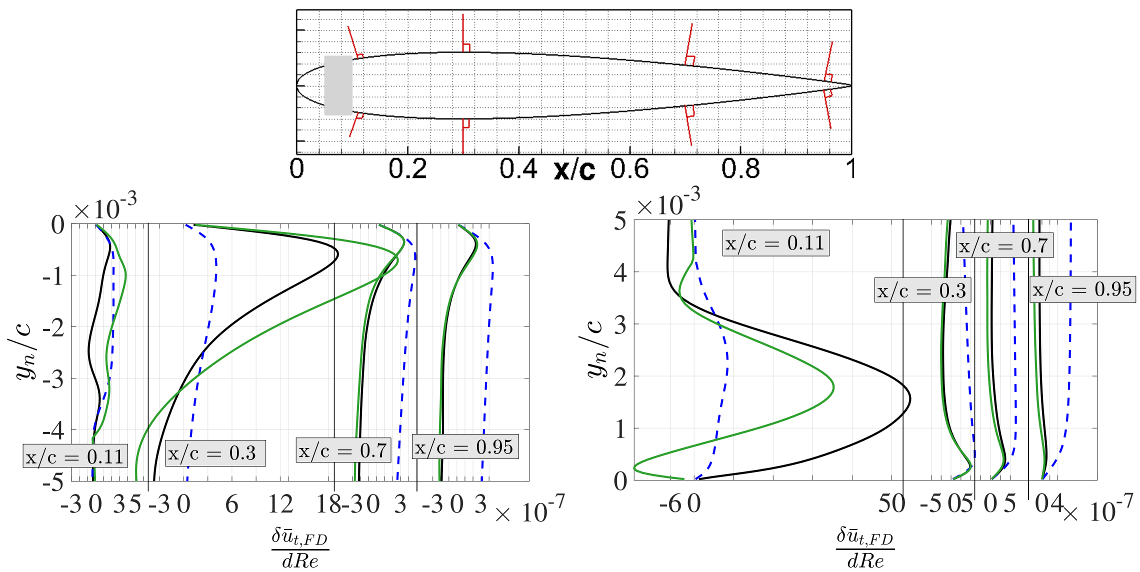

Figure 7.7: Profiles for the variation in the mean velocity with respect to the Reynolds number, $d\overline{u}_{t,FD}/dRe$ at different airfoil locations. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The blue dotted curve is the MFSA prediction using the frozen turbulence assumption, $\delta\nu_t = 0$.

by

$$\delta\nu_{t,Pr} = \nu_t \left[\frac{\overline{S}_{ij}\delta\overline{S}_{ij}}{\overline{S}_{kl}\overline{S}_{kl}} + \frac{d\,\mathrm{sgn}(\tau_w)\exp\left[\frac{-d^+}{A^+}\right]\delta\tau_w}{\nu A^+\sqrt{D|\tau_w|}} - \frac{d^+\exp\left[\frac{-d^+}{A^+}\right]\delta\nu}{\nu A^+\sqrt{D}} \right] . \tag{7.5}$$

After the verification process shown in chapter 6, the model is used to predict the perturbed solution $\delta\overline{\mathbf{v}}$. This model is chosen because it does not require the solution of an extra transport equation, meaning it does not add significant cost to the tangent equation solution, and its implementation is non-instrusive.

Equation 7.5 shows that $\delta\overline{\mathbf{v}}$ is a function of the base eddy viscosity, among other variables. In this study, two different base eddy viscosities were used: one using the inferred eddy viscosity from LES given by equation (7.2) and another using the standard Prandtl's formula given by

$$\nu_t = \kappa^2 d^2 \left| S \right| D \quad . \tag{7.6}$$

Figure 7.8 presents the sensitivities for the skin friction pressure coefficients with respect to the Reynolds number. The MFSA predictions using the perturbed eddy viscosity from Prandtl's zero-equation model, given by equation (7.5), are shown as pink dash-dotted curve and orange dashed curves. The pink dash-dotted curves use ν_t from LES given by equation (7.2) and the orange dashed curves use ν_t from the standard Prandtl's formula given by equation (7.6). One can notice in this figure that the profiles in orange are better than the pink ones. First, when using the inferred eddy viscosity, the sensitivity of the transition point on the lower surface is not captured in both QoIs. Second, the sensitivity of the transition point on the upper surface for the skin friction coefficient is three times higher than the true value. Third, the sensitivity of the pressure coefficient on the lower surface towards the trailing edge is clearly off. Fourth, when comparing the L2 norm for the difference of the skin friction and pressure coefficients sensitivities in the following section, one can see that quantitatively the predictions using the standard Prandtl's formula are better than the pink ones. Thus, using the standard Prandtl's formula is the approach chosen for the proposed MFSA method and used to compare against the fully-RANS based method.

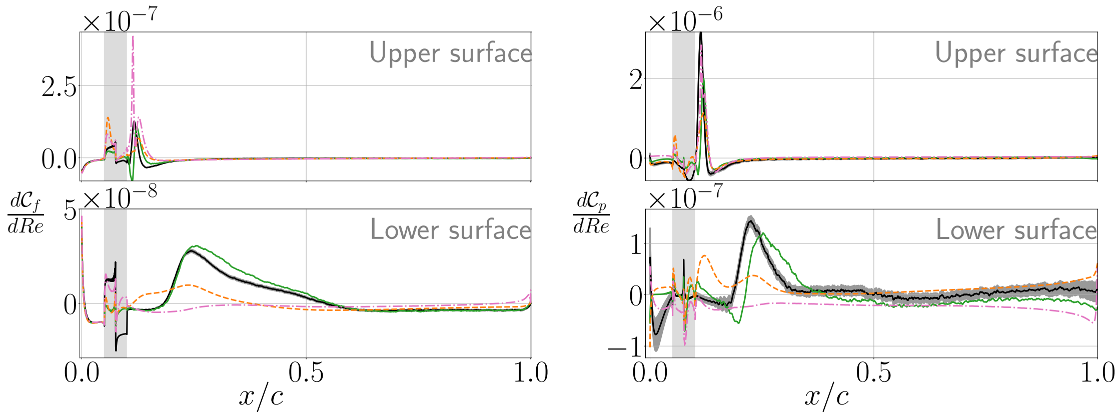

Figure 7.8: Sensitivities for the skin friction and pressure coefficient with respect to the Reynolds number. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The shaded region around this curve represents the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The pink dash-dotted curve is the MFSA prediction using Prandtl's zero-equation $\delta\nu_{t,Pr}$ with ν_t from LES given by equation (7.2). The orange dashed curve is the MFSA prediction using Prandtl's zero-equation $\delta\nu_{t,Pr}$ with ν_t from the standard Prandtl's formula given by equation (7.6). The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

Figures 7.9 and 7.10 present the sensitivities for the skin friction pressure coefficients with respect to the Reynolds number. The MFSA predictions using the perturbed eddy viscosity from Prandtl's zero-equation model, given by equation (7.5), are shown as orange dashed curves. Additionally, the red dash-dotted lines are the finite difference sensitivity obtained using the RANS model with a $\Delta Re = 50$; these results would represent the current state-of-the-art in engineering practice. Figure 7.11 shows the profiles for the variation in the mean velocity and inferred eddy viscosity with respect to the Reynolds number, $d\overline{u}_{t,FD}/dRe$ and

$d\nu_{t,FD}/dRe$. The main observation here is that the results of the proposed MFSA with the perturbed Prandtl model are closer to the true sensitivities in most parts of the airfoil and at least about the same on the lower surface leading edge when compared against the results from the fully RANS-based method.

Comparing results from the RANS-based finite-difference method, which has a computational cost of solving two RANS simulations (one standard forward simulation and either another forward RANS, a linear RANS, or an adjoint RANS simulation), against the MFSA method, which has a cost of solving one LES and one RANS, should not be understood as a direct comparison. As explained in chapter 1, one of the purposes of the MFSA method is to complement, rather than replace, the fully RANS-based method in situations in which sensitivities with higher fidelity are required within the design process.

With this clarification, one can make several observations from the Figs. 7.9 and 7.10. First, the MFSA method significantly improves the sensitivity predictions when compared with the fully RANS-based method on the upper surface of the airfoil for both quantities of interest. In particular, the proposed method is able to accurately predict both sensitivities of the transition point on the upper surface, which would be almost neglected by the fully RANS-based method. On the lower surface of the skin friction and pressure coefficient sensitivities the MFSA predictions are qualitatively the same as the predictions of the fully RANS-based method, except between $0.1 < x/c < 0.15$, where the MFSA prediction is qualitatively worse. Although using $\delta\nu_{t,Pr}$ improves the predictions towards the trailing edge for both sensitivities when compared against the frozen turbulence assumption, they show a similar behavior to the ones obtained by the fully RANS-based method, but are somewhat different to the behavior shown by the true sensitivities. When looking at Fig. 7.11 one can see that both the MFSA and the RANS-based methods overpredict the $d\overline{u}_{t,FD}/dRe$ by 2 and 2.4 times, respectively, on both surfaces. This could explain the different behavior of the skin friction and pressure coefficient sensitivities towards the trailing edge.

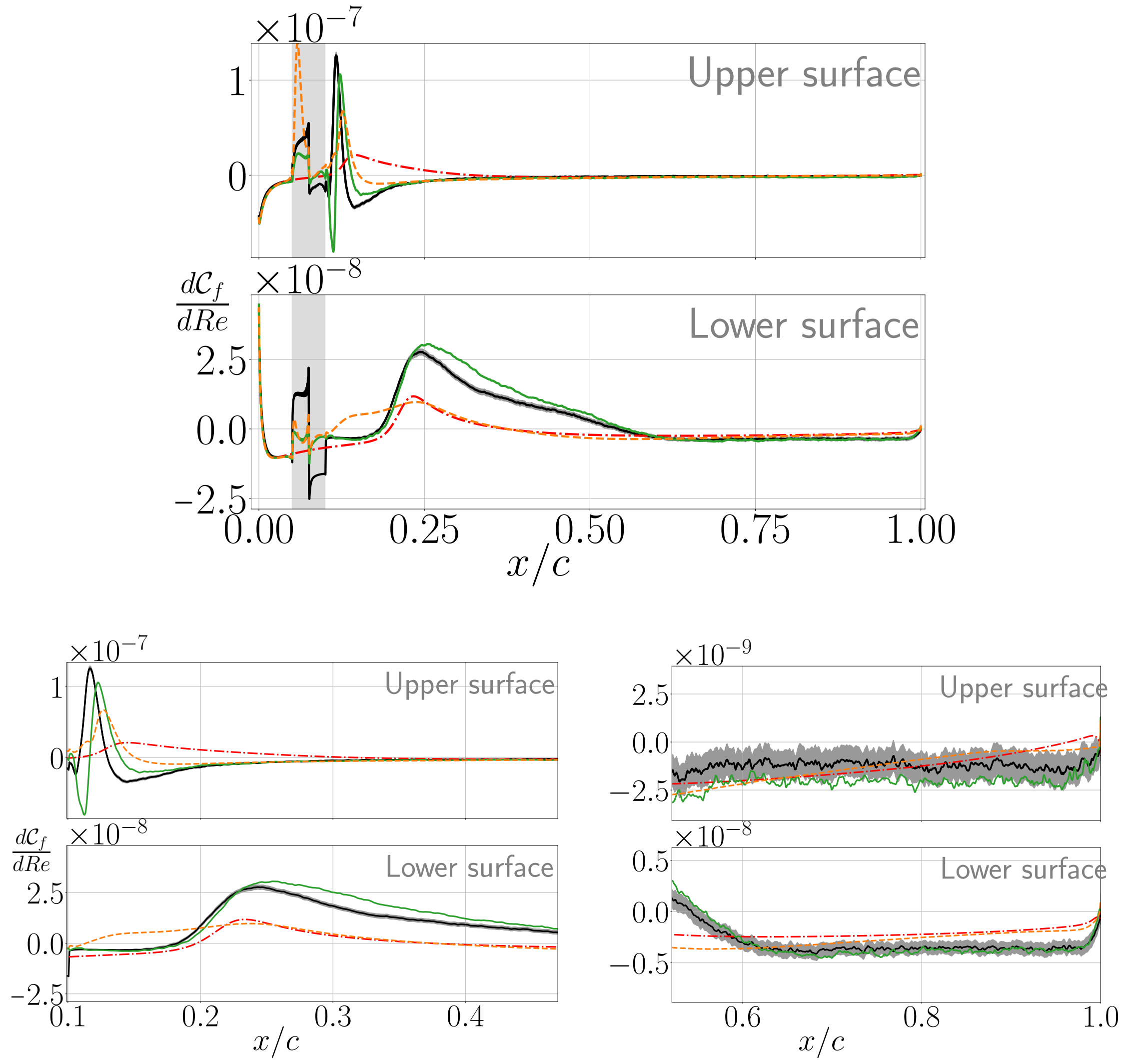

Figure 7.9: Sensitivities for the skin friction coefficient with respect to the Reynolds number. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The shaded region around this curve represents the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The red dash-dotted curve is the finite difference sensitivity obtained using RANS with a $\Delta Re = 50$. The orange dashed curve is the MFSA prediction using Prandtl's zero-equation $\delta\nu_{t,Pr}$. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

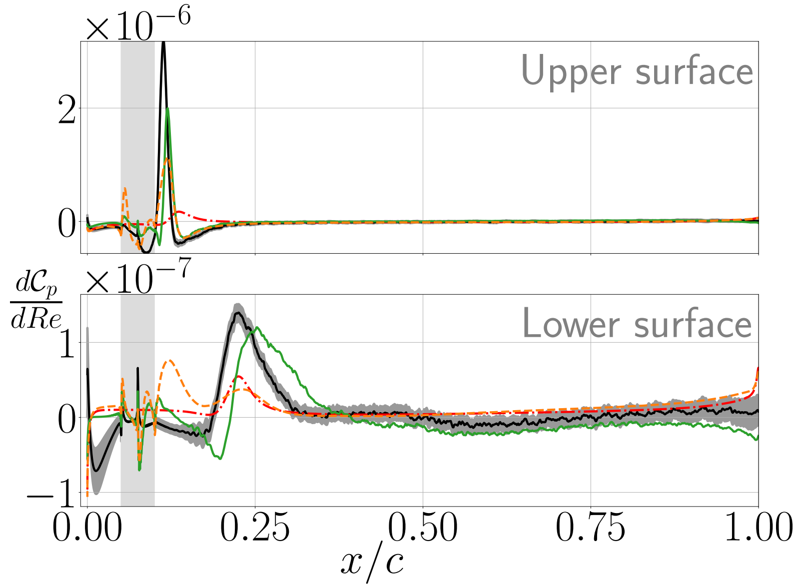

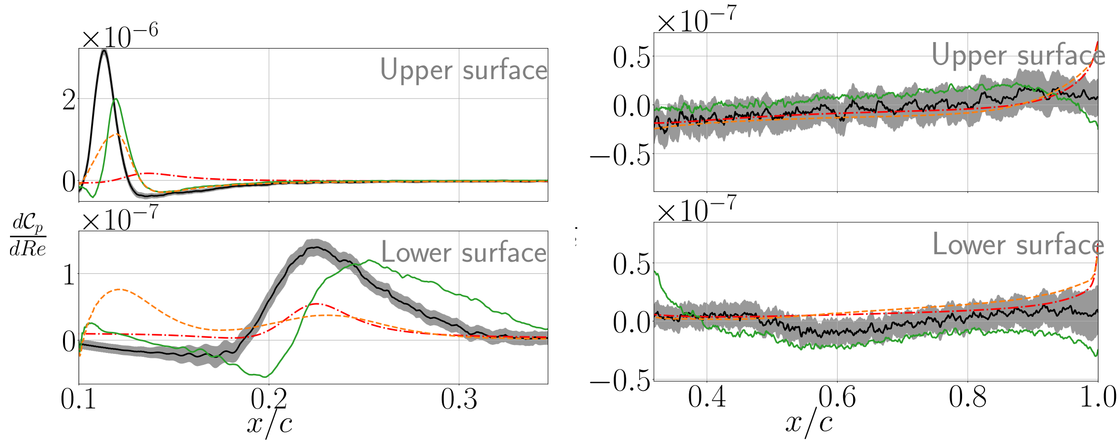

Figure 7.10: Sensitivities for the pressure coefficient with respect to the Reynolds number. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The shaded region around this curve represents the 95 percent confidence interval due to the imperfect averaging. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The red dash-dotted curve is the finite difference sensitivity obtained using RANS with a $\Delta Re = 50$. The orange dashed curve is the MFSA prediction using Prandtl's zero-equation $\delta\nu_{t,Pr}$. The gray rectangular area between $0.05 < x/c < 0.1$ is the region in which the suction and blowing is applied.

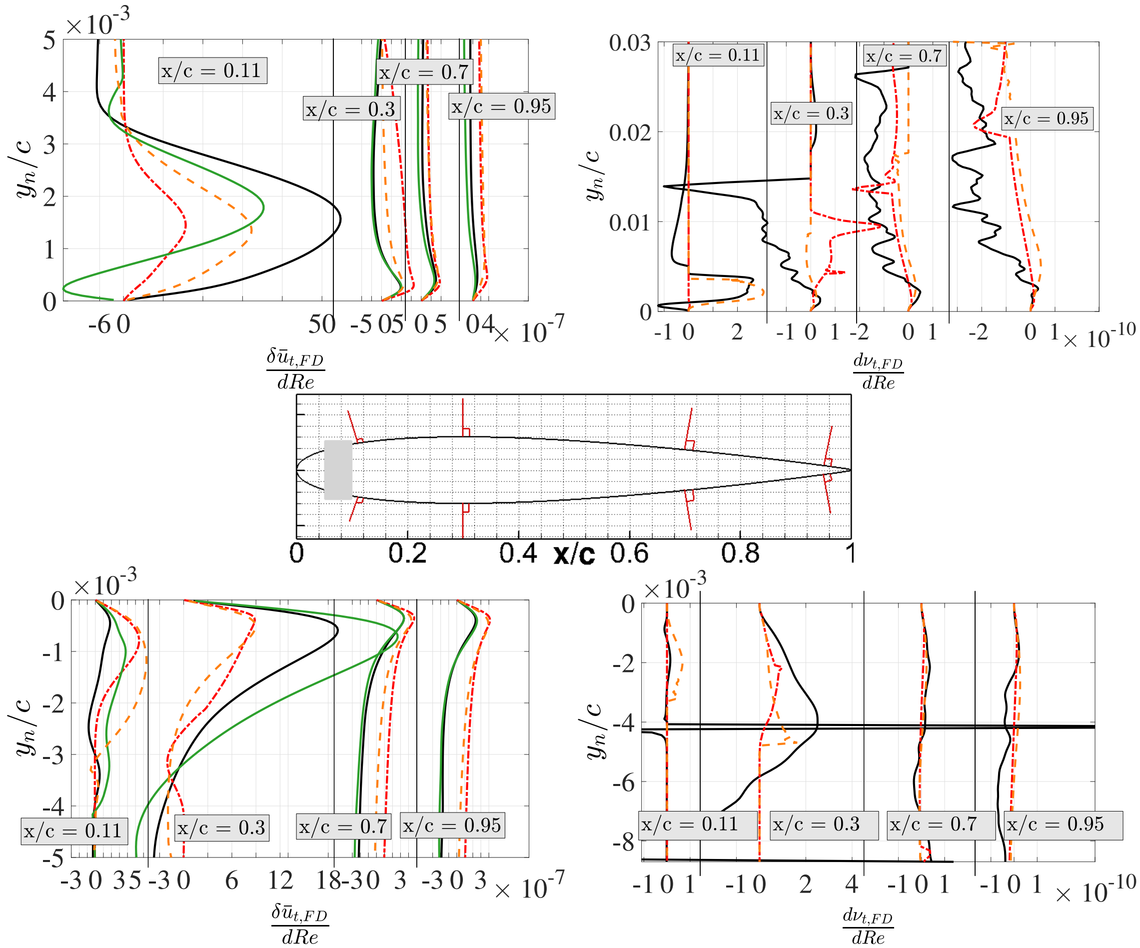

Figure 7.11: Profiles for the variation in the mean velocity and inferred eddy viscosity with respect to the Reynolds number, $d\bar{u}_{t,FD}/dRe$ and $d\nu_{t,FD}/dRe$, at different airfoil locations. The black solid curve is the "true" sensitivity computed using LES with a finite-difference over $\Delta Re = 4 \times 10^4$. The green solid curve is the MFSA prediction using the LES-inferred $\delta\nu_{t,FD}$. The red dash-dotted curve is the finite difference sensitivity obtained using RANS with a $\Delta Re = 50$. The orange dashed curve is the MFSA prediction using Prandtl's zero-equation $\delta\nu_{t,Pr}$.

7.3 Aerodynamic coefficient sensitivities

Table 7.1 shows the sensitivity predictions of the aerodynamic coefficients for all methods analyzed. Table 7.2 shows a comparison of the L2 norm for the difference of the skin friction and pressure coefficients sensitivities, $||d\mathcal{C}_f/dRe_{true} - d\mathcal{C}_f/dRe||$ and $||d\mathcal{C}_p/dRe_{true} - d\mathcal{C}_p/dRe||$, respectively. Both tables are organized showing the method with the closest predictions to the true sensitivities first, followed by the furthest predictions from the true sensitivity predictions. Several observations can be drawn. First and most importantly, assessing the accuracy of proposed method by only comparing the aerodynamic coefficients could be misleading. One can see for example that the MFSA method with Prandtl's model for closure modeling, $\delta\nu_{t,Pr}$, outperforms all the other methods, and that after the results obtained using the frozen turbulence assumption, predictions using the MFSA with $\delta\overline{u_i'u_j'}_{FD}$ are the worst. Second, using the frozen turbulence assumption is clearly wrong, no matter what metric is used, the predictions using this assumption are the worst when compared to the other alternatives. Third, it seems that the MFSA arguably performs better at predicting the quantities of interest that are viscous in nature like drag or skin friction coefficients than the ones that are inviscid in nature like lift or pressure coefficients. This might be because the Reynolds number is the perturbed parameter. Fourth, Table 7.2 presents a better assessment of the predictions obtained by the different closure models. It shows that the MFSA with $\delta\overline{u_i'u_j'}_{FD}$ is the most accurate at predicting the sensitivities and that the MFSA results using the frozen turbulence assumption is the worst.

Table 7.1: Comparison of sensitivity predictions of the aerodynamic coefficients for all methods analyzed.

Method	$\frac{d\mathcal{C}_D}{dRe} \times 10^9$	$\frac{d\mathcal{C}_L}{dRe} \times 10^8$
FD LES $\Delta Re = 4 \times 10^4$	-2.86±0.27	1.62±0.38
MFSA with $\delta\nu_{t,Pr}$	-3.01 (+5%)	2.65 (+64%)
FD RANS $\Delta Re = 50$	-4.99 (+75%)	1.68 (+4%)
FD LES $\Delta Re = 1 \times 10^5$	-4.39±0.15 (+53%)	0.80±0.19 (-51%)
MFSA with $\delta\nu_{t,FD}$	-2.14 (+25%)	-0.479 (-129%)
MFSA with $\delta\overline{u_i'u_j'}_{FD}$	-5.11 (+78%)	6.07 (+270%)
MFSA with $\delta\nu_{t,FD} = 0$	-13.9 (+385%)	15.0 (+826%)

Table 7.2: L2 norm for the difference of the skin friction and pressure coefficients sensitivities, $||d\mathcal{C}_f/dRe_{true} - d\mathcal{C}_f/dRe||$ and $||d\mathcal{C}_p/dRe_{true} - d\mathcal{C}_p/dRe||$ for all methods analyzed.

Method	$\vert\vert d\mathcal{C}_f/dRe_{true} - d\mathcal{C}_f/dRe\vert\vert$	$\vert\vert d\mathcal{C}_p/dRe_{true} - d\mathcal{C}_p/dRe\vert\vert$
MFSA with $\delta\overline{u_i'u_j'}_{FD}$	2.2×10^{-5}	4.4×10^{-4}
MFSA with $\delta\nu_{t,FD}$	6.1×10^{-5}	10.8×10^{-4}
MFSA with $\delta\nu_{t,Pr}$	6.9×10^{-5}	8.9×10^{-4}
FD RANS $\Delta Re = 50$	7.2×10^{-5}	12.8×10^{-4}
MFSA with $\delta\nu_t = 0$	7.3×10^{-5}	13.8×10^{-4}
MFSA with $\delta\nu_{t,Pr}$ ν_t from LES	1.0×10^{-4}	7.6×10^{-4}

Chapter 8: Conclusions and suggestions for future direction

This dissertation proposes and performs the feasibility analysis of a new methodology to estimate the sensitivity of chaotic and turbulence-resolving simulations. The proposed methodology represents a multi-fidelity method for sensitivity analysis, where the low-fidelity model (linearized-RANS in this particular study) is leveraged to reduce the computational cost of computing the sensitiviy due to small perturbations, but recoursing to the high-fidelity model to improve the accuracy of the prediction. Traditionally, the outcome of an LES or DNS is the prediction of the QoIs, for instance, the aerodynamic lift coefficient. The outcome of the proposed multi-fidelity sensitivity analysis (MFSA) is, not only the prediction of the aerodynamic coefficient, but also the sensitivity of the QoI at the extra cost of only one RANS simulation. The proposed method aims to complement, rather than replace, the fully RANS-based method, in situations in which sensitivities with higher fidelity are required.

Mathematically, the MFSA method should be able to predict the exact sensitivities provided the exact infinitesimal perturbed mean Reynolds stress, $\delta\overline{u_i'u_j'}$. However, there is no access to this term and therefore it needs to be modeled. In this study, $\delta\overline{u_i'u_j'}$ is modeled using a standard constitutive model, in which the eddy viscosity, ν_t, is approximated by minimizing the error between the deviatoric part of the Reynolds stress tensor given by high- and low-fidelity models; and the perturbed eddy viscosity, $\delta\nu_t$, is modeled using a simple zero-equation model.

A turbulent flow over a NACA 0012 profile at an angle of attack of $5°$ with a

Reynolds number of 4×10^5 is chosen to assess the proposed method. This profile is chosen because of the numerous studies and data published for different angles of attack and Reynolds numbers, and because it is relevant for industrial applications and engineering design. This angle of attack is chosen to ensure that both aerodynamic coefficients are larger than zero in a situation in which the flow around the airfoil remains attached. A moderate Reynolds number is chosen in order to have computationally affordable wall-resolved large eddy simulations.

A proof-of-concept of the MFSA is done in which the “best available estimate” for the perturbed eddy viscosity is given as an input to the method and its results are compared against a benchmark. A solution verification technique based on statistical analysis and hypothesis testing was developed and implemented to assess the accuracy of the MFSA’s predictions. In particular, the solution verification establishes the sensitivity obtained by computing the finite-difference of two different large eddy simulations as the true value and used to quantify the accuracy of the MFSA’s predictions. The proof-of-concept indicates that the proposed method is able to capture the different hydrodynamic phenomena present along the two airfoil surfaces, provided accurate modeling for the perturbed eddy viscosity, $\delta\nu_t$. However, there were differences between the true benchmark and the MFSA’s predictions.

An error analysis was carried out to identify the sources of error. It was shown that imperfect modeling introduces errors, like a short delay in the transition point sensitivity and a sharp decrement in the sensitivity of the skin friction before transitioning on the upper surface. Likewise, the “best available estimate” for the perturbed eddy viscosity introduces error because it is, after all, a finite-difference approximation. On the other hand, some of the discrepancies were due to the limitations on the benchmark predictions. For example, the truncation error affects strongly the magnitude and location of the transition point sensitivity; and the imperfect averaging affects significantly the predictions at the leading edge.

It is common practice in literature to assume that the eddy viscosity is constant when perturbing the flow. This study found that this assumption is insufficiently accurate. None of the important features on the transition region are captured. The predictions at the leading edge and towards the trailing edge on both surfaces for the pressure coefficient sensitivity have different magnitude and behavior when compared to the the true values. But most importantly, the analysis confirmed that the perturbed eddy viscosity carries the information on how the turbulence changes when design parameters are changed.

Having concluded that the perturbed eddy viscosity is crucial, the next step is modeling it. The eddy viscosity perturbation is modeled using Prandtl's zero-equation model and compared against the fully RANS-based method. This comparison should not be understood as a direct comparison; results from the fully RANS-based method have a computational cost of two RANS simulations, whereas the MFSA method has a cost of one LES and one RANS. One of the purposes of the MFSA method is to complement, rather than replace, the current state-of-the-art method in situations in which sensitivities with higher fidelity are required within the design process.

The MFSA method signficantly improves the sensitivity predictions when compared with the fully RANS-based method on the upper surface of the airfoil for both quantities of interest. In particular, the proposed method is able to accurately predict both sensitivities of the transition point on the upper surface, which would be almost neglected by the state-of-the-art method.

8.1 Publications and presentations

8.1.1 Publications

- Performed the first ever feasibility study and error analysis of the novel multi-fidelity sensitivity analysis (MFSA) proposed in this work. A paper describing this work will be submitted to the *AIAA Journal* at the end of the spring 2022.

- Developed and implemented a way to frame the question of grid sufficiency in turbulence-resolving simulations as a hypothesis test. This work is currently under review for the *International Journal of Computational Fluid Dynamics*.

8.1.2 Presentations

- Two presentations to the Predictive Science Academic Alliance Program III (PSAAP III) committee review: first at the annual tri-lab support team (TST) meeting in December 2020, and second at the alliance strategy team (AST) meeting in September 2021.

- Two presentations at the APS Division of Fluid Dynamics conference in 2019 and 2021.

8.2 Future direction

Although it is clear that the MFSA method signficantly improves the sensitivity predictions when compared with the fully RANS-based (the current state-of-the-art in practice) method, its full potential would be clearly seen when used in situations in which the different models (in this case LES and RANS) predict different base conditions. The method could be used on situations that are close to the edge of the operational envelope, like a turbulent flow past an airfoil near aerodynamic stall. It is known that the LES and RANS models predict different stall angles. Under drastically different base flows, the proposed MFSA is expected to outperforms the state-of-the-art method.

Based on the finding that the perturbed eddy viscosity or the perturbed Reynolds stress are crucial to the MFSA's predictions, it may be worthwhile exploring the possibility of developing eddy viscosity and/or Reynolds stress models that explicitly target the perturbed turbulence quantities. It is clear that all eddy viscosity and Reynolds stress models are tuned to predict base quantities and therefore might not be suited to predict changes in turbulence.

It seems that the MFSA arguably performs better at predicting the quantities of interest that are viscous in nature like drag or skin friction coefficients than the ones that are inviscid in nature like lift or pressure coefficients. This might be because the Reynolds number is the perturbed parameter. It may be worthwhile exploring the behavior of the MFSA method when an inviscid design parameter is perturbed, for example the angle of attack for the same case studied in this thesis.

Bibliography

[1] Juan C. del Álamo and Javier Jiménez. Spectra of the very large anisotropic scales in turbulent channels. *Physics of Fluids*, 15(6):L41–L44, 2003.

[2] A. Tanarro, R. Vinuesa, and P. Schlatter. Effect of adverse pressure gradients on turbulent wing boundary layers. *Journal of Fluid Mechanics*, 883:A8, 2020.

[3] David C. Wilcox. *Turbulence Modeling for CFD*. DCW industries, 2006.

[4] Christopher Rumsey. Turbulence modeling resources. Website: https://turbmodels.larc.nasa.gov/.

[5] Antony Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3(3):233–260, Sep 1988.

[6] Edward Tinoco. *The changing role of computational fluid dynamics in aircraft development*, pages 161 – 174. AIAA, 1998.

[7] Forrester T. Johnson, Edward N. Tinoco, and N. Jong Yu. Thirty years of development and application of cfd at boeing commercial airplanes, seattle. *Computers & Fluids*, 34(10):1115 – 1151, 2005.

[8] Johan Larsson and Qiqi Wang. The prospect of using large eddy and detached eddy simulations in engineering design, and the research required to get there. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2022):20130329, 2014.

[9] Benjamin. Peherstorfer, Karen. Willcox, and Max. Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591, 2018.

[10] Walters R.W. and Huyse L. Uncertainty analysis for fluid mechanics with applications. Technical report, ICASE-NASA, 2002.

[11] Krzysztof J. Fidkowski and David L. Darmofal. Review of output-based error estimation and mesh adaptation in computational fluid dynamics. *AIAA Journal*, 49(4):673–694, 2011.

[12] Patrick J. Blonigan, Qiqi Wang, Eric J. Nielsen, and Boris Diskin. Least-squares shadowing sensitivity analysis of chaotic flow around a two-dimensional airfoil. *AIAA Journal*, 56(2):658–672, 2018.

[13] Qiqi Wang. Convergence of the least squares shadowing method for computing derivative of ergodic averages. *SIAM Journal on Numerical Analysis*, 52(1):156–170, 2014.

[14] Patrick J. Blonigan, Qiqi Wang, Eric J. Nielsen, and Boris Diskin. *Least Squares Shadowing Sensitivity Analysis of Chaotic Flow around a Two-Dimensional Airfoil*, chapter 1, pages 658–672. AIAA, 2018.

[15] Angxiu Ni and Qiqi Wang. Sensitivity analysis on chaotic dynamical systems by non-intrusive least squares shadowing (nilss). *Journal of Computational Physics*, 347:56 – 77, 2017.

[16] Angxiu Ni. Approximating ruelle's linear response formula by shadowing methods, 2020.

[17] Nisha Chandramoorthy and Qiqi Wang. A computable realization of ruelle's formula for linear response of statistics in chaotic systems, 2020.

[18] Mathelin L. and Hussaini Y. A stochastic collocation algorithm for uncertainty analysis. Technical report, ICASE-NASA, 2003.

[19] Serhat Hosder, Robert Walters, and Rafael Perez. *A Non-Intrusive Polynomial Chaos Method For Uncertainty Propagation in CFD Simulations*, chapter 1, pages 1–19. 44th AIAA Aerospace Sciences Meeting and Exhibit, 2006.

[20] Habib N. Najm. Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 41(1):35–52, 2009.

[21] Dongbin Xiu and George Em Karniadakis. Modeling uncertainty in flow simulations via generalized polynomial chaos. *Journal of Computational Physics*, 187(1):137 – 167, 2003.

[22] Steven A. Orszag and L. R. Bissonnette. Dynamical properties of truncated wiener-hermite expansions. *The Physics of Fluids*, 10(12):2603–2613, 1967.

[23] Alexandre Joel Chorin. Gaussian fields and random flow. *Journal of Fluid Mechanics*, 63(1):21–32, 1974.

[24] Leo Wai-Tsun Ng and Michael Eldred. *Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation*, chapter 1, pages 1–17. 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, 2012.

[25] Michael S. Eldred, Gianluca Geraci, Alex Gorodetsky, and John Jakeman. Multilevel-multidelity approaches for forward uq in the darpa sequoia project. In *2018 AIAA Non-Deterministic Approaches Conference*, 2018.

[26] Christopher Critzos, Harry Heyson, and Robert Boswinkle. Aerodynamic characteristics of naca 0012 airfoil section at angles of attack from 0 to 180 degrees. Technical Report 3361, National Advisory committee for aeronautics, 1955.

[27] N. Gregory and C. L. OReilly. Low-speed aerodynamic characteristics of naca 0012 aerofoil section, including the effects of upper-surface roughness simulating hoar frost. Technical Report 3726, Aerodynamic division, N.P.L., 1970.

[28] Paul Klimas Robert Sheldahl. Aerodynamic characteristics of seven symmetrical airofil section through 180-degree angle of attack for use in aerodynamic analysis of vertical axis wind turbines. Technical Report SAND80-2114, Sandia National Laboratories, 1981.

[29] A. Michos, G. Bergeles, and N. Athanassiadis. Aerodynamic characteristics of naca 0012 airfoil in relation to wind generators. *Wind Engineering*, 7(4):247–262, 1983.

[30] W.J. McCroskey. A critical assessment of wind tunnel results for the naca 0012 airfoil. Technical Report 100019, National Aeronautics and Space Administration, 1987.

[31] Nikhil Oberoi. Eddy-viscosity inference. *TBD*, TBD(TBD):TBD, TBD.

[32] A. W. Vreman. An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. *Physics of Fluids*, 16(10):3670–3681, 2004.

[33] Yan Zang, Robert L. Street, and Jeffrey R. Koseff. A non-staggered grid, fractional step method for time-dependent incompressible navier-stokes equations

in curvilinear coordinates. *Journal of Computational Physics*, 114(1):18 – 33, 1994.

[34] Alan A. Wray. Minimal storage time advancement schemes for spectral methods, 1990.

[35] LLNL. Hypre: Scalable linear solvers and multigrid methods, 2020. Available at: https://computing.llnl.gov/projects/hypre-scalable-linear-solvers-multigrid-methods.

[36] William R. Wolf, Joao Luiz F. Azevedo, and Sanjiva K. Lele. Convective effects and the role of quadrupole sources for aerofoil aeroacoustics. *Journal of Fluid Mechanics*, 708:502–538, 2012.

[37] National Academy of Sciences. Assessing the reliability of complex models: Mathematical and statistical foundations of verification, validation, and uncertainty quantification. Technical report, NRC, 2012.

[38] C. Mockett, T. Knacke, and F. Thiele. Detection of initial transient and estimation of statistical error in time-resolved turbulent flow data. In *Conference: 8th International Symposium on Engineering Turbulence Modelling and Measurements*, ETMM8 N147, 2010.

[39] M. Bergmann, C. Morsbach, G. Ashcroft, and E. Kgeler. Statistical error estimation methods for engineering-relevant quantities from scale-resolving simulations. *J. Turbo.*, 21-1038, 2021.

[40] D. M. Israel. A new approach for turbulent simulations in complex geometries. PhD thesis, University of Arizona, 2005.

[41] T. A. Oliver, N. Malaya, R. Ulerich, and R. D. Moser. Estimating uncertainties in statistics computed from direct numerical simulation. *pof*, 26:035101, 2014.

[42] S. Russo and P. Luchini. A fast algorithm for the estimation of statistical error in DNS (or experimental) time averages. *jcp*, 347:328–340, 2017.

[43] P. Fischer, J. LOttes, and S. Kerkemeier. Nek5000: Open source spectral element cfd solver, 2008. Available at: https://nek5000.mcs.anl.gov.

[44] W.Kyle Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *Computers & Fluids*, 28(4):443–480, 1999.

[45] D.I. Papadimitriou and K.C. Giannakoglou. A continuous adjoint method with objective function derivatives based on boundary integrals, for inviscid and viscous flows. *Computers & Fluids*, 36(2):325 – 341, 2007.

[46] C. Othmer. A continuous adjoint formulation for the computation of topological and surface sensitivities of ducted flows. *International Journal for Numerical Methods in Fluids*, 58(8):861–877, 2008.

[47] A. Jameson, L. Martinelli, and N.A. Pierce. Optimum aerodynamic design using the navier–stokes equations. *Theoretical and Computational Fluid Dynamics*, 10(1):213–237, Jan 1998.

[48] Georgi Kalitzin, Gorazd Medic, Gianluca Iaccarino, and Paul Durbin. Near-wall behavior of rans turbulence models and implications for wall functions. *Journal of Computational Physics*, 204(1):265–291, 2005.