UCLA
UCLA Electronic Theses and Dissertations

Title

Resolvent Analysis Techniques for Turbulent Flows: Application to an SUV Wake

Permalink

https://escholarship.org/uc/item/73v2d7n7

Author

House, Dylan

Publication Date

2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Resolvent Analysis Techniques for Turbulent Flows: Application to an SUV Wake

A thesis submitted in partial satisfaction

of the requirements for the degree

Master of Science in Aerospace Engineering

by

Dylan Christopher House

2021

ABSTRACT OF THE THESIS

Resolvent Analysis Techniques for Turbulent Flows: Application to an SUV Wake

by

Dylan Christopher House

Master of Science in Aerospace Engineering

University of California, Los Angeles, 2021

Professor Kunihiko Taira, Chair

Resolvent analysis allows for the extraction of the dominant input-output behavior of a fluid flow near a mean state, which can be used to advise potential applications to flow control. A significant hurdle in the adoption of resolvent analysis is the singular value decomposition (SVD) of the large linear operators involved. A matrix sketching algorithm is used to extract the primary forcing and response modes, with their associated gain. The formulation of an iterative algorithm is shown to be able to calculate the SVD of the resolvent operator with greater accuracy. The sources of error due to the selection of a test vector are discussed and it is shown that an accurate calculation of the forcing and response modes can be obtained by utilizing a test vector corresponding to a single point. The strength of this algorithm is shown by calculating the resolvent modes for a flow over a NACA 0012 airfoil at a Reynolds number of 23,000. This method is shown to converge for an arbitrary selection of test vector, obtaining results in agreement with past studies of this flow.

This method is used to perform windowed resolvent analysis on a two-dimensional turbulent flow of over a Honda sports utility vehicle (SUV) at $Re \approx 2.5 \times 10^4$ to propose a flow control strategy for drag reduction. The force characteristics of the

SUV are highly dependent on the wake structures and large structures in the wake region near the rear end of the vehicle are responsible for increase in drag of an SUV. We characterize the flow physics and modal analysis of the shear layer on the roof and the role of shear layer physics leading to the large structures in the wake dynamics. A moving window approach to investigate the windowed response modes of the resolvent operator over the roof of SUV is performed. The location of maximum gain shifts to the rear of the car as the spanwise wavenumer is reduced indicating the transition from small structures in the shear layer to large structures in the wake region. We use a body constrained window for forcing modes to investigate the optimum location of forcing to achieve desired response along the roof of the vehicle. An optimal forcing strategy about the placement of actuators and the frequency and direction of forcing is suggested after investigating the forcing modes at different frequencies and wave numbers.

The thesis of Dylan Christopher House is approved.

Jeff Eldredge

Ann Karagozian

Kunihiko Taira, Committee Chair

University of California, Los Angeles

2021

To my family, without whom I would not have made it to where I am today.

TABLE OF CONTENTS

LIST OF FIGURES

ix

ACKNOWLEDGMENTS

CHAPTER 1

Background

Modal analysis has become a widespread tool in fluid mechanics (BHL93; HLB12; Sch10; Kut13; The11; TBD17; THB19). Among these modal decomposition techniques, resolvent analysis reveals the input-output relationship with respect to a mean flow of interest (FI93; TTR93; MS10). Resolvent analysis is based on the concept of a particular solution under forcing, as opposed to global stability analysis which focuses on the homogeneous solution (Sch07). For this reason, resolvent analysis provides insights on how the flow responds to harmonic input to the system as well as the transient dynamics (JB05; TE05). As flows are exposed to perturbations and forcing at all times, resolvent analysis becomes a very attractive tool for characterizing, modeling, and controlling high-dimensional flows. Noteworthy here is that resolvent analysis can be performed about time-averaged turbulent base flows if the fluctuations about the base states are statistically stationary (FI94; MS10). For these reasons, resolvent analysis stands out as essentially the only operator-based modal analysis technique that can handle practical engineering flows at high Reynolds numbers.

As a resolvent operator is spatially discretized, it can be represented by a matrix whose dimension is proportional to the number of spatial discretization points (grid size) and the number of state variables. The size of this matrix becomes very large for high-Reynolds number turbulent flows and especially for those without homoge-

neous directions (e.g., spatial periodicity). In such a case, computational challenges are encountered as the singular value decomposition of large-scale matrices demands significant computational resources and memory allocations.

Although the task of performing the resolvent analysis may seem daunting for large-scale flow problems, it should be noted that most fluid flow studies require only the dominant features of the flow to be revealed. That means that the amount of insights to be revealed is relatively small compared to the amount of data held by the large resolvent matrix. This allows for the use of low-rank approximation of the large matrices to be considered to reduce the computational efforts necessitated by the singular value decomposition.

In particular, we consider a technique referred to as *sketching* to extract the dominant characteristics of the original large-scale matrix (EVB19). Sketching involves passing a test matrix comprised of a few linearly independent column vectors through the operator to be analyzed. The product of the original matrix and the test matrix holds important information about the action of the original matrix on the test matrix. In other words, the product should capture the dominant vectors in the range of the original matrix. The use of a random test matrix for determining appropriate rank reduction has resulted in the establishment of the randomized numerical linear algebra, which has become an active area of research in recent years.

In the present work, we place our focus on the choice of test matrix or test vector for the resolvent analysis and how it can be chosen based on the base flow. Furthermore, we demonstrate that in certain cases the singular value decomposition may be eliminated from the resolvent analysis, which significantly reduces the computational burden.

2

We first examine the randomized resolvent algorithm and its effect on single test vectors. Through this analysis we are able to derive error expressions for the forcing and response modes and the gain. We then perform resolvent analysis using this algorithm for flow over a NACA 0012 airfoil using test random vectors, and vectors corresponding to a single point in the flow. We discuss how the error is affected by the choice of test vector as well as how the behavior described by the error expression derived earlier manifests in modes obtained through this method.

This method of performing resolvent analsysis is then used to develop a potential flow control strategy for a Honda SUV. The primary goal of such a strategy is to reduce the drag force on the vehicle, improving fuel efficiency and limiting the negative environmental impacts of gas or electric vehicles. While there have been efforts made to perform active and passive flow control over cars in the past, active flow control studies have been primarily limited to low Reynolds number flows. Resolvent analysis is used to extract the dominant forcing and response structures for the flow over a 2D slice of a Honda SUV. Past work has shown that large low pressure structures in the wake are the dominant contributor to drag, with such structures responsible for at least 91% of the total pressure drag for flow over an Ahmed body (AR84) (This section has been adapted from a section originally written by Vedasri Godavarthi for a co-authored paper "Open loop flow control strategy for drag reduction over an SUV using windowed resolvent analysis", which is currently in preparation).

In order to gain insights into the wake dynamics, we utilize windowed resolvent analysis to study a focused region of the shear layer over the vehicle. By moving this window from the front of the roof to the rear, we can gain valuable insights into how the response of the flow transitions from small scale three-dimensional structures in the shear layer, to large two-dimensional structures in the wake. Generating

3

three-dimensional structures corresponding to these response modes is proposed as a potential flow control method by delaying the formation of large wake structures, limiting their effect on the body of the car.

A forcing side response window is applied restricting the potential region of actuation to the body of the car to identify optimal actuator locations. We show that the flow is most sensitive to disturbance near the front of the roof, however small scale structures are not very persistent and cannot be always actuated from this point. Instead the introduction of small-scale structures with actuation must be performed near to the desired response, while the actuation of larger structures can be performed from the front of the roof.

1.1 General Resolvent Analysis

For an unsteady fluid flow, we can decompose the state variable $\boldsymbol{q}$ into a time-invariant base flow $\bar{\boldsymbol{q}}$ and the time-variant fluctuations $\boldsymbol{q}'$ as

$$\boldsymbol{q} = \bar{\boldsymbol{q}} + \boldsymbol{q}'. \tag{1.1}$$

The time-variant and invariant state vectors are related by the following equation

$$\frac{\partial \boldsymbol{q}'}{\partial t} = \boldsymbol{L}_{\bar{\boldsymbol{q}}}\boldsymbol{q}' + \boldsymbol{f}, \tag{1.2}$$

where $\boldsymbol{L}_{\bar{\boldsymbol{q}}}$ is the linearized Navier–Stokes operator about the base state $\bar{\boldsymbol{q}}$, and $\boldsymbol{f}$ represents the collection of nonlinear terms, which can be interpreted as an internal forcing (MS10). We can analyze the frequency response of the flow through analyzing the Fourier transform of equation (1.2)

$$i\omega\hat{\boldsymbol{q}}_{\omega} = \boldsymbol{L}_{\bar{\boldsymbol{q}}}\hat{\boldsymbol{q}}_{\omega} + \hat{\boldsymbol{f}}_{\omega}. \tag{1.3}$$

where

$$\boldsymbol{q}(x,t) = \hat{\boldsymbol{q}}_\omega e^{i\omega t} \tag{1.4}$$

$$\boldsymbol{f}(x,t) = \hat{\boldsymbol{f}}_\omega e^{i\omega t} \tag{1.5}$$

The frequency ω can be chosen to be real for stable $\boldsymbol{L}_{\bar{q}}$ or complex as $\omega = \omega_r + i\beta$ for unstable $\boldsymbol{L}_{\bar{q}}$, where ω_r and β are both real and β discounts the modal growth rate of $\boldsymbol{L}_{\bar{q}}$ (Jov04; YBT20). By rearranging (1.3) we can obtain an operator that maps a forcing input to a response of the state vector $\boldsymbol{q}$ in the form of

$$\hat{\boldsymbol{q}}_\omega = \boldsymbol{A}\hat{\boldsymbol{f}}_\omega, \tag{1.6}$$

where

$$\boldsymbol{A} = [i\omega\boldsymbol{I} - \boldsymbol{L}_{\bar{q}}]^{-1}, \tag{1.7}$$

is the resolvent operator. We can perform a singular value decomposition (SVD) of $\boldsymbol{A}$ to yield

$$\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*. \tag{1.8}$$

Written in this form, $\boldsymbol{V}$ represents the primary directions in which forcings are most effective, referred to as forcing modes, $\boldsymbol{U}$ represents the responses these forcings will induce, referred to as response modes, and $\boldsymbol{\Sigma}$ represents the associated gains mapping between the forcing modes and the associated responses.

1.2 Randomized Resolvent Analysis

The size of the resolvent operator is proportional to the number of grid points used to discretize the flow field. For higher-Reynolds number flows, the grid needs to be well-refined making the resolvent operator large in its size, and hence makes it difficult to perform a SVD.

The randomized resolvent analysis can be used to reduce the computational cost and memory requirement associated with performing the SVD while producing accurate resolvent modes and gains (RYT20). The reduction in computational costs is accomplished by performing the SVD on a low-rank approximation $\boldsymbol{B}$ of the resolvent operator $\boldsymbol{A}$ instead of the full $\boldsymbol{A}$ itself. The first step towards obtaining the low-rank approximation is called sketching. In this most crucial part of the randomized resolvent analysis, a sketch matrix $\boldsymbol{Y}$ is generated by passing a tall and skinny test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{m \times k}$ ($m \gg k$) through the resolvent operator as

$$\boldsymbol{Y} = \boldsymbol{A}\boldsymbol{\Omega}. \tag{1.9}$$

Oftentimes, this test matrix $\boldsymbol{\Omega}$ is chosen to be a random matrix, $\boldsymbol{\Omega} = \texttt{rand}(m, k)$. In the present study, we take a detailed look at this sketching process with respect to the accuracies of the recovered resolvent modes and gains, and complement the randomized algorithm with physical insights.

The sketched matrix $\boldsymbol{Y}$ should contain the dominant effect of the operator $\boldsymbol{A}$. We perform the QR decomposition of $\boldsymbol{Y}$ to obtain a low-rank basis aligned with the primary actions (directions) of the resolvent operator $\boldsymbol{A}$

$$\boldsymbol{Y} = \boldsymbol{Q}\boldsymbol{R}. \tag{1.10}$$

We then project the resolvent operator $\boldsymbol{A}$ onto this low-rank basis, to obtain a low-

rank approximation of the resolvent

$$\boldsymbol{B} = \boldsymbol{Q}^*\boldsymbol{A} \in \mathbb{C}^{k \times m} \tag{1.11}$$

We can then perform the SVD of this projected matrix $\boldsymbol{B}$

$$\boldsymbol{B} = \widetilde{\boldsymbol{U}}\boldsymbol{\Sigma}\boldsymbol{V}^*. \tag{1.12}$$

While $\widetilde{\boldsymbol{U}}$ is the left singular vectors, we can obtain more accurate response modes of $\boldsymbol{A}$ by right multiplying the resolvent operator by the forcing modes $\boldsymbol{V}$ yielding

$$\boldsymbol{A}\boldsymbol{v} = \sigma\boldsymbol{u}, \tag{1.13}$$

Since the left singular vectors are unitary ($\|\boldsymbol{u}\| = 1$), we obtain

$$\sigma = \|\boldsymbol{A}\boldsymbol{v}\|, \tag{1.14}$$

$$\boldsymbol{u} = \sigma^{-1}\boldsymbol{A}\boldsymbol{v}. \tag{1.15}$$

This process is outlined in algorithm 1 and discussed in detail by (RYT20).

First the test matrix $\boldsymbol{\Omega}$ is passed through the resolvent in line 2 to form the sketch matrix $\boldsymbol{Y}$, whose columns represent responses to the k forcing vectors $\boldsymbol{\Omega}$. In this step, the forcings excite response modes associated with the largest gains. However, these excited response modes appear in every column of the sketch matrix, or equivalently the response vectors $\boldsymbol{Y}$. To extract the response modes, the response vectors $\boldsymbol{Y}$ are orthonormalized via a QR decomposition in line 3 to provide a set of orthogonal bases $\boldsymbol{Q}$ for the response modes. The resolvent operator $\boldsymbol{A}$ is then projected onto those bases $\boldsymbol{Q}$ in line 4 to produce a low-rank projected operator $\boldsymbol{B}$, which holds the principal characteristics of $\boldsymbol{A}$. The forcing modes $\boldsymbol{V}$ are then extracted in line 5 through the SVD of forcing bases $\boldsymbol{B}$. These forcing modes are left multiplied by $\boldsymbol{A}$ to yield $\sigma\boldsymbol{u}$ in line 6 (viz. equation 1.13) and subsequently separated into σ and $\boldsymbol{u}$ by

---

Algorithm 1: Randomized Resolvent Analysis (RYT20)

Require: Discrete resolvent operator $\boldsymbol{A} \in \mathbb{C}^{m \times m}$

    Function `randomized_resolvent(`k`):`

1      $\boldsymbol{\Omega} \leftarrow$ `randn(`m,k`)`   // generate random forcing vectors $\boldsymbol{\Omega}$

2      $\boldsymbol{Y} \leftarrow \boldsymbol{A\Omega}$ // sketching: obtain response vectors $\boldsymbol{Y}$ by passing $\boldsymbol{\Omega}$ to $\boldsymbol{A}$

3      $(\boldsymbol{Q}, \sim) \leftarrow$ `qr(`$\boldsymbol{Y}$`,0)`   // obtain orthonormal response bases $\boldsymbol{Q}$ from $\boldsymbol{Y}$

4      $\boldsymbol{B} \leftarrow \boldsymbol{Q}^{*}\boldsymbol{A}$ // obtain forcing bases $\boldsymbol{B}$ by passing $\boldsymbol{Q}$ to $\boldsymbol{A}^{*}$

5      $(\sim, \sim, \boldsymbol{V}) \leftarrow$ `svd(`$\boldsymbol{B}$`)`   // obtain forcing modes via reduced SVD of $\boldsymbol{B}$

6      $\boldsymbol{U}^{\boldsymbol{\Sigma}} \leftarrow \boldsymbol{AV}$ // obtain gain-scaled response modes $\boldsymbol{U}^{\boldsymbol{\Sigma}}$ by passing $\boldsymbol{V}$ to $\boldsymbol{A}$

    for $j \leftarrow 1$ *to* k **do**

7         $\boldsymbol{\Sigma}_{j,j} \leftarrow$ `norm(`$\boldsymbol{U}^{\boldsymbol{\Sigma}}_{1:m,j}, 2$`)`   // recover gain

8         $\boldsymbol{U}_{1:m,j} \leftarrow \boldsymbol{U}^{\boldsymbol{\Sigma}}_{1:m,j} / \boldsymbol{\Sigma}_{j,j}$ // obtain unit-norm response modes $\boldsymbol{U}$

    end

    return $(\boldsymbol{U}, \boldsymbol{\Sigma}, \boldsymbol{V})$

---

using the unitary nature of the singular vectors in lines 7 and 8 (viz. equation 1.15). This randomized resolvent analysis algorithm has been demonstrated to perform well even with $k \ll m$ to determine the forcing modes, response modes, and gain while retaining its numerical accuracy.

CHAPTER 2

Methodology

2.1 Iterative Sketching Algorithm

Let us take a detailed look at sketching within the randomized resolvent analysis, algorithm 1. In line 1, we form a random test matrix $\mathbf{\Omega} \in \mathbb{R}^{m \times k}$, which can be considered as being comprised of k random forcing vectors in the length of the state variable m. We take this perspective to derive an iterative formulation, in which we utilize the forcing modes obtained by the algorithm as test matrix for the next iteration. In what follows, we seek an algorithm that finds the primary forcing and response modes along with the largest singular value through an iterative formulation.

Sketching plays the most important role in the algorithm of the randomized resolvent analysis. The test matrix $\mathbf{\Omega}$ initiates the sketching process with the attempt of exciting modes with the largest gains. If this attempt is successful, the rest of the algorithm shall produce the optimal gains and the associated forcing and response modes, yielding an accurate SVD. The test matrix $\mathbf{\Omega}$ is the only tunable input we have control over in the randomized resolvent analysis. The choice of $\mathbf{\Omega}$ is crucial in producing accurate resolvent modes and gains, and we can rely on physical insights of the problem to improve the design of $\mathbf{\Omega}$.

A random matrix is intuitively the first choice for the test matrix $\boldsymbol{\Omega}$, if one has little knowledge about the physics of the operator (HMT11). Since a random $\boldsymbol{\Omega}$ can be interpreted as white forcing, it excites all response modes and allows those with highest gains to be amplified when being passed through the resolvent operator. However, better choices for $\boldsymbol{\Omega}$ can be made with the physical knowledge of the problem. (RYT20) considered a physics-informed test matrix based on the velocity gradient of the base flow and showed that the use of such test matrices can further enhance the accuracy of the randomized analysis with fewer number of vectors. Their demonstration showed great potential of leveraging flow physics in the design of test matrices.

We develop algorithm 2 by replacing the test matrix of algorithm 1 with a test vector, that is, a test matrix where $k = 1$. In this case, the sketch vector $\boldsymbol{y}$ contains the response bases for the forcing vector $\boldsymbol{\Omega}$. Since we only have a single vector to represent the response, there is no need for orthogonalization and we normalize the response $\boldsymbol{y}$ to obtain an approximation to the primary response mode. Similar to algorithm 1, we can determine the primary forcing basis by projecting the resolvent operator onto the response mode based on

$$\boldsymbol{A}^{*}\boldsymbol{u}_i = \sigma_i \boldsymbol{v}_i. \tag{2.1}$$

Due to the usage of a test vector, performing the SVD is not necessary and reduces to a simple normalization in algorithm 2. From here, we can once again utilize equation 1.13 to extract the response mode with greater accuracy. Since the response and forcing modes may not be accurate with a single evaluation, we then resort to an iterative procedure for lines 1 to 4. For the iterative process, the forcing mode from the previous iteration is used as the test vector for the following iteration, establishing algorithm 2.

10

Algorithm 2: Resolvent Analysis with a test vector $\boldsymbol{\Omega}$ (iterative)

Require: Discrete linear operator $\boldsymbol{L}_{\bar{q}} \in \mathbb{C}^{m \times m}$, test vector $\boldsymbol{\Omega} \in \mathbb{C}^{m \times 1}$

Function `test_vector_resolvent`(ω)**:**

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

1      $\boldsymbol{y} \leftarrow [-i\omega\boldsymbol{I} - \boldsymbol{L}_{\bar{q}}] \backslash \boldsymbol{\Omega}$ // solve linear system for $\boldsymbol{y}$

2      $\boldsymbol{q} \leftarrow \boldsymbol{y}/\|\boldsymbol{y}\|_2$ // normalize $\boldsymbol{y}$

3      $\boldsymbol{b} \leftarrow \boldsymbol{q}^*/[-i\omega\boldsymbol{I} - \boldsymbol{L}_{\bar{q}}]$ // solve linear system for $\boldsymbol{b}$

4      $\boldsymbol{\Omega} \leftarrow \boldsymbol{b}^*/\|\boldsymbol{q}\|_2$ // normalize $\boldsymbol{b}$ and transpose

end

5      $\boldsymbol{v} \leftarrow \boldsymbol{\Omega}$ // set $\boldsymbol{V}$ to $\boldsymbol{\Omega}$

6      $\boldsymbol{u}^{\boldsymbol{\Sigma}} \leftarrow [-i\omega\boldsymbol{I} - \boldsymbol{L}_{\bar{q}}] \backslash \boldsymbol{v}$ // solve linear system to recover $\boldsymbol{u}^{\boldsymbol{\Sigma}}$

7      $\sigma \leftarrow \|\boldsymbol{u}^{\boldsymbol{\Sigma}}\|_2$ // recover gain

8      $\boldsymbol{u} \leftarrow \boldsymbol{u}^{\boldsymbol{\Sigma}}/\sigma$ // recover response mode

return $(\boldsymbol{u}, \sigma, \boldsymbol{v})$

Algorithm 2 considers a user-selected test vector as the test matrix $\boldsymbol{\Omega}$. If we define the inner product $\alpha_i \equiv \langle \boldsymbol{v}_i, \boldsymbol{\Omega} \rangle$, we can show that

$$\boldsymbol{y} = \sum_{i=1}^{m} \sigma_i \alpha_i \boldsymbol{u}_i^{\text{exact}} \tag{2.2}$$

where $\boldsymbol{u}_i^{\text{exact}}$ is the exact ith response mode. From this, we can determine an expression for $\boldsymbol{b}^*$

$$\boldsymbol{b}^* = \frac{1}{\|\boldsymbol{y}\|} \sum_{i=1}^{m} \sigma_i^2 \alpha_i \boldsymbol{v}_i^{\text{exact}}. \tag{2.3}$$

The forcing mode can be found by simply normalizing $\boldsymbol{v}_i$

$$\boldsymbol{v}_1 = \boldsymbol{b}^*/\|\boldsymbol{b}^*\|. \tag{2.4}$$

We recover the response mode and its gain by multiplying $\boldsymbol{v}$ by the resolvent operator

11

to obtain

$$\sigma_1 \boldsymbol{u}_1 = \frac{1}{||\boldsymbol{b}||} \sum_{i=1}^{m} \sigma_i^3 \alpha_i \boldsymbol{u}_i^{\text{exact}}. \tag{2.5}$$

The main feature of this test vector based algorithm is that we make only a single attempt to initially excite the optimal responses. This prompts the question of how this test vector should be chosen. As previously discussed, (RYT20) suggested that the mean-flow physics should inform the choice. In this study, we will take this idea to its extreme by considering the case where the test vector $\boldsymbol{\Omega}$ is comprised of forcing at a single spatial point. We will later examine the effectiveness of this test vector with respect to the location of forcing. With the single-point test vector, we will use the iterative approach in algorithm 2 where the right singular vector obtained by the algorithm is then reused as the test vector for N iterations of this procedure.

We can apply equations 2.2-2.5 to the iterative form of the algorithm and substitute the values of $||\boldsymbol{y}||$ and $||\boldsymbol{b}||$ to obtain

$$\mathbf{v} = \frac{\sum_{i=1}^{m} \sigma_i^{2N} \alpha_i^* \boldsymbol{v}_i}{\sqrt{\sum_{i=1}^{m} \sigma_i^{4N} |\alpha_i|^2}}, \tag{2.6}$$

$$\sigma \boldsymbol{u} = \frac{\sum_{i=1}^{m} \sigma_i^{2N+1} \alpha_i^* \boldsymbol{u}_i}{\sqrt{\sum_{i=1}^{m} \sigma_i^{4N} |\alpha_i|^2}}, \tag{2.7}$$

$$\sigma^2 = \frac{\sum_{i=1}^{m} \sigma_i^{4N+2} |\alpha_i|^2}{\sum_{i=1}^{m} \sigma_i^{4N} |\alpha_i|^2}. \tag{2.8}$$

These expressions show the dependence of the output of the algorithm on all the singular values, and also on the initial test vector via α_i. Picking the test vector to perfectly align with the leading forcing mode would give $\alpha_i = \delta_{i1}$ and these expressions would then give the singular triplet $(\sigma_1, \boldsymbol{u}_1, \boldsymbol{v}_1)$ with no error. To examine the effect of a non-perfect choice we will make the assumption that the resolvent operator is

12

low-rank. This will mean that $\sigma_1 \gg \sigma_2 \gg ... \gg \sigma_N$. This assumption allows us to Taylor expand the above expressions to give an estimate of the error in the obtained modes.

$$\|\boldsymbol{v} - \boldsymbol{v}_1\|^2 \approx \sum_{i=2}^{m} \frac{\sigma_i^{4N}|\alpha_i|^2}{\sigma_1^{4N}|\alpha_1|^2}, \tag{2.9}$$

$$\|\boldsymbol{u} - \boldsymbol{u}_1\|^2 \approx \sum_{i=2}^{m} \frac{\sigma_i^{4N+2}|\alpha_i|^2}{\sigma_1^{4N+2}|\alpha_1|^2}, \tag{2.10}$$

$$\sigma \approx \sigma_1 - \frac{\sum_{i=2}^{m} \sigma_i^{4N}|\alpha_i|^2}{2\sigma_1^{4N-1}|\alpha_1|^2}. \tag{2.11}$$

While it is impractical to obtain an approximation of the error prior to performing the SVD, these expressions still provide valuable insight on the sources of error. From this we can see that at leading order the relative error for the singular values $|\sigma - \sigma_1|/\sigma_1$ is $\mathcal{O}(\sigma_2^{4N}/\sigma_1^{4N})$, for the forcing $\|\boldsymbol{v} - \boldsymbol{v}_1\|$ is $\mathcal{O}(\sigma_2^{2N}/\sigma_1^{2N})$ and for the response $\|\boldsymbol{u} - \boldsymbol{u}_1\|$ is $\mathcal{O}(\sigma_2^{2N+1}/\sigma_1^{2N+1})$.

This algorithm will produce much more accurate results for systems where $\sigma_1 \gg \sigma_2, \sigma_3, \sigma_4, \ldots$. This agrees with the observations in (RYT20). For cases where σ_1/σ_2 is sufficiently large, this algorithm will produce accurate results regardless of the test vector used, including random test vectors. While the magnitude of the singular values are purely system dependent, the test vector can be selected to minimize error. An ideal test vector is aligned perfectly with with the first order forcing mode. Realistically this is very difficult since the structure of the forcing modes cannot generally be known in advance. Instead we rephrase the problem as attempting to find a vector which is nearly orthogonal to all other modes except for the first, and due to the assumption that the resolvent operator is low rank, we are primarily interested in ensuring orthogonality to specifically the second order mode. By localizing the forcing in a region where the second order mode is weak we can minimize this source of error. As a result, we can obtain a more accurate approximation for flows where

13

the first two modes are spatially separated.

Forcing a single point of the flow means that α_1 will be smaller than what could be obtained by using a global forcing; however, if the point is chosen using some physical insight as to where the first order forcing mode should appear then α_1 should be significantly larger than α_2. Through careful selection of the test vector, accurate results can be obtained using a single point test vector.

2.2 Validation

2.2.1 Problem setup

We consider the application of the present formulation to uncover the response characteristics of turbulent separated flow over a NACA0012 airfoil (YT19). For the resolvent analysis, we consider the flow response with respect to the time-averaged turbulent base flow, which is obtained by a high-fidelity large-eddy simulation. The computation is performed for the flow over an airfoil at an angle of attack of $\alpha = 9°$, a chord-based Reynolds number of $Re = U_\infty c/\nu = 23,000$, and a free-stream Mach number of $M_\infty = 0.3$. The flow is taken to be spanwise-periodic with an extent of $0.2c$.

A finite-volume compressible flow solver CharLES (KNH11; BHN17) is used to perform the large-eddy simulation of the turbulent flow over the wing. This solver is second and third-order accurate in space and time, respectively, with the Vreman's model (Vre04) implemented for the subgrid-scale model. The spatial domain is chosen to be $x/c \in [-19, 26]$, $y/c \in [-20, 20]$ and $z/c \in [-0.1, 0.1]$ in the streamwise, transverse, and spanwise directions, respectively, with the leading edge positioned
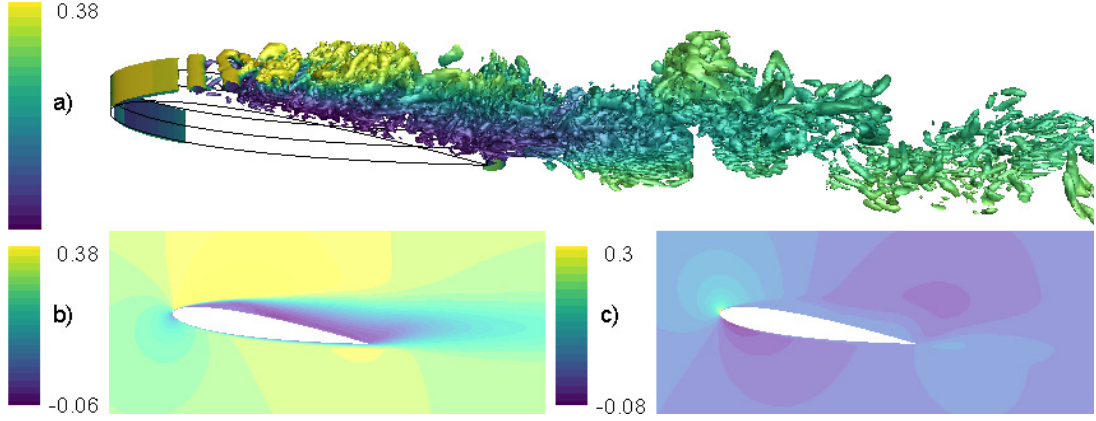
Figure 2.1: Flow over a NACA 0012 airfoil at a Reynolds number of 23,000 and and a Mach number of 0.3 visualized by the Q-criterion isosurfaces (a). Time-averaged flows for the u (b) and v (c) components of velocity.

at the origin. This computational domain is discretized with a C grid. We specify the freestream velocity and temperature at the far field and prescribe a no-slip adiabatic boundary condition on the wing surface. A sponge layer is applied over the outlet to enable outgoing disturbances to exit the domain without reflecting back towards the near wake of the wing.

The base flow visualized in figure 2.1 is obtained by taking the time and spanwise average of the turbulent flow. As the whole computational domain is not required for the resolvent analysis, we extract the base flow over $x/c \in [-15, 16]$ and $y/c \in [-12, 12]$. The discretization of the resolvent operator over this sub-domain results in a resolvent operator of size 0.75 million by 0.75 million, which will be used in the subsequent analysis. The computational setup and the simulated base flow have been verified and validated extensively, as reported in our past studies (YT19; NYK19; RYT20).

2.2.2   Resolvent analysis

We first examine results from the resolvent sketching algorithm for two choices of sparse test vectors, and compare the results obtained to results obtained using two different random test vectors. The first sparse test vector contains non-zero values only for the entries corresponding to the point of with the highest ratio α_1/α_2, and zeros otherwise. All non-zero elements in the vector are equal such that $||\boldsymbol{\alpha}|| = 1$. The second sparse test vector is similar in construction, with non-zero elements corresponding to point in the freestream. The points effectively represent the best and worst case scenarios for the selection of test vector. The random test vectors used are both generated using a random normal distribution, however one is additionally weighted according to the velocity gradient as described in (RYT20).

Results from resolvent analysis at $St = 1, \beta = 0$ and $St = 6, \beta = 20\pi$ are summarized in figures 2.2 and 2.3. Both of these are cases in which the algorithm can potentially struggle as σ_1/σ_2 is not very large, ≈ 6 and 20 respectively. For the former case, all response modes converge by ten iterations, with results showing excellent agreement. A poorly chosen test point results in the forcing mode converging more slowly, not converging by the tenth iteration. However selecting a test point at the location of the maximum velocity gradient results in the solution converging, and once again obtains excellent agreement with the exact solution. Of all the approaches, a single test point near the leading edge of the airfoil performs best after a single iteration, while a single test point in the wake performs the worst, as it is unable to excite the first order forcing mode at all.

For the $St = 6, \beta = 20\pi$ case, the results follow a similar pattern. After the first iteration, the random test vector, gradient weighted random test vector, and
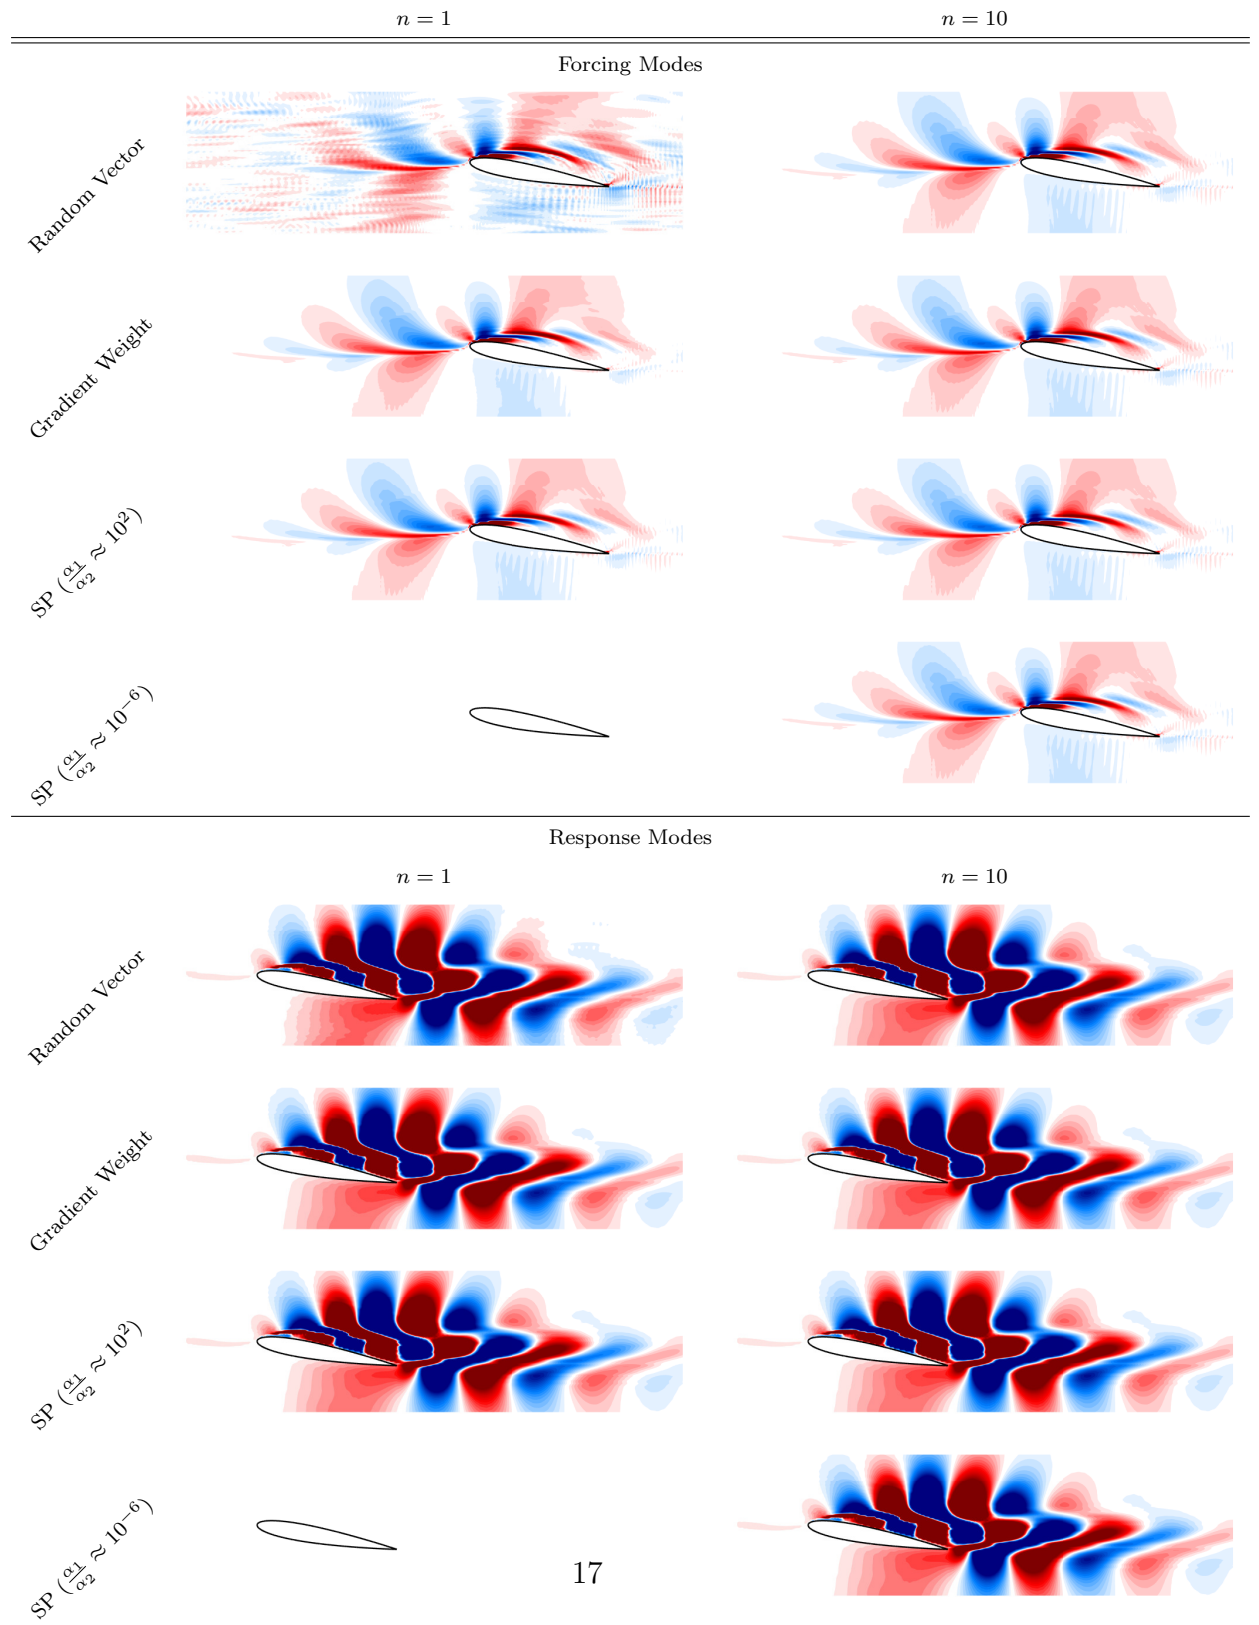
Figure 2.2: Forcing and Response modes at $St = 1, \beta = 0$ for different types of test vectors after 1 and 10 iterations
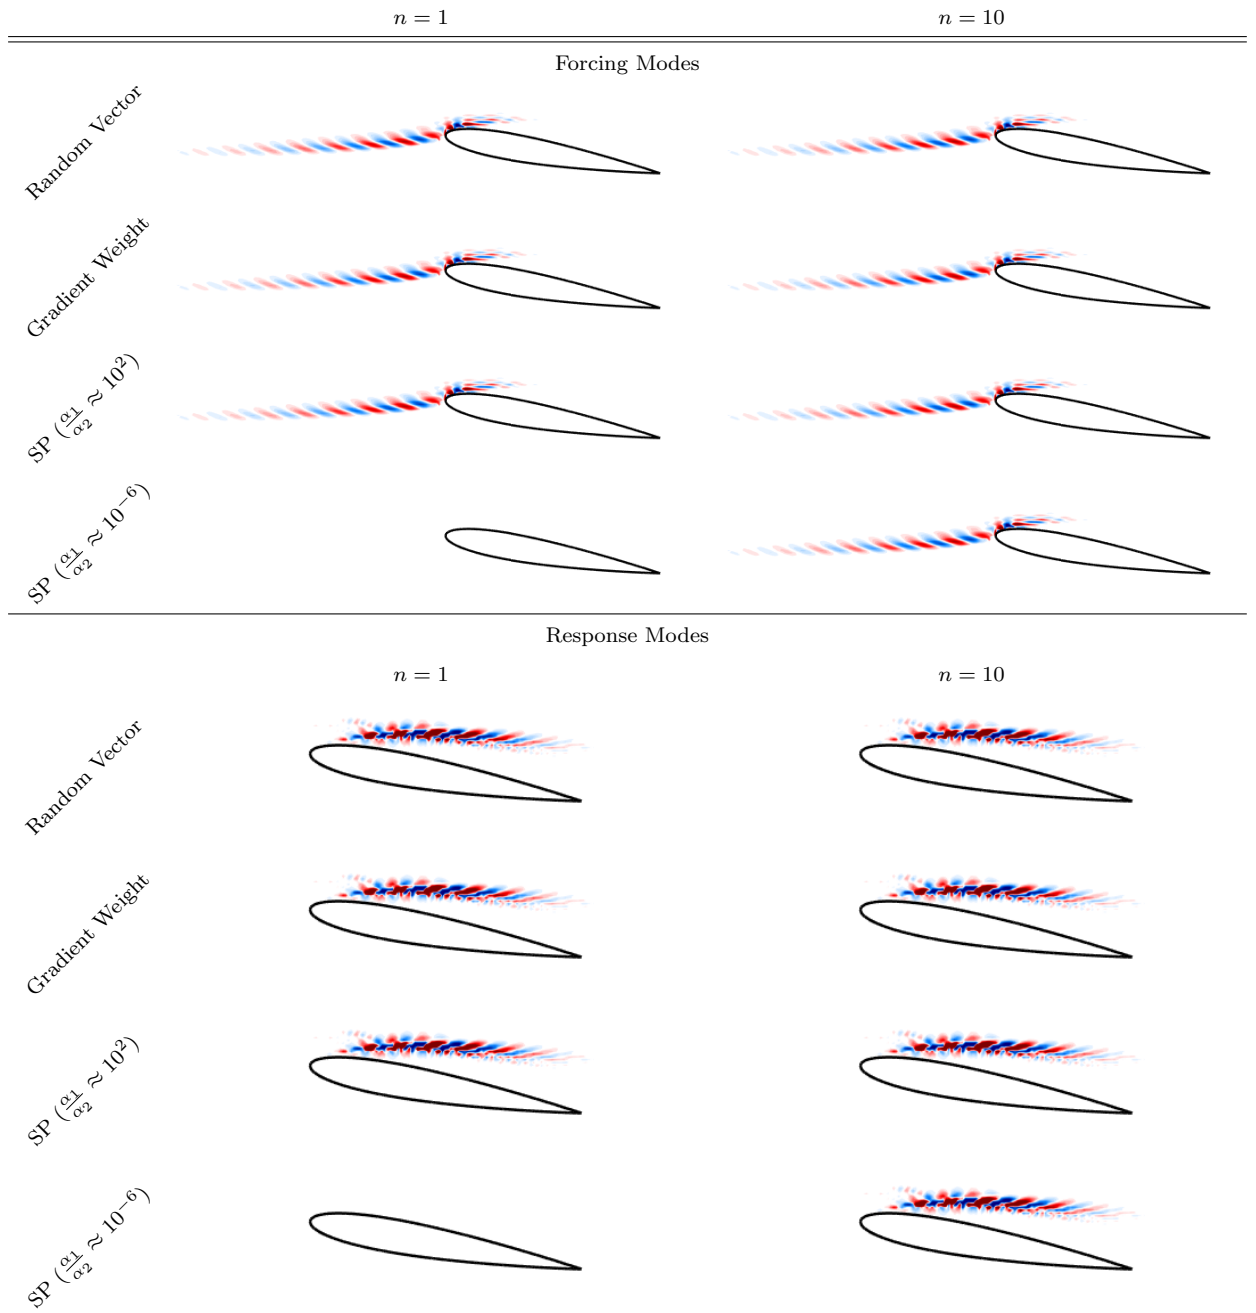
Figure 2.3: Forcing and Response modes at $St = 6, \beta = 20\pi$ for different types of test vectors after 1 and 10 iterations

the single point near the leading edge are already very difficult to distinguish, in agreement with the results of (RYT20) The single point forcing near the wake is unable to excite the forcing mode, however by the tenth iteration also results in an accurate solution. While the random test vector has been shown to work consistently, a single point test vector can obtain even better performance when chosen correctly as it can more precisely target purely the first order mode.

Convergence behavior of the iterative algorithm is shown in figures 2.4 and 2.5. The convergence rate for all methods follows the theoretical predictions, depending only on the ratio σ_1/σ_2. The convergence rate for some high error test vectors demonstrates odd behavior for the initial iterations, since the error is so large in these cases that the simplifying assumption used in the Taylor expansion of the error formula is no longer valid. For the $St = 6, \beta = 20\pi$ case, the wake test vector actually converges slower than the $St = 1, \beta = 0$ case despite a higher σ_1/σ_2. Early on the error for this case is dominated by the spatial discrepancy between the test vector and the forcing mode resulting in a value of $\alpha_1/\alpha_2 \approx 10^{-6}$. In order for the solution to converge the test vector in future iterations must overlap with the structure of the forcing mode, which is more difficult in the $\beta = 20\pi$ case as the size of the forcing mode is much smaller.

While the convergence rate is indeed similar for all methods, the choice of sketch vector can have a dramatic impact on the initial error, with a single sketch point placed at the leading edge of the airfoil has an error of order $< 10^{-5}$ while a randomized vector has an error of 10^2 for $St = 1, \sigma_1/\sigma_2 \approx 6$. For cases with ratios $\sigma_1/\sigma_2 \approx 1$, poorly chosen test vectors can struggle to converge, however proper selection of the test vector can result in convergence after only 1 or 2 iterations.

We additionally examine the error and convergence behavior of the algorithm for
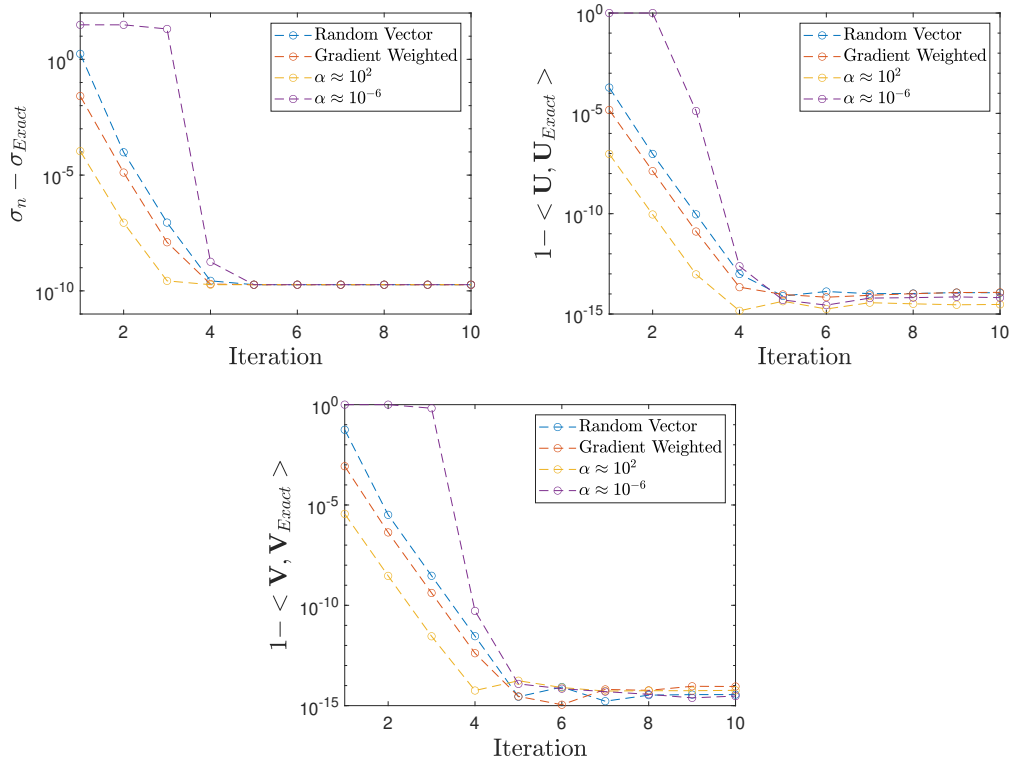
Figure 2.4: Singular value error; Cosine similarity error of response modes; Cosine similarity error of forcing modes for $St = 1, \beta = 0$.
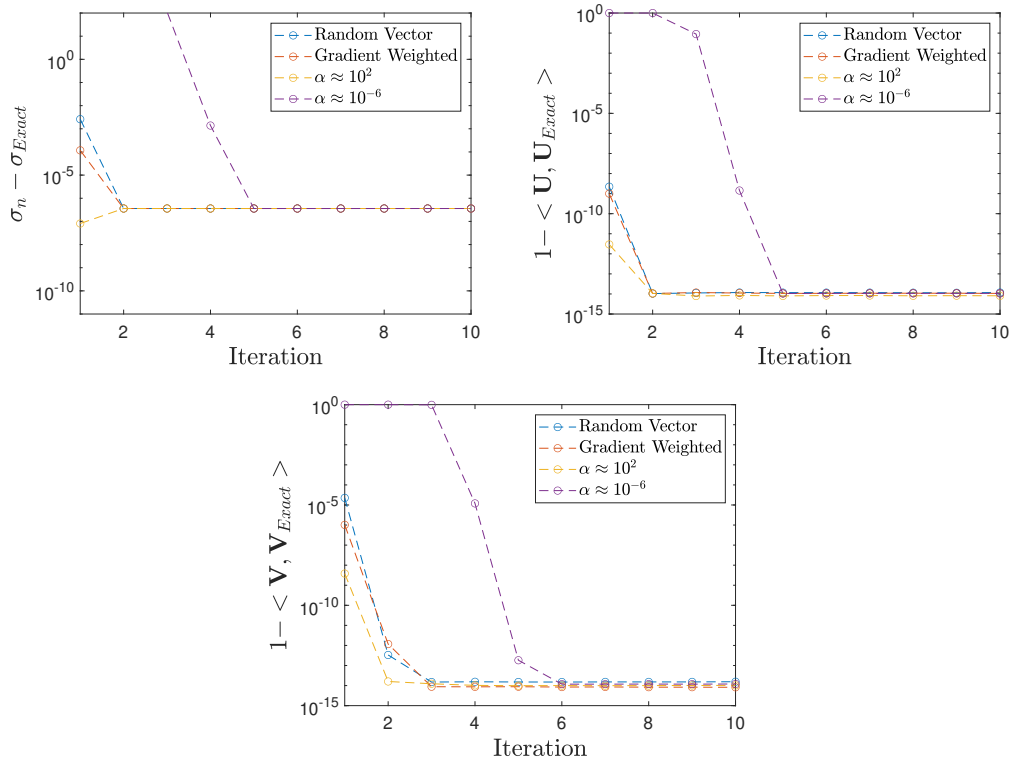
Figure 2.5: Singular value error; Cosine similarity error of response modes; Cosine similarity error of forcing modes for $St = 6, \beta = 20\pi$.
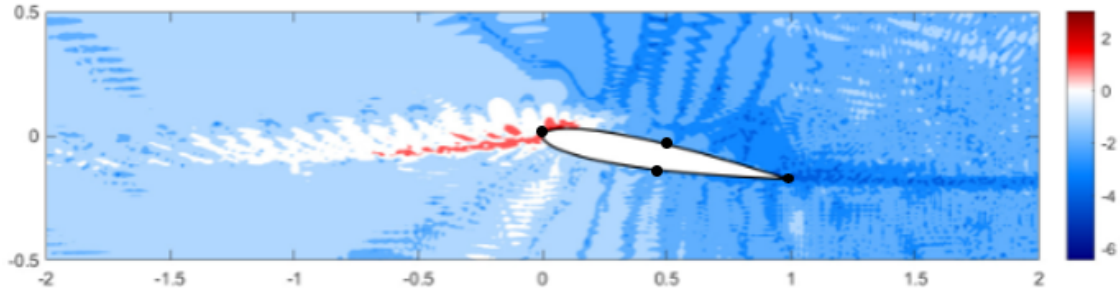
Figure 2.6: Plot of $\log(\alpha_1/\alpha_2)$ for single point forcing.

a selection of points on the surface of the NACA 0012 airfoil located at the leading and trailing edge, as well as the half chord on the pressure and suction side. The first iteration error depends on the ratio α_1/α_2. The spatial distribution of this ratio, as well as the test points examined on the surface are presented in figure 2.6. The first order forcing mode is primarily located near the leading edge of the airfoil, extending over the suction side, resulting in the test points in those regions outperforming the pressure side and trailing edge cases by several orders of magnitude, as presented in figure 2.7. While the initial error strongly depends on the selection of test point, all cases converge at the same rate due to the independence of the convergence rate from α_1/α_2. The convergence rate is instead problem specific, determined by the ratio in singular values σ_1/σ_2.

We also examine the structure of the forcing modes obtained after the first iteration in figure 2.8. The forcing mode calculated using the point placed at the leading edge is practically indistinguishable from the one obtained through the full SVD. The forcing mode calculated using the point at the trailing edge captures the primary features of the forcing mode, but also erroneous structures near the trailing edge. The additional structures are formed due the this point also exciting the second order
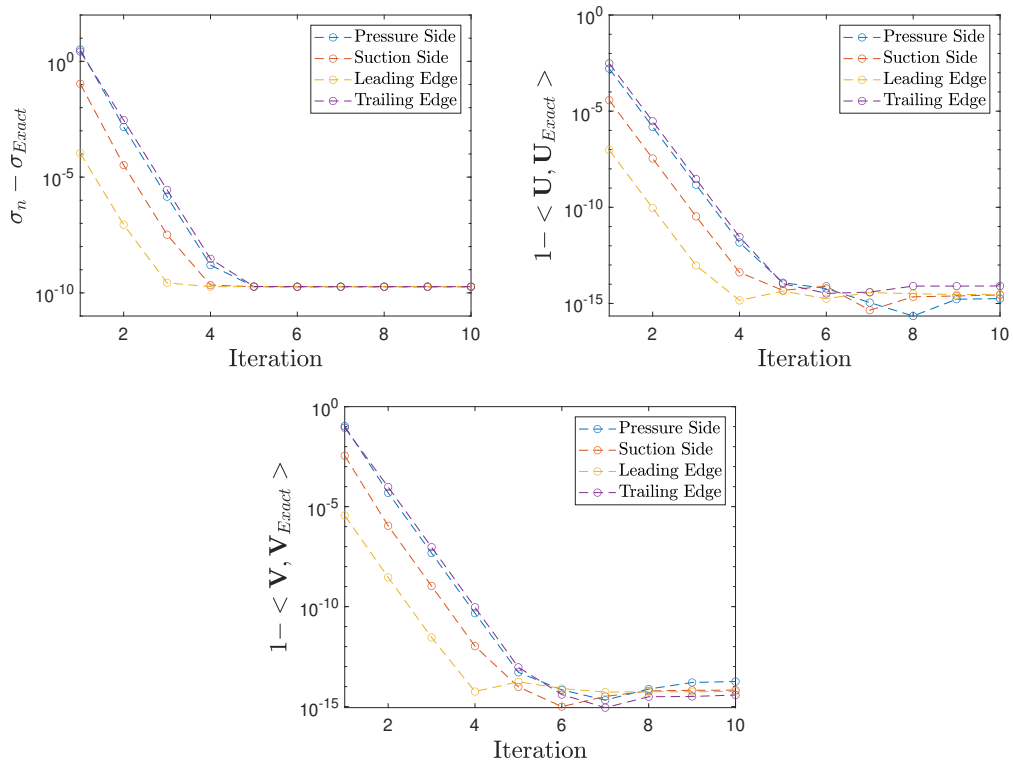
22

Figure 2.7: Singular value error; Cosine similarity error of response modes; Cosine similarity error of forcing modes for a collection of test vectors on the airfoil surface.
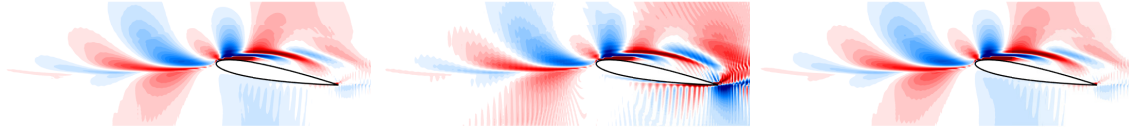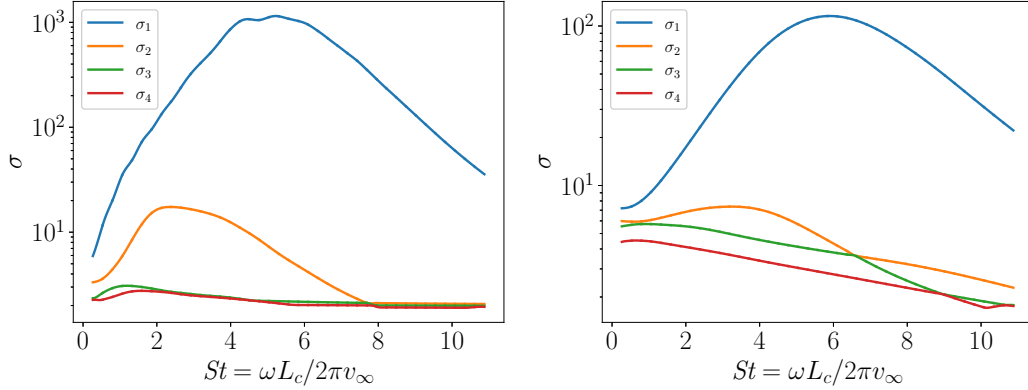
Figure 2.8: Forcing Modes calculated using a point at the leading edge, a point at the trailing edge, and the full SVD respectively

(a) Optimal gains for $k_z L_c = 0$.

(b) Optimal gains for $k_z L_c = 20\pi$.

Figure 2.9: The leading singular values against Strouhal number for two values of the spanwise wavenumber.

mode, which is located primarily in the wake. The forcing modes calculated using this point are therefore a superposition of the first two modes. The vector containing the leading edge point is nearly orthogonal to the second order forcing mode while also aligning much more closely with the first mode, dramatically improving the accuracy of the result.

CHAPTER 3

Application

3.1 Computational Setup

We consider compressible turbulent flow over an extruded two-dimensional cross section of a Honda SUV which is periodic along the z-axis. Simulation results are provided by Honda R&D Co., Ltd. in the form of 2D slices of the 3D flowfield. The streamwise velocity (U_∞), density (ρ_∞), dynamic viscosity (μ_∞) at the free stream are set as 27.77 m/s, 1.205 kg/m^3, 1.822$\times 10^{-5}$ m^2/s, respectively. The free stream Mach number and cross-section based Reynolds number are computed as $M_\infty \equiv U_\infty/a_\infty = 0.1$ and $Re_\infty \equiv \rho U_\infty \sqrt{S}/\mu = 2.5 \times 10^4$, respectively, where S is the streamwise cross-sectional area of the vehicle. The flow is numerically simulated using Large Eddy Simulation (LES) using the Vreman subgrid scale model in the CharLES CFD software. The linearized Navier–Stokes operator $L_{\bar{q}}$ is constructed with respect to the 2D mean flow state $\bar{q}$. The flowfield at the mid slice is shown in figure 3.1. The mean flow is calculated by averaging over a collection of 47 z-slices containing the time averaged flow variables $\bar{q}(x, y, z)$. The resulting mean flow is used for the construction of the linear operator (This section has been adapted from a section originally written by Vedasri Godavarthi for a co-authored paper "Open loop flow control strategy for drag reduction over an SUV using windowed resolvent analysis",

which is currently in preparation).

The resolvent analysis is performed on a seperate grid from that used by the flow solver. The resolvent grid is generally coarser than that need for the LES, however care is taken to resolve the region upstream of the vehicle in order to ensure that forcing modes can be accurately captured. The mesh fills the domain $x/L \in [-27.34, 56.89]$ and $y/L \in [0, 7.19]$ where L is the height of the vehicle. The resulting mesh contains 127959 cells resulting in a linear operator of dimensions 639795×639795. In the process of constructing the linear operator, the surface of the vehicle and the ground are treated as adiabatic walls with the no-slip boundary condition, however the ground has a velocity equal to the free stream. All other boundaries utilize a sponge region to prevent numerical reflections.

3.2   Resolvent Analysis over Honda SUV

We perform resolvent analysis of the base flow over a Honda SUV to understand the dominant input-output mechanisms of the flow. The resolvent operator is formulated for a two dimensional mean flow, under the assumption that the flow over the middle portion of the SUV will be predominantly two dimensional. While there are undoubtedly significant three dimensional contributions to the wake and drag characteristics of the vehicle, this two dimensional approximation should allow for the extraction of key flow behaviors, akin to studying the flow over an airfoil instead of a finite wing. Construction of the resolvent operator for a two dimensional mean flow results in an operator size of $n = 10^5$, which already incurs significant computational expense. The assumption of a two dimensional flow is made for resolvent analysis with hopes that the insights will carry over to a significant part of the overall flow.
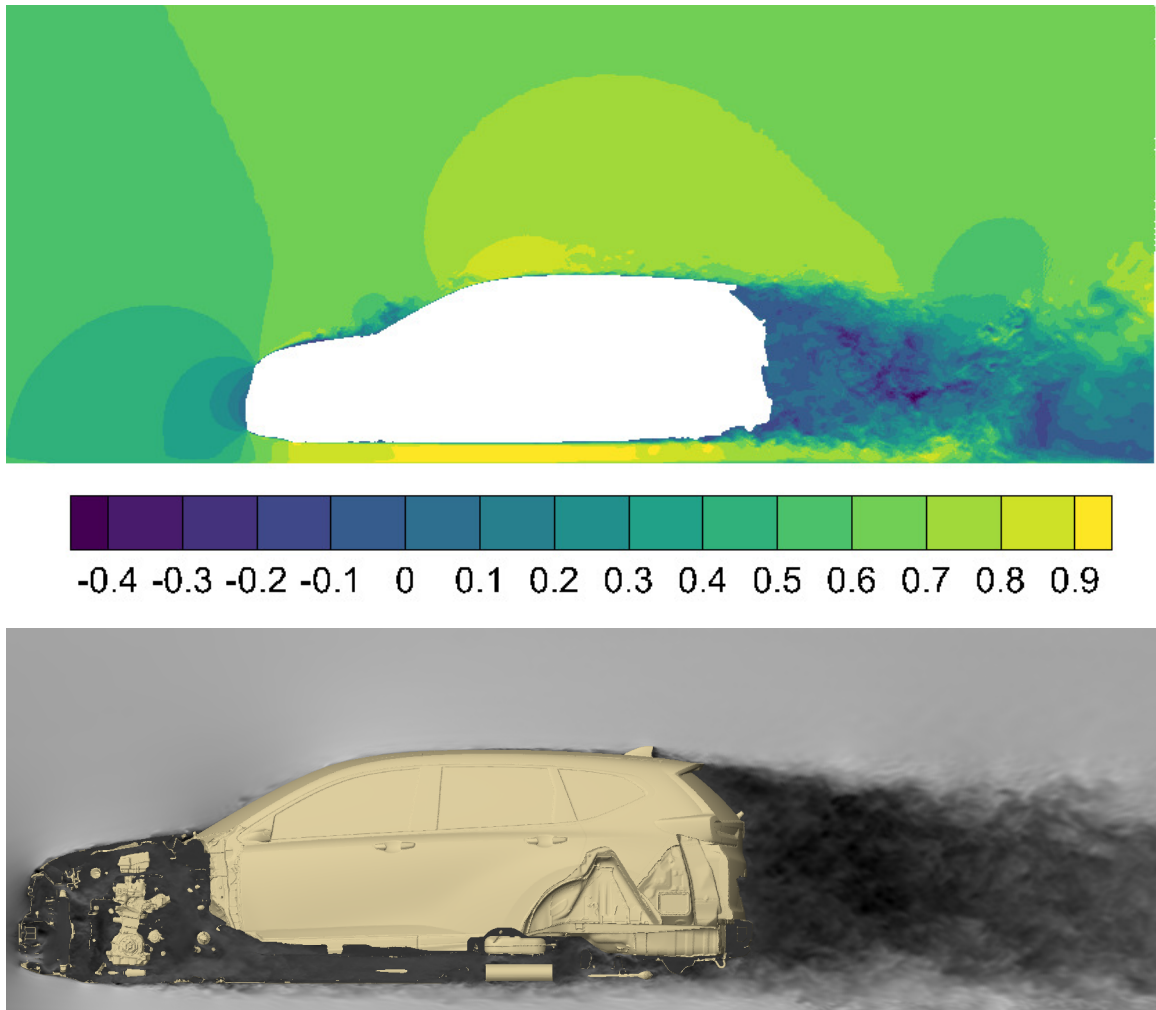
Figure 3.1: Instantaneous u velocity component normalized by the maximum value for flow over a simplified body used for resolvent analysis (top) and the full 3D vehicle geometry (bottom).
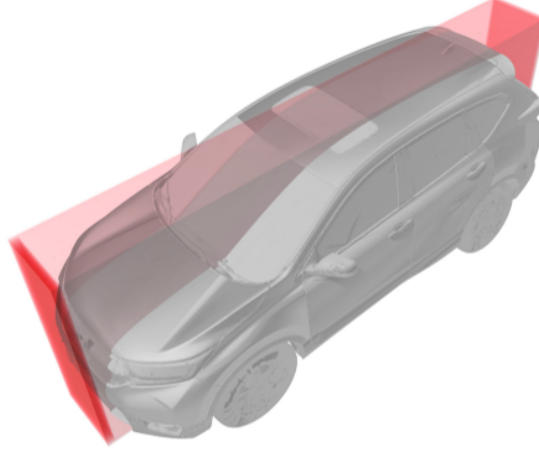
Figure 3.2: The region of focus of this study is highlighted by the red region (middle portion of the vehicle).

To develop guidelines for active flow control, we compute the SVD of the resolvent operator to obtain the forcing and response modes with a corresponding gain. The forcing and response modes are shown in figures 3.3 and 3.4. By examining the forcing modes, we can understand how the flow can be most efficiently actuated. The response mode is also significant, as simply forcing in the optimal manner described in the forcing mode does not guarantee a response which will be aerodynamically beneficial, only one with a large amplification. In order to effectively reduce drag we must induce a response which will nonlinearly alter the mean state of the flow (YT19). Figures 3.3 and 3.4 show a wide range of potential responses, and the effects of these responses on the mean flow can be difficult to predict. Further examination is required to determine which ones may contribute to drag reduction.

Also of great importance is the gain associated with each forcing and response mode pair. The gain distributions for the modes presented in figures 3.3 and 3.4 are presented in figure 3.5. We are generally not concerned with modes with low gains as
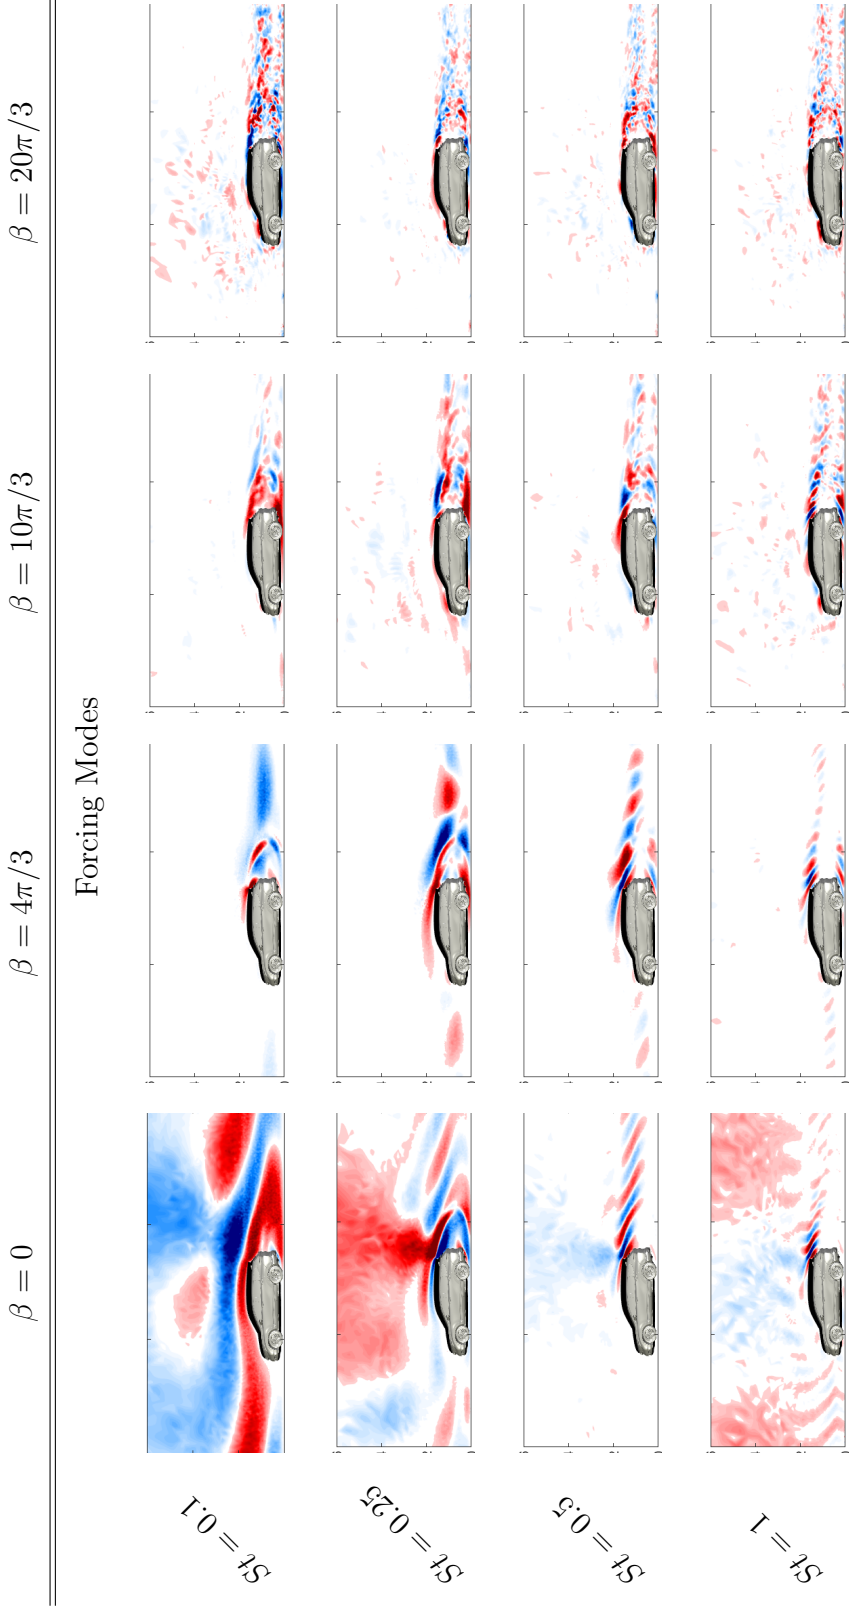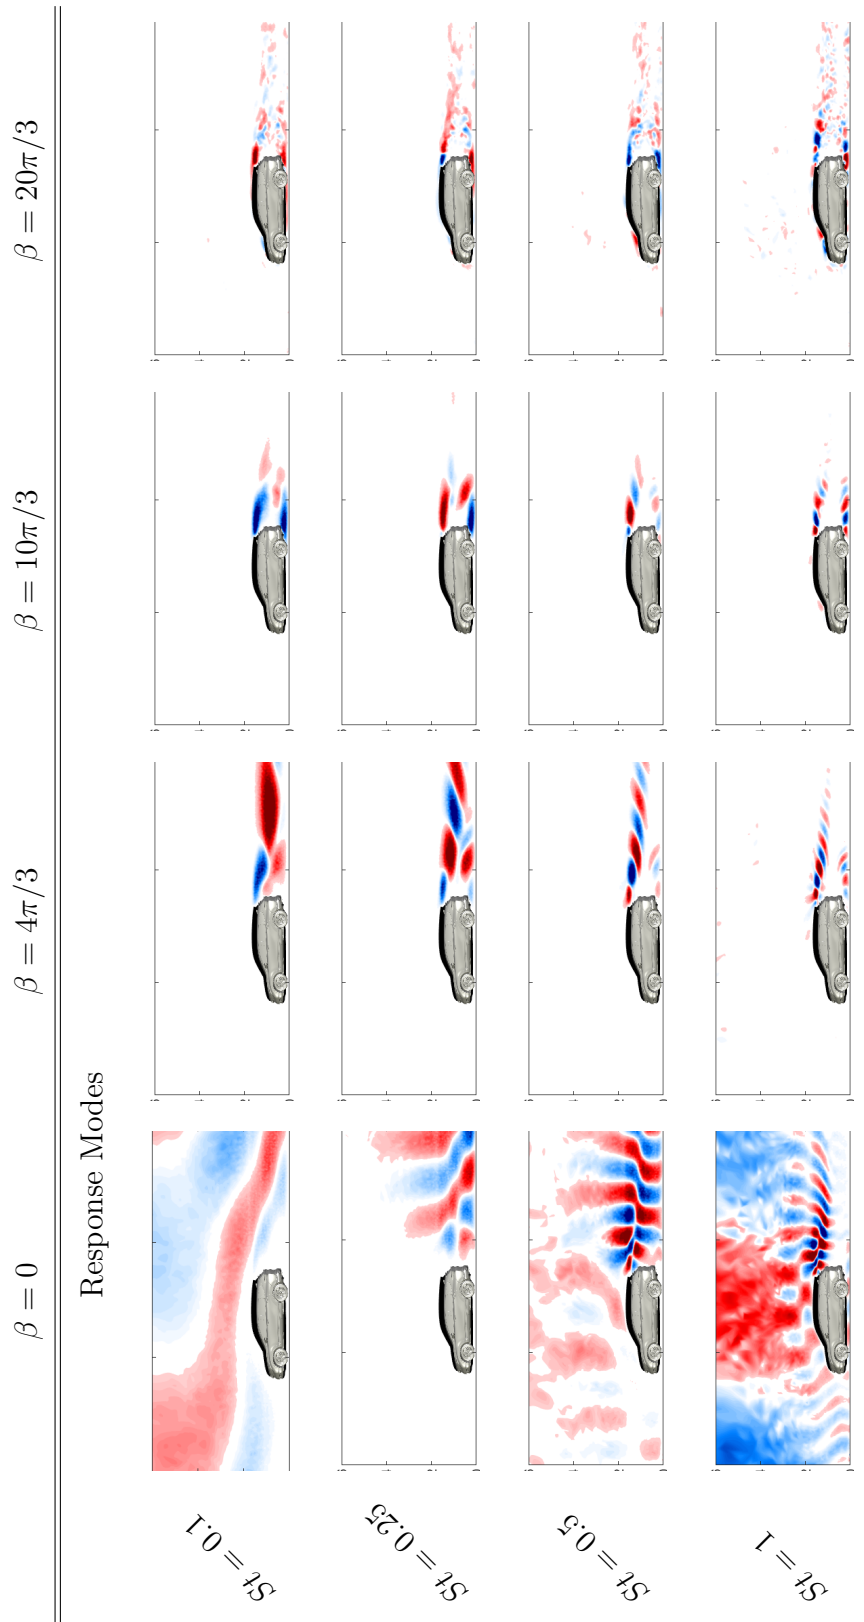
Figure 3.3: Forcing modes for flow over a Honda SUV, for a range of frequencies and wavenumbers. Lower frequencies and wavenumber correspond to larger forcing structures while high frequencies and wavenumbers correspond to much smaller structures

Figure 3.4: Response modes for flow over a Honda SUV, for a range of frequencies and wavenumbers. The response modes follow the same trend as the forcing modes, larger structures are generated for low frequencies and wavenumbers, while the opposite is true for high frequencies and wavenumbers.

actuating according to such a forcing mode will not produce a significant response. The gain, as well as the forcing and response modes, are frequency and wavenumber dependent. The gain plotted against frequency and wavenumber also provides some valuable insight into the flow. Figure 3.6 shows that the flow is most sensitive to two dimensional disturbances at approximately $St = 0.25$, corresponding to the dominant frequency and scale in the wake. Figure 3.7 shows the forcing and response modes corresponding to this gain, highlighting the generation of large wake structures due to a forcing focused near the rear of the vehicle.
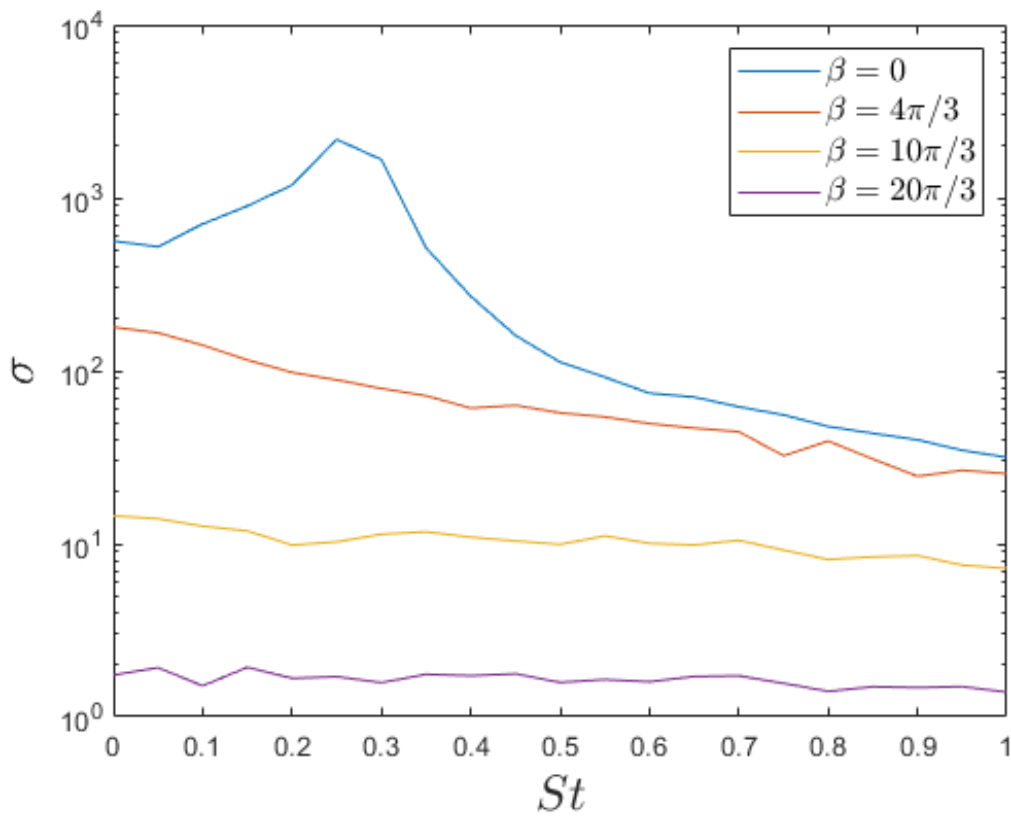

Figure 3.5: The primary gain distributions for St at the wavenumbers presented in figures 3.3 and 3.4
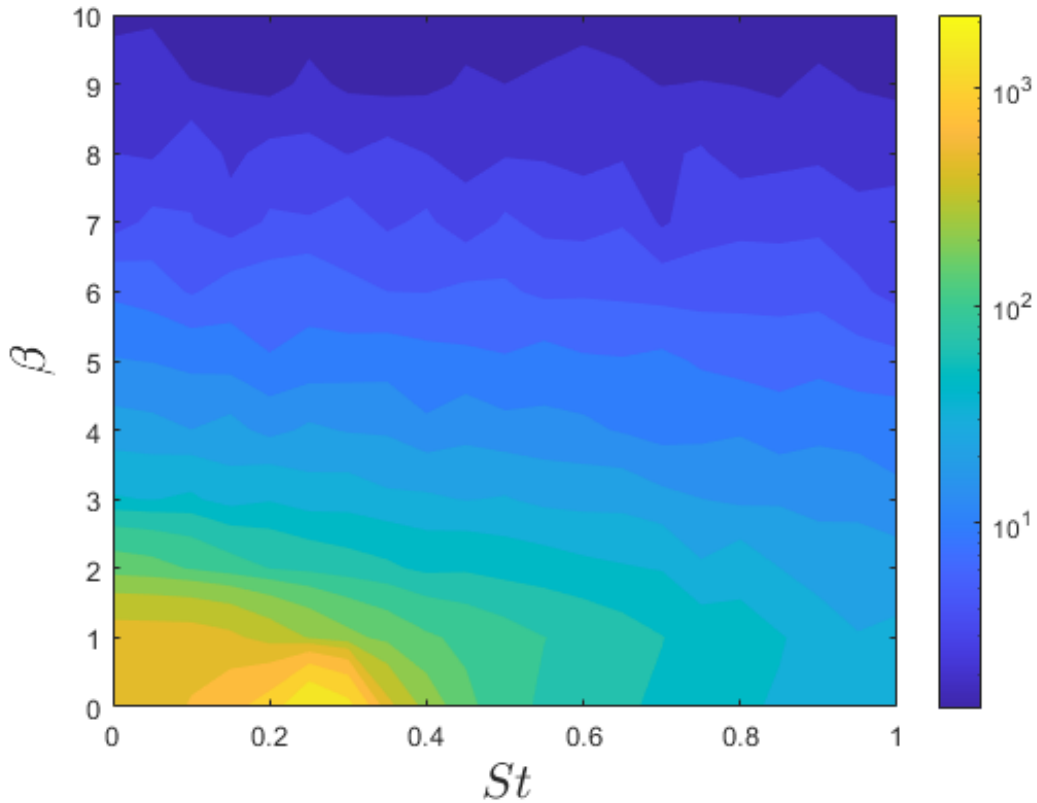
Figure 3.6: Peak in the gain distribution at a peak of $St = 0.25$ corresponding to the dominant vortex shedding frequency. This peak occurs for a wavenumber of 0, indicating the two dimensionality of the wake structures.

The drag characteristics of the vehicle strongly depend on the wake dynamics. Large low pressure structures in the wake of the vehicle are a major contributor to its overall drag. We seek to understand how these structures develop as a result of shear layer interactions over the surface of the vehicle. To accomplish this, we utilize windowed resolvent analysis to examine the sensitivity of the flow to developing structures in the shear layer region. We construct a windowing matrix C corresponding to a 1×0.75 meter region in the flowfield. This window is positioned at various
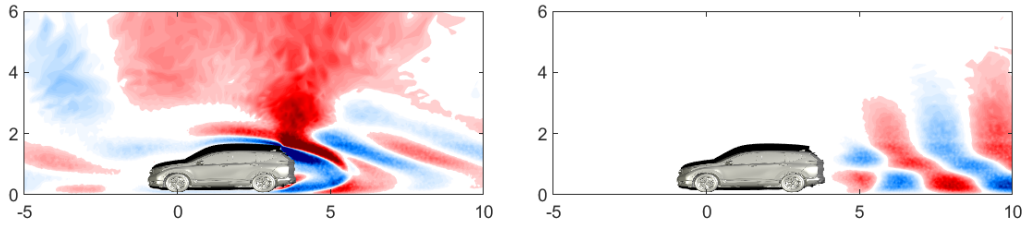
32

Figure 3.7: Global response modes highlight the large amplification occurring in the wake. Global forcing modes reveal that large wake structures can be actuated from the rear of the vehicle. While this is likely undesirable for flow control, this information is useful for understanding where wake structures originate.

regions on top of the car as shown in figure 3.8. Moving this response side window reveals the frequency and structures that can be most effectively generated in the windowed region. Figure 3.9 shows the resulting resolvent gain for several wavenumbers at two frequencies. At the wake frequency of approximately $St = 0.25$, there is a clear trend in the how the gain is affected by the position of the window for different wavenumbers. Higher wavenumbers correspond to smaller structures, which peak in their amplification earlier on in the shear layer. Lower wavenumbers, corresponding to larger structures, peak in their amplification near the back of the vehicle. The wake region is dominated by 2D structures which see the largest overall amplification. At significantly higher and lower frequencies the resulting gain is much smaller. Previous analysis suggested that the modal structures generated by high frequency actuation in the shear layer may be beneficial for drag reduction. While this is true, the gain associated with these frequencies is so much smaller than the gain at frequencies near $St = 0.25$ that we do not consider these higher frequency modes to be suitable for control.
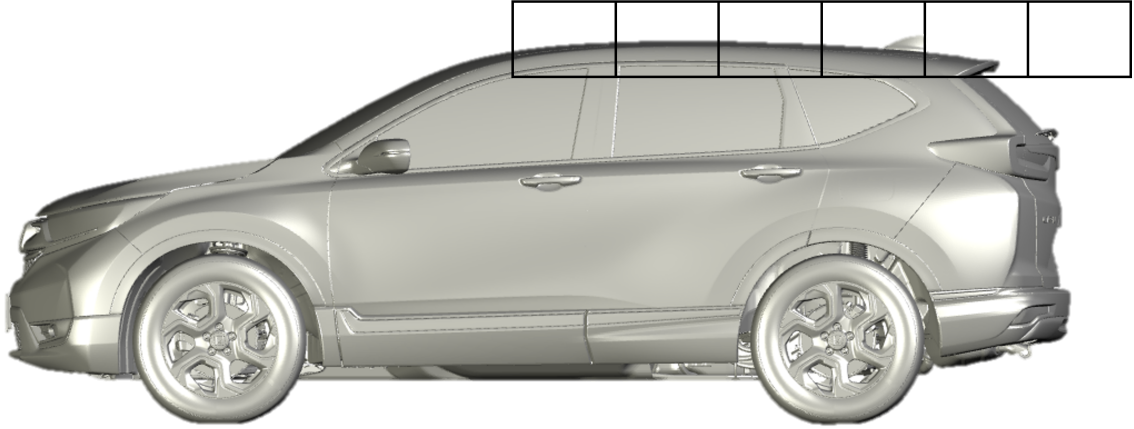
Figure 3.8: Positions of resolvent window over the surface of the vehicle. These windows limit the space in which response modes can be generated, allowing us to isolate the behavior of the flow in a specific region.

This result inspires a potential approach to active flow control. Delaying the formation of large scale structures in the wake could result in drag reduction (GTN22). By utilizing the natural tendency of the flow to amplify smaller three dimensional structures in the shear layer, we may be able to delay the formation of large two dimensional structures which create large drag forces. Generating additional smaller structures earlier in the shear layer and larger structures later, where they will be most greatly amplified, could slow the process by which these structures naturally merge into larger structures in the wake. Delaying these structures until later in the wake can increase the pressure on the rear surface of the body, reducing drag. In order to determine how to generate these structures, we can utilize forcing side windowing to determine where and how to actuate the flow.

We apply a forcing side windowing matrix B corresponding to points on the surface of the vehicle to reflect possible positions for actuator placement. We can see from
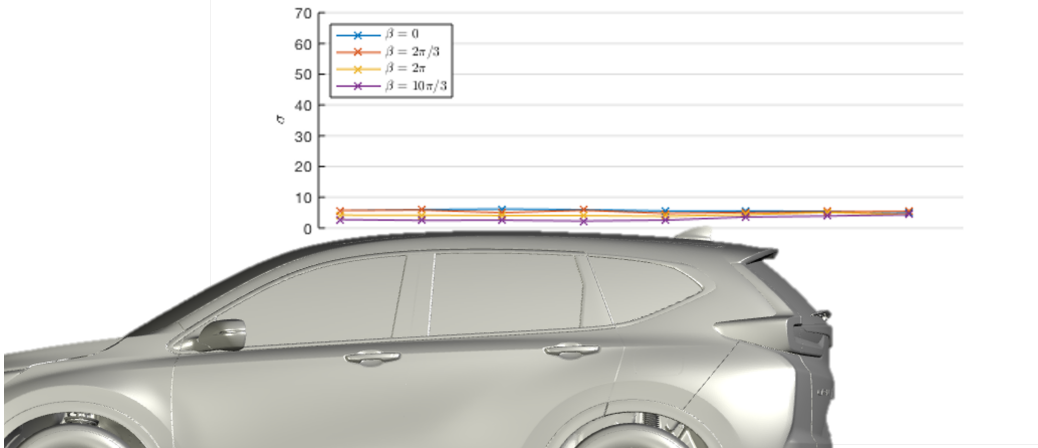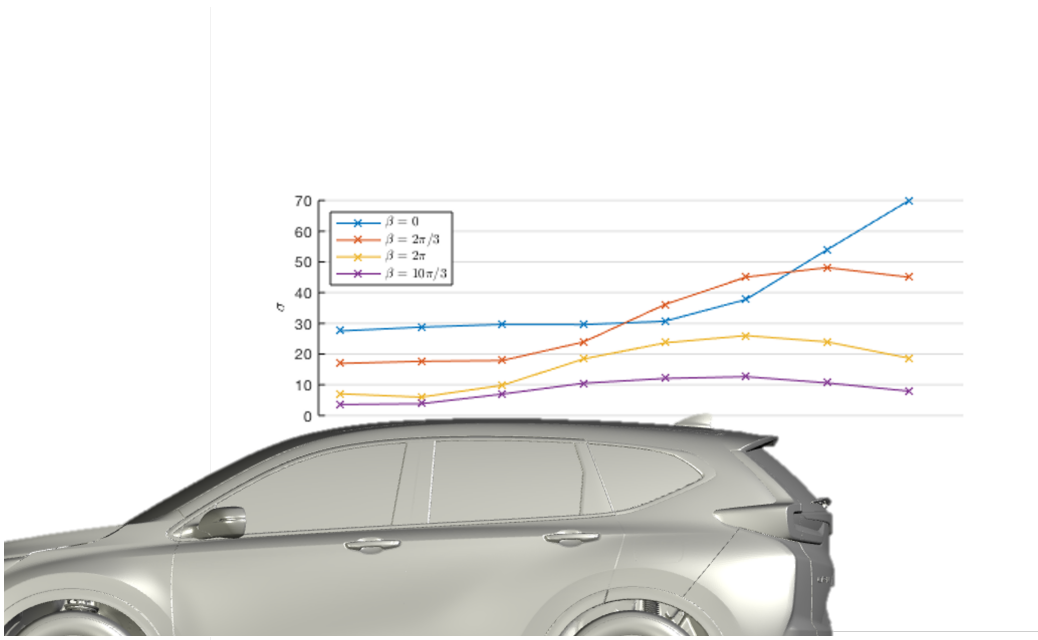
34

Figure 3.9: Resolvent gain for response side windows at several positions over the top of the vehicle at frequency of $St = 0.25$ (left) and $St = 2.5$ (right). The gain for frequencies other that the wake frequency are relatively small.

figure 3.10 that in order to produce structures early in the shear layer, near the front of the roof, we would need to actuate the flow near the top of the windshield. As we move the response window along the surface of the roof, the optimal forcing positions for higher wavenumbers shifts to follow the response window. For lower wavenumbers, the optimal forcing position moves very little if at all, remaining near the top of the windshield. This behavior can be intuitively understood as larger scale structures being more persistent in the flow, making it possible to generate them further upstream and allowing them more time to be amplified by the natural mechanisms of the flow. Smaller structures are less persistent, and must be generated near where we want these structures to appear.

We also examine the direction of the forcing modes, which can be used to inform the optimal direction for a synthetic jet actuator. The forcing directions illustrated in figure 3.10 show the direction and relative strength for actuators placed at the point where the forcing mode is largest. The actuation direction for almost all modes is at an approximately 45 degree angle from the surface of the car. Modes actuated near the front of the roof are actuated with blowing in the same direction as the free stream flow while higher wavenumber modes near the rear of the vehicle are actuated by blowing in the opposite direction. We note that all of the modes of resolvent analysis are sinusoidal with respect to time. While it is possible for the x and y components of the velocity to be out of phase with one another, for all cases studied here the phase difference is very small, on the order of 10^{-5}. This is fortunate for potential real world applications, as any potential method of actuation only need be able to apply forcing along one dimension instead of two.

The combined information from the windowed forcing and response modes results in a comprehensive recommendation for potential flow control. Actuation should be

36

performed near the dominant wake frequency to obtain the largest possible amplification for any input. We hope that by introducing three-dimensionality into the flow, the formation of two-dimensinoal structures in the wake can be delayed, potentially leading to drag reductino. We wish to generate small structures early in the shear layer, and larger structures later, corresponding to the regions of the flow where they will be most greatly amplified according to the windowed resolvent analysis. Larger structures should be actuated near the top of the windshield since they can persist in the flow over the entire length of the vehicle, allowing them more time to be amplified. Smaller structures which do not last long in the flow must instead be actuated much later in the shear layer, near the corresponding response mode. These results are summarized in figure 3.12.

3.3   Conclusion

Resolvent analysis is a powerful tool, capable of providing valuable insights into the input-output behaviour of turbulent flows. While the computational burden of this technique can be significant, sketching the resolvent operator can significantly redcuce the compuation time and memory requirements of the SVD. While a randomized test matrix is suitable for many cases, performance can be improved through careful selection of the a test vector when higher order modes are not necessary. Using a test vector makes it relatively simple to examine how exactly the sketching process affects the accuracy of the algorithm. The error analysis of the algorithm suggest that the test matrix can be sparse without a loss in performance. We show that when properly selected, a single point test vector can outperform a full test vector.

An important term in the error analysis is the inner product $\alpha_1 = \langle V_1, \Omega \rangle$. The

improvement of the accuracy for test vectors more closely aligned with the forcing mode leads us to formulate this algorithm as an iterative process, where the test vector in the next step is the forcing mode calculated in the current step. Utilizing this iterative algorithm we can obtain results in excellent agreement with traditional SVD methods using arbitrary test vectors, with the number of iterations necessary for convergence significantly reduced by careful selection of the initial test vector.

This method is then used to analyze the flow over a Honda SUV. The global forcing and response modes grant insight into the behavior of the flow. Specifically, in order to reduce drag we wish to alter the wake dynamics, as large structures in the wake are responsible for a large amount of the total drag on the vehicle. To understand how to alter the wake dynamics we perform windowed resolvent analysis, utilizing a moving window to study how the response changes as the window moves from the shear layer into the wake.

Windowed resolvent analysis reveals the dependence of the gain on the scale of the structures. Higher wavenumber structures corresponding to small three dimensional structures peak in their gain over the vehicles surface. As we proceed to larger scale structures with lower wavenumbers, we find that the maximum gain for these structures occurs at the very rear of the vehicle. Large two dimensional structures dominate the wake itself. These results inspire a potential flow control method, by generating a smaller scale structures earlier in the shear layer, and larger structures later, we can introduce a large amount of three dimensionality to the flow.

In order to understand how to actuate these structures we study the body windowed forcing modes. We find that the flow is generally sensitive near the top of the windshield, however smaller scale structures will not persist long in the flow and must be actuated nearer to where we desire the response to be amplified. Larger structures

38

can also be actuated at the front of the roof, allowing them more time to be amplified by the flow.
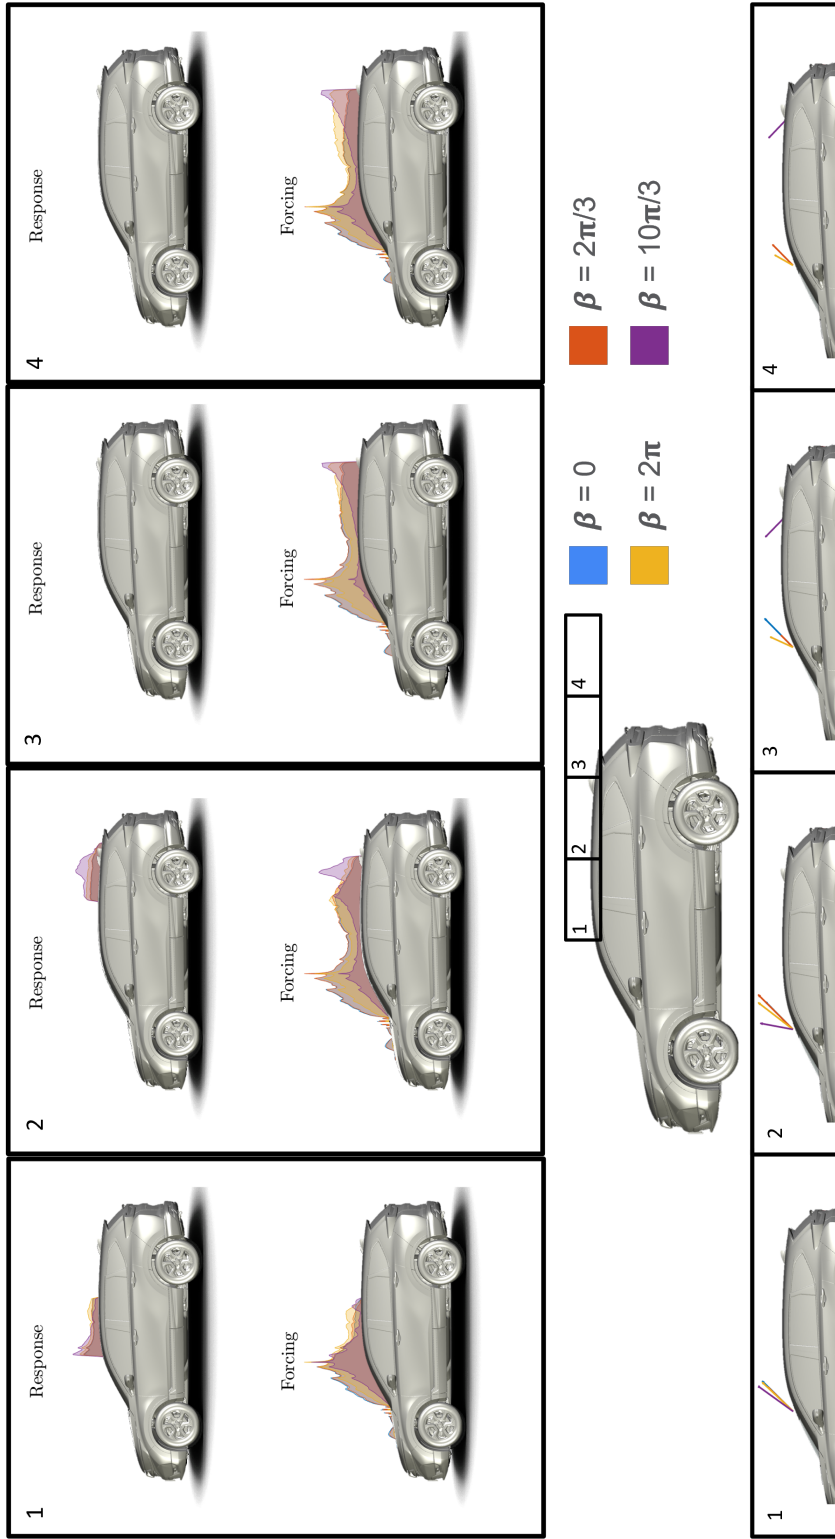
Figure 3.10: Forcing and response mode strength over the surface of the vehicle. The forcing and response modes are shown for a collection of response windows. The forcing window is constant, corresponding to the surface of the vehicle. 50 centimeters has been truncated from the forcing window in the front and back of the car to prevent any effect from the singularity at the beginning and end of the boundary layer.
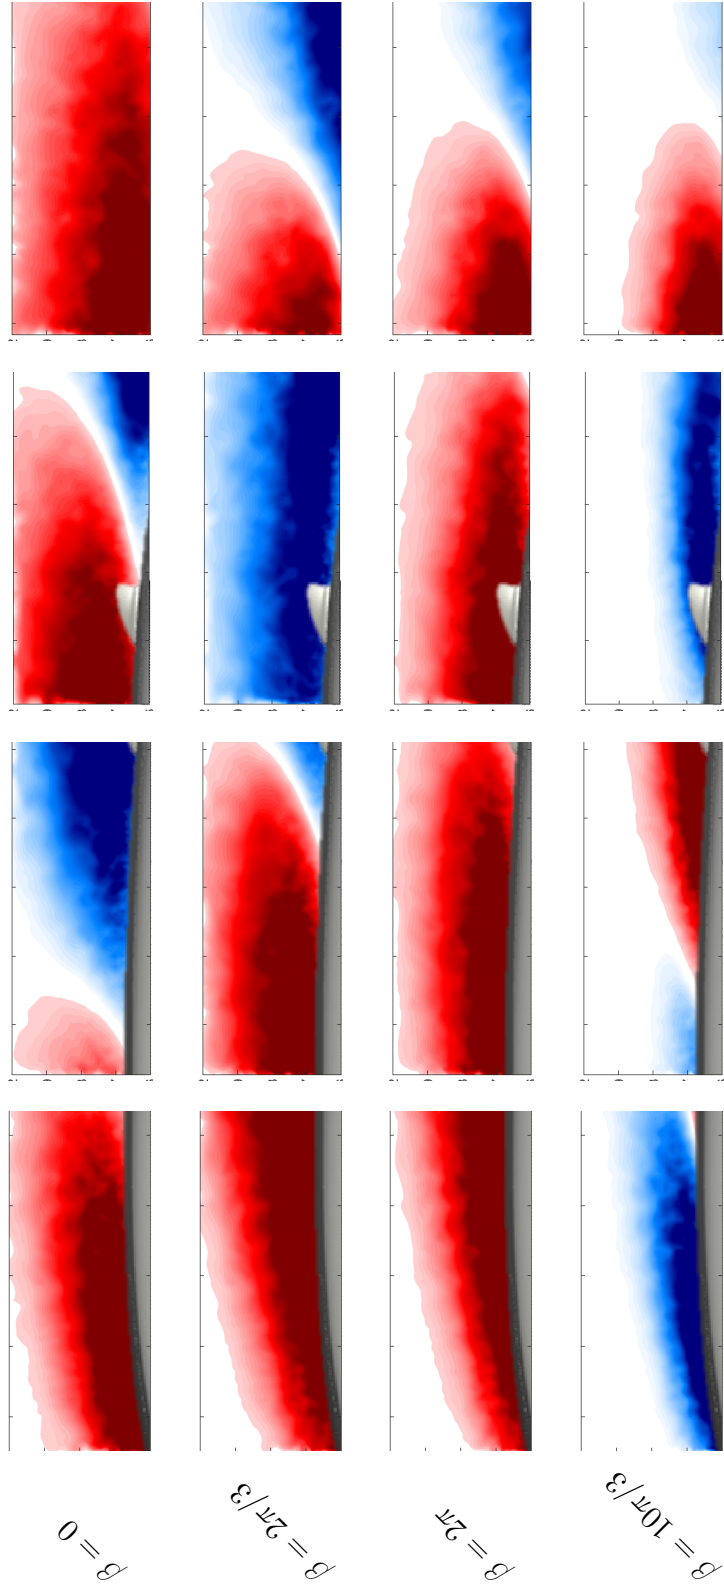
Figure 3.11: Response modes corresponding to the forcing modes shown in figure 3.10
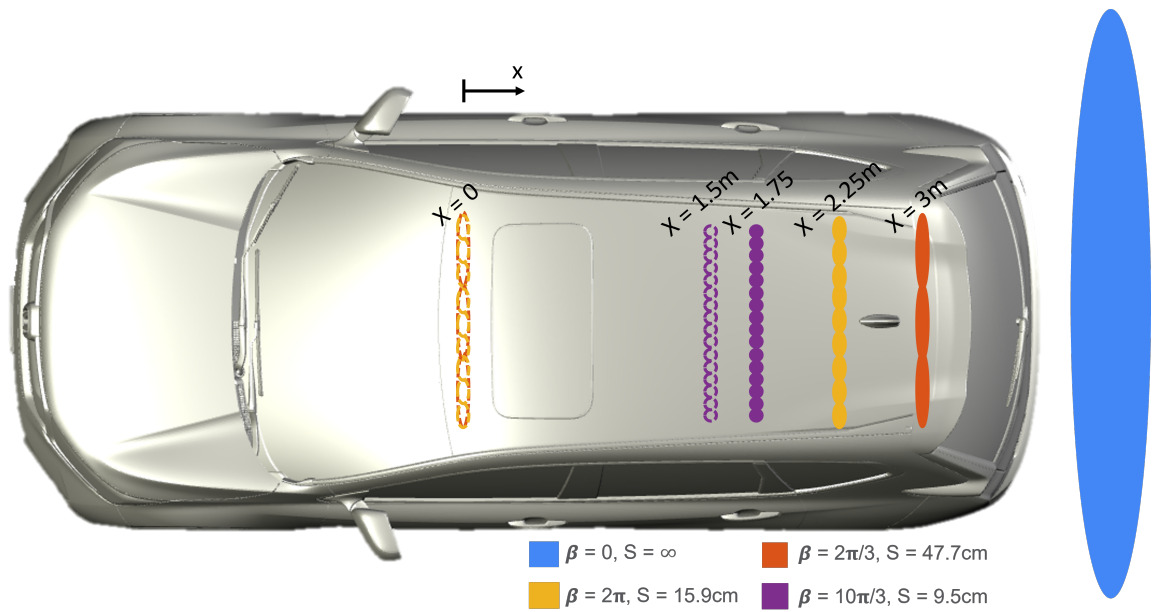
Figure 3.12: Recommendations for flow control over a Honda SUV. Dashed structures corresponding to regions of actuation, while solid structures correspond to the regions in which they will be amplified. Colors are labeled by wave number β, as well as the corresponding spacing S. All x positions are measured with respect to the top of the windshield.

42

Bibliography

[AR84]   S. R. Ahmed and G. Ramm. "Some Salient Features of the Time-Averaged Ground Vehicle Wake." SAE-Paper 840300, 1984.

[BHL93]   G. Berkooz, P. Holmes, and J. L. Lumley. "The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows." *Annual Review of Fluid Mechanics*, **25**(1):539–575, 1993.

[BHN17]   G. A. Brès, F. E. Ham, J. W. Nichols, and S. K. Lele. "Unstructured Large-Eddy Simulations of Supersonic Jets." *AIAA J.*, **55**(4):1164–1184, 2017.

[EVB19]   N. Benjamin Erichson, Sergey Voronin, Steven L. Brunton, and J. Nathan Kutz. "Randomized Matrix Decompositions Using R." *Journal of Statistical Software*, **89**(11), 2019.

[FI93]   B. F. Farrell and P. J. Ioannou. "Stochastic forcing of the linearized Navier–Stokes equations." *Physics of Fluids A: Fluid Dynamics*, **5**(11):2600–2609, 1993.

[FI94]   B. F. Farrell and P. J. Ioannou. "Variance maintained by stochastic forcing of non-normal dynamical systems associated with linearly stable shear flows." *Physical Review Letters*, **72**:1188–1191, Feb 1994.

[GTN22]   V. Godavarthi, K. Taira, Y. Nakamura, and K. Ambo. "Cluster-based analysis of a SUV wake." AIAA 2022 SciTech meeting, submitted, 2022.

[HLB12]   P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. *Turbulence,*

coherent structures, dynamical systems and symmetry. Cambridge Univ. Press, 2nd edition, 2012.

[HMT11] N. Halko, P.-G. Martinsson, and J. A. Tropp. "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions." *SIAM review*, **53**(2):217–288, 2011.

[JB05] M. R. Jovanovic and B. Bamieh. "Componentwise energy amplification in channel flows." *J. Fluid Mech.*, **534**:145–183, 2005.

[Jov04] M. R. Jovanović. *Modeling, analysis, and control of spatially distributed systems.* PhD thesis, University of California at Santa Barbara, Dept. of Mechanical Engineering, 2004.

[KNH11] Y. Khalighi, J. W. Nichols, F. Ham, S. K. Lele, and P. Moin. "Unstructured Large Eddy Simulation for Prediction of Noise Issued from Turbulent Jets in Various Configurations." 17th AIAA/CEAS Aeroacoustics Conference, 2011.

[Kut13] J. N. Kutz. *Data-driven modeling and scientific computation.* Oxford Univ. Press, 2013.

[MS10] B. J. McKeon and A. S. Sharma. "A critical-layer framework for turbulent pipe flow." *J. Fluid Mech.*, **658**:336–382, 2010.

[NYK19] Aditya G. Nair, Chi-An Yeh, Eurika Kaiser, Bernd R. Noack, Steven L. Brunton, and Kunihiko Taira. "Cluster-based feedback control of turbulent post-stall separated flows." *Journal of Fluid Mechanics*, **875**:345–375, 2019.

[RYT20] J. H. M. Ribeiro, C.-A. Yeh, and K. Taira. "Randomized resolvent analysis." *Phys. Rev. Fluids*, **5**:033902, Mar 2020.

[Sch07] P. J. Schmid. "Nonmodal Stability Theory." *Annual Review of Fluid Mechanics*, **39**(1):129–162, 2007.

[Sch10] P. J. Schmid. "Dynamic mode decomposition of numerical and experimental data." *J. Fluid Mech.*, **656**:5–28, 2010.

[TBD17] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley. "Modal analysis of fluid flows: An overview." *AIAA J.*, **55**(12):4013–4041, 2017.

[TE05] L. N. Trefethen and M. Embree. *Spectra and pseudospectra.* Princeton Univ. Press, 2005.

[THB19] K. Taira, M. S. Hemati, S. L. Brunton, Y. Sun, K. Duraisamy, S. Bagheri, S. Dawson, and C.-A. Yeh. "Modal Analysis of Fluid Flows: Applications and Outlook." *AIAA J.*, **58**(3):998–1022, 2019.

[The11] V. Theofilis. "Global linear instability." *Annu. Rev. Fluid Mech.*, **43**:319–352, 2011.

[TTR93] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. "Hydrodynamic stability without eigenvalues." *Science*, **261**(5121):578–584, 1993.

[Vre04] A. W. Vreman. "An eddy-viscosity subgrid-scale model for turbulent shear flow: algebraic theory and applications." *Phys. Fluids*, **16**(10):3670–3681, 2004.

[YBT20] C.-A. Yeh, S. I. Benton, K. Taira, and D. J. Garmann. "Resolvent Analysis

of an Airfoil Laminar Separation Bubble at $Re = 500,000$." *Physical Review Fluids*, **5**:083906, Aug 2020.

[YT19] C.-A. Yeh and K. Taira. "Resolvent-analysis-based design of airfoil separation control." *J. Fluid Mech.*, **867**:572–610, 2019.