AFRL-AFOSR-VA-TR-2024-0234

Resolvent-based estimation for control of turbulent aerodynamic flows

Towne, Aaron
REGENTS OF THE UNIVERSITY OF MICHIGAN
1109 GEDDES AVE, SUITE 3300
ANN ARBOR, MI, 48109
USA

05/30/2024
Final Technical Report

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
20240530	Final	START DATE 20200806	END DATE 20240205

4. TITLE AND SUBTITLE
Resolvent-based estimation for control of turbulent aerodynamic flows

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
	FA9550-20-1-0214	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Aaron Towne

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
REGENTS OF THE UNIVERSITY OF MICHIGAN
1109 GEDDES AVE, SUITE 3300
ANN ARBOR, MI 48109
USA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
Air Force Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203	AFRL/AFOSR RTA1	AFRL-AFOSR-VA-TR-2024-0234

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
In this project, we developed, implemented, and demonstrated an optimal resolvent-based estimation and control framework for turbulent aerodynamic flow fields using limited measurements on the surface of an aircraft. While state estimation and control are classical
topics in dynamical systems and control theory, standard methods have several disadvantages when applied to turbulent flows, including high costs and restrictions in their ability to incorporate key physics. We overcame these limitations by leveraging state-of-the-art models
for coherent structures based on resolvent analysis. Specifically, we derived an optimal resolvent-based estimator and controller with several advantages over standard methods. When equivalent assumptions are made, the resolvent-based estimator and controller
reproduce the classical Kalman filter and LQG controller, respectively, but at substantially lower computational cost using an efficient time-stepping method for constructing the kernels. Unlike these standard methods, the resolvent-based approach can naturally accommodate
forcing terms (nonlinear terms from Navier-Stokes) with colored-in-time statistics, significantly improving the accuracy of the methods. When desired, causality is optimally enforced via a Wiener-Hopf decomposition.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU	31

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
GREGG ABATE	425-1779

Standard Form 298 (Rev.5/2020)
Prescribed by ANSI Std. Z39.18

Final Research Performance Progress Report:

AFOSR Award FA9550-20-1-0214

Resolvent-based estimation for control of turbulent aerodynamic flows

Principal Investigator:

Aaron Towne

Assistant Professor

Department of Mechanical Engineering

University of Michigan

Tel: (734) 647-5278

Email: towne@umich.edu

Program Officer:

Dr. Gregg Abate

Unsteady Aerodynamics & Turbulent Flows

Air Force Office of Scientific Research (AFOSR)

Tel: (703) 696-7917

Email: gregg.abate@us.af.mil

Grant period:

6 August 2020 – 5 February 2024

Abstract

In this project, we developed, implemented, and demonstrated an optimal resolvent-based estimation and control framework for turbulent aerodynamic flow fields using limited measurements on the surface of an aircraft. While state estimation and control are classical topics in dynamical systems and control theory, standard methods have several disadvantages when applied to turbulent flows, including high costs and restrictions in their ability to incorporate key physics. We overcame these limitations by leveraging state-of-the-art models for coherent structures based on resolvent analysis.

Specifically, we derived an optimal resolvent-based estimator and controller with several advantages over standard methods. When equivalent assumptions are made, the resolvent-based estimator and controller reproduce the classical Kalman filter and LQG controller, respectively, but at substantially lower computational cost using an efficient time-stepping method for constructing the kernels. Unlike these standard methods, the resolvent-based approach can naturally accommodate forcing terms (nonlinear terms from Navier-Stokes) with colored-in-time statistics, significantly improving the accuracy of the methods. When desired, causality is optimally enforced via a Wiener-Hopf decomposition.

We implemented the resolvent-based estimation and control methods within a high-fidelity solver, CharLES. All routines are fully parallelized, allowing us to apply our tools to any flow that can be simulated in CharLES. Additionally, a data-driven implementation expands the applicability of these methods to experimental settings and simulations that lack adjoint capabilities. Moreover, the method admits an approach to efficiently explore sensor and actuator placement by avoiding redundant calculations and obtaining an a priori estimate of the control performance without applying the controller to the flow. After several proof-of-concept test cases, we applied the resolvent-based tools to a series of flows over airfoils, including the laminar flow over an airfoil at Reynolds number 5000 immersed in clean and noisy freestreams and the turbulent flow over an airfoil at Reynolds number 23000. In all cases, the optimal causal resolvent-based estimator outperforms a previously employed method of truncating a noncausal kernel.

Additionally, we extended the time-stepping method used to compute the resolvent-based kernels and paired it with a randomized singular value decomposition to create a new algorithm to efficiently compute resolvent modes for large systems. This combination of methods overcomes the primary computational bottlenecks of previous methods and yields an algorithm that scales linearly with problem size, drastically reducing CPU and memory costs for large systems and extending the applicability of resolvent analysis to previously intractable engineering-scale problems.

Overall, our accomplishments during this project met the original project objective of developing a resolvent-based estimator and applying it to aerodynamic flow and exceeded it by also addressing the control problem and introducing the new algorithm for resolvent analysis. Moving forward, these capabilities will aid in the Air Force's mission to maintain its technological superiority by enabling physics-based prediction and closed-loop control leading to aerodynamic performance improvements.

1 Project goals and objectives

1.1 Overview

Practical limitations restrict the number and types of measurements that can be performed in practical aerodynamic applications. For example, measurements of the flow around an aircraft may be limited to sparse, localized measurements of pressure or shear stress along selected surfaces [1, 2]. Flow estimation is the process of using limited measurements to approximate the state of the flow at other locations and/or for other quantities that are not measured.

The objective of this project was to develop an improved framework for obtaining real-time estimates of turbulent aerodynamic flow fields using limited measurements by leveraging organized motions within the flow and state-of-the-art models for describing them. State estimation is a classical topic in controls theory, but many standard approaches are poorly suited for turbulent flows and fail to take advantage of physical insights into the flow physics. We proposed to develop an improved method by leveraging the existence of organized motions within the flow, often called coherent structures. These coherent structures exhibit correlation over long time and length scales and thus provide a principled, physics-based means to estimate the state of a flow at one location using measurements taken at another location.

Motivated by this insight, we proposed to develop an optimal estimation method based on resolvent analysis of the Navier-Stokes equations. Resolvent analysis identifies modes that are optimal in terms of their gain between the nonlinear terms in the Navier-Stokes equations and the flow field and has been shown to efficiently capture coherent structures in a variety of different flows, including flows around aerodynamic bodies. Our method builds on previous work on resolvent analysis but introduces a new approach for using it to estimate the state of the flow from sparse measurements. We proposed to evaluate the performance of our resolvent-based estimation method by applying it to the separated flows around airfoils at several angles of attack and comparing our results to those obtained from standard estimation methods.

The following subsections briefly introduce resolvent analysis and recap the proposed objectives to provide a baseline against which our accomplishments, reported in Section 2, can be measured.

1.2 Introduction to resolvent analysis

Resolvent analysis is derived from the Navier-Stokes equations, represented here as

$$\frac{\partial \boldsymbol{q}}{\partial t} = \mathcal{F}(\boldsymbol{q}), \tag{1}$$

where $\boldsymbol{q}$ is a state vector of flow variables that describe the flow, i.e., velocities and thermodynamic variables. Applying the Reynolds decomposition

$$\boldsymbol{q}(x,t) = \bar{\boldsymbol{q}}(x) + \boldsymbol{q}'(x,t) \tag{2}$$

to Eq. (1) and isolating the terms that are linear in $\boldsymbol{q}'$ yields an equation of the form

$$\frac{\partial \boldsymbol{q}'}{\partial t} - \boldsymbol{A}(\bar{\boldsymbol{q}})\boldsymbol{q}' = \boldsymbol{f}(\bar{\boldsymbol{q}}, \boldsymbol{q}'), \tag{3}$$

where $\boldsymbol{A}$ is the linearized Navier-Stokes operator and $\boldsymbol{f}$ contains the remaining nonlinear terms, which are conceptualized as a forcing on the linear dynamics.

In the frequency domain, Eq. (3) can be manipulated to give an input/output relationship between the nonlinear forcing term and the state [3, 4],

$$\hat{\boldsymbol{q}} = \boldsymbol{R}_q \hat{\boldsymbol{f}}, \tag{4a}$$

where

$$\boldsymbol{R}_q(x,\omega) = (i\omega \boldsymbol{I} - \boldsymbol{A})^{-1} \tag{5a}$$

is called the resolvent operator. Typically, resolvent analysis proceeds by computing the singular value decomposition (SVD) of the resolvent operator,

$$\boldsymbol{R}_q = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*. \tag{6}$$

The columns of $\boldsymbol{U}$ and $\boldsymbol{V}$ constitute orthogonal bases for the state and the forcing, respectively, ordered by the gain between them, which is given by the singular values within $\boldsymbol{\Sigma}$.

Resolvent analysis has become the dominant tool for studying energy amplification mechanisms within turbulence flows and is arguably the most important operator-theoretic tool in the field [3, 5, 6]. The leading left singular vectors of the resulting resolvent operator have been shown to provide an excellent model for organized motions within turbulent flows called coherent structures [7]. These coherent structures are of central interest within the study of turbulence due to their significant impact on quantities of engineering interest (e.g., drag, heat transfer, momentum transfer, noise emissions) as well as their potential to serve as building blocks toward an improved theoretical understanding of turbulence [5]. Resolvent analysis has been used as a starting point for achieving numerous objectives, including design optimization [8], flow reconstruction [9, 10], receptivity analysis [11], and open-loop control [12, 13]. In this project, we use it to derive optimal estimation and closed-loop control methods.

1.3 Objectives

Here, we briefly review the specific objectives specified in the project proposal.

Objective 1: Develop an optimal resolvent-based estimation method

The primary objective of the project is to develop a new, optimal approach for using resolvent analysis for flow estimation. While the basic resolvent formulation is the same as in previous work, the way in which it will be used to estimate the flow using sparse measurements is entirely new. Estimation problems implicitly involve solving an under-determined inverse problem due to there being more unknown values than there are known values from the measurements. As a result, solving the inverse problem entails selecting a regularization to make the solution unique. We will explore several possible regularizations and also attack the problem directly by formulating and solving appropriate optimization problems to minimize the expected value of the estimation error. This will include careful consideration of the influence of measurement noise to ensure robust solutions, which is critically important for the practical deployment of the method.

Objective 2: Determine the maximum accuracy of resolvent-based models for aerodynamic flows

Upper bounds on the accuracy of any resolvent-based description of coherent structures can be determined using SPOD, which serves as a best-case point of comparison, following the theory developed in [7]. Additionally, we will derive an approach for choosing the types and locations of sensors to maximize the accuracy of the resolvent-based estimation method.

Objective 3: Assess the performance of resolvent-based estimation for aerodynamic flows

We will apply our estimation methodology to flows relevant to aerodynamic applications, specifically the separated, turbulent flow around an airfoil matching the conditions previously studied by Yeh & Taira [13] in their study of resolvent-informed open-loop control. The chord-based Reynolds number and free-stream Mach number are 23,000 and 0.3, respectively. Following their computational approach, the flow will be computed via large-eddy simulation (LES) using the CharLES code from Cascade Technologies, Inc. We will use measurements along the airfoil surface to predict the surrounding flow field, with emphasis on the wake. The quality of our estimation results will be assessed via comparisons with simulation data and other estimation methods.

2 Accomplishments and Technical Updates

In this section, we report the main accomplishments achieved during this project. To clearly link these accomplishments with the proposed project, the objective associated with each accomplishment is indicated in parentheses. Further details about each accomplishment can be found in the cited publications (see also Section 4). In aggregate, we believe we have met or exceeded all project objectives.

Accomplishment 1: Derived an optimal noncausal resolvent-based estimator (Objectives 1 & 2)

The noncausal estimator is obtained by minimizing the cost function

$$J = \int_{-\infty}^{\infty} E\{\boldsymbol{e}^*(t)\boldsymbol{e}(t)\}\,dt = \frac{1}{2\pi}\int_{-\infty}^{\infty} E\{\hat{\boldsymbol{e}}^*(\omega)\hat{\boldsymbol{e}}(\omega)\}\,d\omega, \tag{7}$$

where $\boldsymbol{e}$ is the error between the true and estimated state [14]. To estimate the flow from limited measurements, we seek a transfer function between the flow state $\tilde{\boldsymbol{q}}(t)$ and the measurement $\boldsymbol{y}(t)$ of the form

$$\tilde{\boldsymbol{q}}(t) = \int_{-\infty}^{\infty} \boldsymbol{T}_q(t-\tau)\boldsymbol{y}(\tau)\,d\tau, \tag{8}$$

which can equivalently be expressed in the frequency domain as

$$\hat{\tilde{\boldsymbol{q}}}(\omega) = \hat{\boldsymbol{T}}_q(\omega)\hat{\boldsymbol{y}}(\omega). \tag{9}$$

Here, $\boldsymbol{y} = \boldsymbol{C}\boldsymbol{q}$ is the measurement that is extracted from the flow state via the measurement operator $\boldsymbol{C}$. Minimizing the cost function leads to the optimal frequency domain transfer function (or kernel)

$$\hat{\boldsymbol{T}}_q = \boldsymbol{R}_q\hat{\boldsymbol{F}}\boldsymbol{R}_y^*\left(\boldsymbol{R}_y\hat{\boldsymbol{F}}\boldsymbol{R}_y^* + \hat{\boldsymbol{N}}\right)^{-1}. \tag{10}$$

Here, $\boldsymbol{R}_q$ is the standard resolvent operator defined in Eq. (4), $\boldsymbol{R}_y = \boldsymbol{C}\boldsymbol{R}_q$ is a modified resolvent operator linking the forcing and the measurement, and $\hat{\boldsymbol{F}}$ and $\hat{\boldsymbol{N}}$ are the cross-spectral density tensors for the nonlinear fluctuation terms in the Navier-Stokes equations and the measurement noise, respectively. This expression optimally accounts for the space-time statistics of the forcing and is clearly formulated in terms of resolvent operators. Time-domain estimates of the flow state are obtained by taking the inverse Fourier transform of Eq. (10) and using the result in Eq. (8).

As a first test of the noncausal resolvent-based estimator, we applied it to channel flows at several Reynolds numbers [15]. In all cases, measurements of shear stress and/or pressure at the wall were used to estimate the flow state on the interior of the channel. Figure 1 shows comparisons between DNS data and instantaneous estimates of streamwise velocity at several wall-normal planes. The estimator faithfully

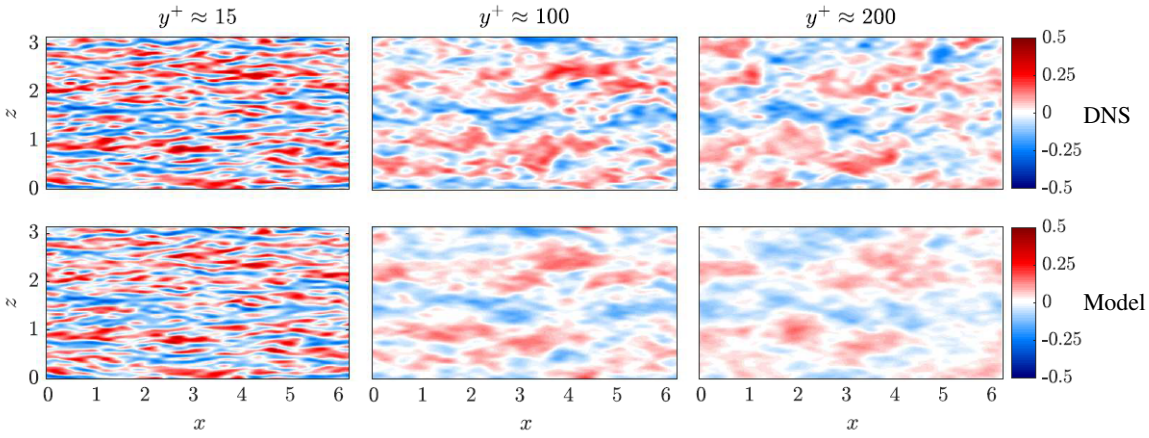

Figure 1: DNS (top row) and model estimates (bottom row) of instantaneous streamwise velocity at several wall-normal planes in an $Re_\tau = 550$ channel flow based on measurements of wall shear stress [15].

captures the most energetic organized motions within the flow using only measurements at the wall. We also showed that leveraging the ability of the resolvent-based estimator to account for colored forcing statistics leads to large improvements compared to a standard implementation of a Kalman smoother, the noncausal version of the well-known Kalman filter.

Accomplishment 2: Derived an optimal causal resolvent-based estimator (Objectives 1 & 2)

The estimator derived in Accomplishment 1 is noncausal, i.e., the estimate potentially depends on future measurements. This can be seen in the doubly infinite integration limits in Eq. (8). Such an estimator is useful for reconstructing a flow state from previously measured data but cannot, in general, be used for real-time applications such as flow control.

One option to obtain a causal estimate would be to use the noncausal estimation kernel and simply truncate the integral such that only past and present measurements are considered. However, this destroys the optimality of this noncausal estimation kernel since it was derived assuming that future measurements would be available. Instead, we wish to derive a new estimation kernel that is optimal under the constraint of causality [16, 17, 18]. This can be accomplished by augmenting our cost function to be minimized with additional terms meant to enforce causality,

$$J = \int_{-\infty}^{\infty} \left(E\{\boldsymbol{e}^*(t)\boldsymbol{e}(t)\} + \boldsymbol{\Lambda}_-(t)\boldsymbol{T}_{q,c}(t) + \boldsymbol{\Lambda}_-^*(t)\boldsymbol{T}_{q,c}^*(t) \right) dt. \tag{11}$$

Compared to the cost function in Eq. (7), we have added Lagrange multipliers that are forced to be zero for times greater than zero multiplied by the estimation kernel. The extra c subscript here is meant to indicate that these kernels will be causal.

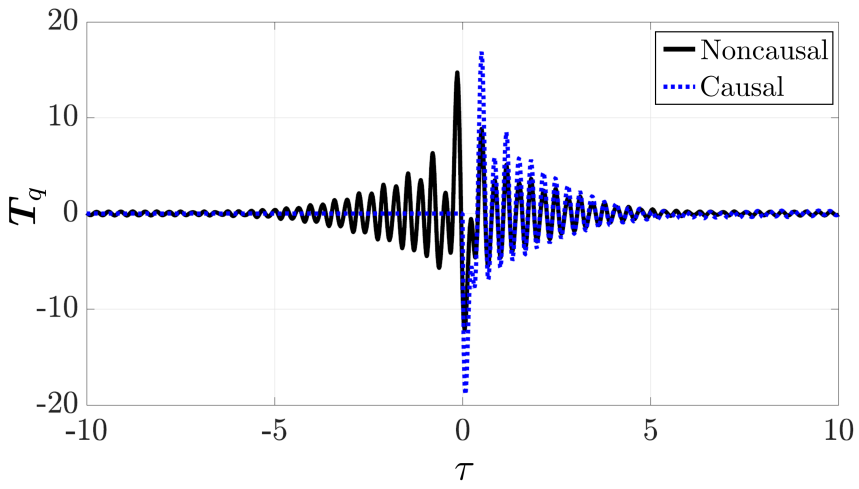

Figure 2: Comparison between the noncausal (black) and optimal causal (blue) control kernels for one of the sensor-target pairs for the laminar airfoil estimation problem shown in Figure 11.

Minimizing this new cost function leads to the condition

$$\hat{\boldsymbol{T}}_{q,c}\hat{\boldsymbol{G}}_l + \hat{\boldsymbol{\Lambda}}_- = \hat{\boldsymbol{G}}_r, \tag{12}$$

where

$$\hat{\boldsymbol{G}}_l = \boldsymbol{R}_y \hat{\boldsymbol{F}} \boldsymbol{R}_y^* + \hat{\boldsymbol{N}}, \tag{13a}$$

$$\hat{\boldsymbol{G}}_r = \boldsymbol{R}_q \hat{\boldsymbol{F}} \boldsymbol{R}_y^*, \tag{13b}$$

under the constraint that $\boldsymbol{\Lambda}_-(t>0)=0$ and $\boldsymbol{G}_l(t<0)=0$. This is a matrix Wiener-Hopf problem that must be solved to obtain the optimal causal kernel $\hat{\boldsymbol{T}}_{q,c}$. Analogous to Eq. (10), the matrix coefficients $\hat{\boldsymbol{G}}_l$ and $\hat{\boldsymbol{G}}_r$ in the Wiener-Hopf problem are defined in terms of resolvent operators and space-time statistics of the nonlinear forcing and noise. An efficient method to solve these matrix Wiener-Hopf problems is introduced in Ref. [17]. Figure 2 shows an example (from the laminar airfoil problem described in Accomplishment 10) of the outcome of this procedure. The noncausal control kernel (black line) is nonzero for $\tau<0$, indicating that it requires future measurements to be evaluated. The causal kernel (blue dashed line) is optimized under the constraint that it must be zero for $\tau<0$. That the two kernels are not equal for $\tau>0$ exemplifies that simply truncating the noncausal kernel and using its causal part is not optimal.

The causal resolvent-based estimator can be shown to be equivalent to a Kalman filter when equivalent assumptions are made about the forcing and noise statics, i.e., that they are temporally uncorrelated. However, unlike the Kalman filter, the resolvent-based estimator can accommodate colored space-time forcing

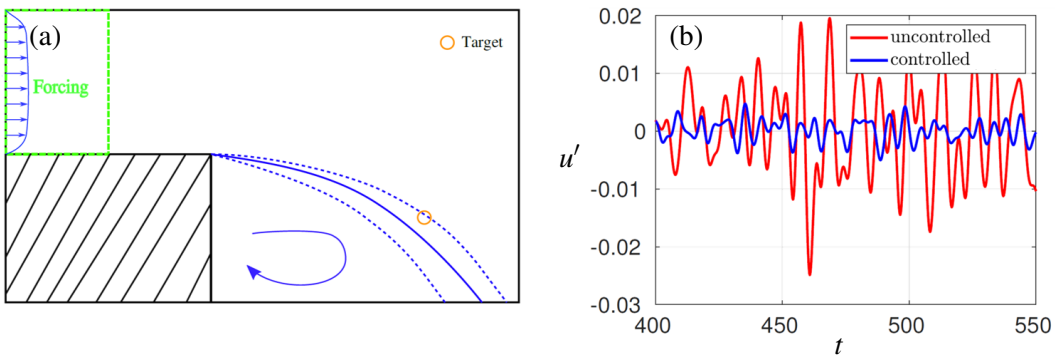

Figure 3: Control of velocity perturbations in the flow over a backward-facing step using the resolvent-based formulation: (a) problem setup; (b) streamwise velocity perturbation at the target location shown at left before and after control is applied. The velocity perturbations at this location are reduced by an order of magnitude.

and noise statistics [17]. Within the context of Kalman filtering, the forcing is simply process noise, but for us, it represents the nonlinear terms in Navier-Stokes. Numerous studies have shown that accounting for these space-time colored statistics is essential for obtaining accurate models [19, 20, 21], and this is true for estimation as well. Additionally, the causal resolvent-based estimator can be computed far more efficiently than a Kalman filter for fluid mechanics problems (see Accomplishment 4).

Accomplishment 3: Derived an optimal resolvent-based controller (Bonus!)

While control is beyond the original scope of the project objectives, deriving an optimal resolvent-based controller was straightforward once we learned how to formulate the causal estimator [17]. In short, we define a cost function analogous to Eq. (11),

$$J = \int_{-\infty}^{\infty} \left(E\{\boldsymbol{z}^*(t)\boldsymbol{z}(t) + \boldsymbol{a}^*(t)\boldsymbol{P}\boldsymbol{a}(t)\} + \boldsymbol{\Lambda}_-(t)\boldsymbol{\Gamma}(t) + \boldsymbol{\Lambda}_-^*(t)\boldsymbol{\Gamma}^*(t) \right) dt. \tag{14}$$

Here, $\boldsymbol{\Gamma}(t)$ is the control kernels, the target flow perturbations $\boldsymbol{z}(t)$ and control effort $\boldsymbol{a}(t)$ are penalized, and causality is again enforced using Lagrange multipliers. Minimizing the cost function leads to another matrix Wiener-Hopf problem, which is solved using the same approach as for the controller. When white-noise forcing is assumed, the resolvent-based controller reproduces an LQG control law, but our approach can also accommodate arbitrary forcing color and, again, can be computed more efficiently.

As an initial test, we applied the full resolvent-based estimation and control framework to the flow over a backward-facing step, as shown in Figure 3. Controlling the flow using sensors and actuators above the step leads to an order-of-magnitude reduction in the amplitude of velocity fluctuations in the shear layer that forms between the free stream and the recirculation bubble behind the step [16].

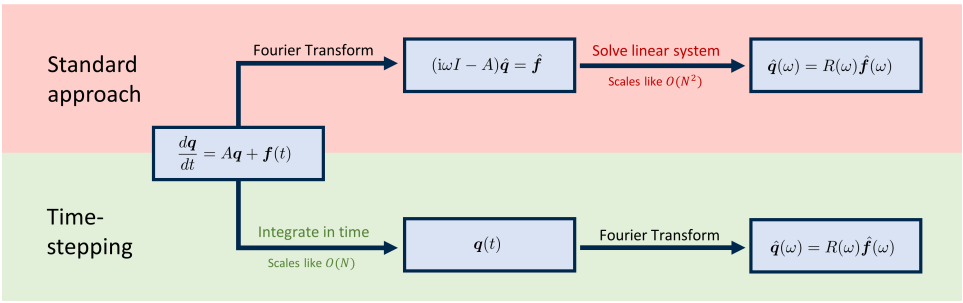

Figure 4: Comparison between the standard frequency domain approach and the time-stepping approach for obtaining the action of the resolvent operator on a vector. The time-stepping approach can be used to efficiently compute the estimation and control kernels (Accomplishment 4) or the singular value decomposition of the resolvent operator (Accomplishment 5) or Harmonic resolvent operator (Accomplishment 6).

In summary, we have developed an optimal resolvent-based estimation and control framework with improved accuracy (by accounting for forcing color) and reduced cost (using an efficient time-marching method described next) compared to standard control methods like the Kalman filter and LQG control.

Accomplishment 4: Formulated a time-marching method to efficiently compute the kernels (Objective 1)

Obtaining a Kalman filter or LQG controller requires the solution of an algebraic Riccati equation, the computational cost of which scales cubicly with the dimension of the problem. This poor scaling makes these methods intractable for large systems typical of turbulent flows, except in cases where symmetries of the problem can be used to effectively reduce the size of the state. The usual approach for circumventing this high cost is to first use a model reduction method to obtain a smaller system to which to apply a Kalman filter and/or LQG controller, but doing so forfeits their optimality when applied to the full-order system.

As we have seen, the resolvent-based estimation and control kernels are defined in terms of resolvent operators. Computing these resolvent operators is also a computationally intensive task, and modern methods for doing so scale (roughly) quadratically with problem size. This is already an improvement compared to the Kalman filter but is still problematic for the large systems of interest in this project.

To make our resolvent-based estimator applicable to large systems, we formulated a time-marching method to efficiently compute the estimation kernels [17]. The central idea, depicted in Figure 4, is to return to the linearized time-domain equations from which the resolvent operators are defined, and integrate the direct and adjoint linear equations forward and backward in time to obtain the action of the resolvent operator or its adjoint on a specified forcing vector. By strategically selecting the forcing vectors, we can efficiently obtain the desired estimation and control kernels. Assuming the availability of a well-implemented

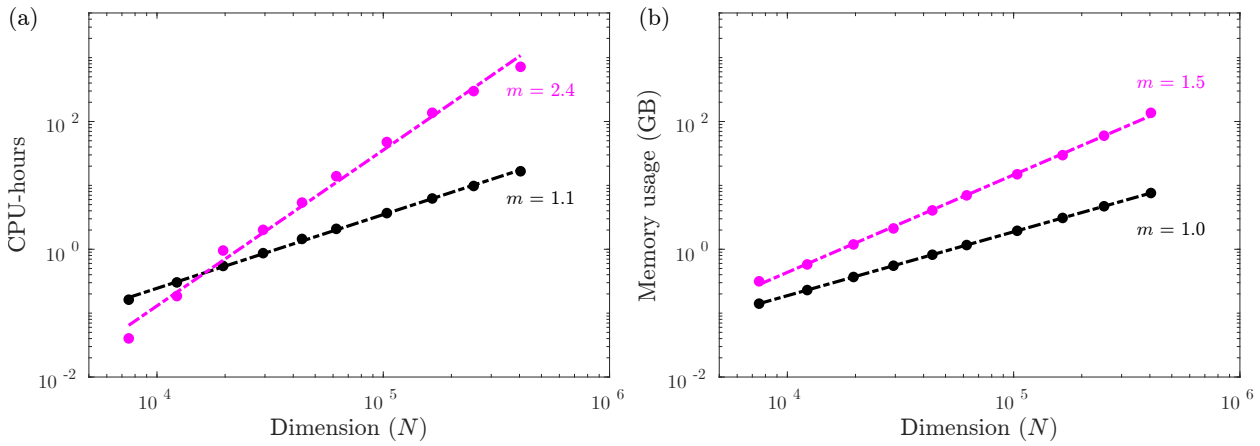

Figure 5: Computational cost of resolvent analysis as a function of the state dimension N for a three-dimensional jet: (a) CPU-hours and (b) memory usage for the RSVD-LU (pink) and RSVD-Δt (black) algorithms.

simulation code for a given problem of interest, this approach achieves linear cost scaling with problem size, leading to enormous cost savings for large problems such as those of interest for this project.

Accomplishment 5: Extended the time-stepping method to efficiently compute resolvent modes for large systems (Bonus!)

The same time-stepping method, paired with randomized singular value decomposition (RSVD), can be used to efficiently compute resolvent modes, i.e., singular vectors of the resolvent operator, rather than estimation and control kernels [22, 23]. Despite its popularity and demonstrated value, the application of resolvent analysis beyond canonical, academic flows to turbulent flows of practical engineering interest has been seriously hindered by the computational cost of computing resolvent modes [24]. In particular, the cost of computing resolvent modes increases rapidly with the state dimension N (the number of grid points times the number of physical variables). For three-dimensional turbulent flows of practical interest, $N \sim O(10^6 - 10^8)$ and the cost scales like N^2, making computation of resolvent modes intractable (except in some idealized cases). Modern algorithms such as RSVD reduce the absolute cost (typically by about 50% compared to previous Arnoldi-based methods) but do not overcome the fundamental scaling problem, which originates from an LU-decomposition used to apply the action of the resolvent operator to candidate forcing vectors with algorithms such as RSVD.

We have eliminated this bottleneck by combining RSVD with a variant of the efficient time-stepping method described in Accomplishment 4 [22, 23]. Critically, we exploit the direct and adjoint time-domain equations underlying the resolvent system to avoid the problematic LU decomposition bottleneck. Paired with a strategy to minimize the duration of the time stepping by removing unwanted transients and a stream-

ing approach to minimize memory consumption, our algorithm (which we call RSVD-Δt) drastically reduces CPU and memory costs for large systems.

Figure 1 shows a proof-of-concept demonstration of the new algorithm. Figure 1(a) compares the CPU cost of the previous state-of-the-art algorithm (RSVD with LU decomposition) and the new time-stepping algorithm as a function of the state dimension for a three-dimensional jet. The improved scaling is evident and leads to much lower costs. Figure 1(b) shows the peak memory requirement for each algorithm, which restricts the largest problem that can be solved on a given machine. Again, the new algorithm shows superior scaling and lower memory requirements. The largest values considered in this test case are still relatively modest (to make the problem tractable for the LU-based algorithm). If we extrapolate these results to a problem with 40 million degrees of freedom (a typical value for a three-dimensional grid), the new algorithm provides a 30,000x speedup and a 200x memory reduction.

Accomplishment 6: Extended the time-stepping method to efficiently compute harmonic resolvent modes (Bonus!)

Our RSVD-Δt algorithm for efficiently computing resolvent modes of large systems can be further extended to reduce the cost of computing harmonic resolvent modes [25]. Harmonic resolvent analysis is an extension of resolvent analysis that investigates flows characterized by a time-varying, periodic base flow, such as vortex shedding behind a cylinder [26]. The frequency content of the base flow determines the number of triadic interactions between the forcing, base flow, and response. The singular value decomposition (SVD) of the harmonic resolvent operator is employed to determine the optimal forcing and identify the most amplified response. Standard resolvent analysis is recovered as a special case when the base flow becomes time-independent. However, the computational requirements for computing the leading harmonic-resolvent modes are even more demanding than the resolvent case due to the higher dimensionality of the linearized operator caused by the aforementioned frequency coupling. This increase in dimensionality is a consequence of working in the frequency domain and can be avoided by applying our time-domain RSVD-Δt algorithm to the time-periodic system.

We tested our algorithm using the flow around a NACA0012 airfoil at an angle of attack of $\alpha = 20^\circ$ and Reynolds number of $Re = 200$, similar to the case considered by Padovan et al. [26]. This flow exhibits strong vortex shedding at Strouhal number $St_f = 0.114$, and the harmonic resolvent operator is constructed by linearizing about this periodic solution. The CPU and memory costs are reported in Figure 6 as a function of the number of coupled frequencies, N_ω, retained in the harmonic resolvent analysis. Whereas the costs of using the standard RSVD-LU algorithm increase quickly with N_ω, they are independent of N_ω for the

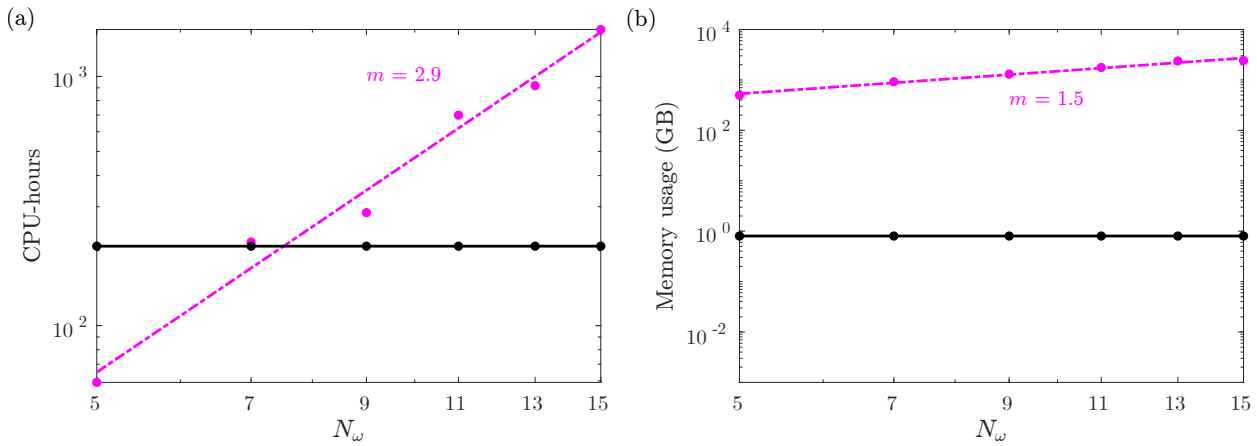

Figure 6: Computational cost of harmonic resolvent analysis as a function of the number of retained frequencies N_ω for the airfoil test case: (a) CPU-hours, and (b) memory usage for RSVD-LU (pink) and RSVD-Δt (black).

RSVD-Δt algorithm. The leading input and output modes are shown in Figure 7. The intrinsic low-rank structure of this harmonic system results in a striking resemblance between the vortical structures observed in the first output mode and the vorticity patterns seen in the nonlinear simulation driven by a sinusoidal perturbation [26]. The forcing modes predominantly occur near the trailing edge, indicating the sensitive region of the airfoil for control and design purposes.

Accomplishment 7: Formulated a data-driven approach for computing the kernels (Objectives 1)

In the approach described so far, the coefficients for the Wiener-Hopf problem, which contain resolvent operators and forcing and noise statistics, are obtained via direct and adjoint time marching of the linearized equations. As an alternative, we have formulated an approach to compute the Wiener-Hopf coefficients directly from data [17, 27].

Specifically, the coefficients in the estimation Wiener-Hopf problem can be computed from correlations between the sensor measurements and the target using data from the uncontrolled system. Specifically, Eq. (13) can be rewritten as

$$\hat{\boldsymbol{G}}_l = \boldsymbol{S}_{yy}, \tag{15a}$$

$$\hat{\boldsymbol{G}}_r = \boldsymbol{S}_{qy}, \tag{15b}$$

where $\boldsymbol{S}_{yy}$ and $\boldsymbol{S}_{qy}$ are the cross-spectral density of the measurement $\boldsymbol{y}$ and between the state $\boldsymbol{q}$ and measurement $\boldsymbol{y}$, respectively. That is, we simply need to take measurements at the locations where we want to place sensors and at the locations where we aim to control the flow, and from these measurements, we can compute the Wiener-Hopf coefficients and, therefore, the optimal estimation kernels. The control Wiener-Hopf

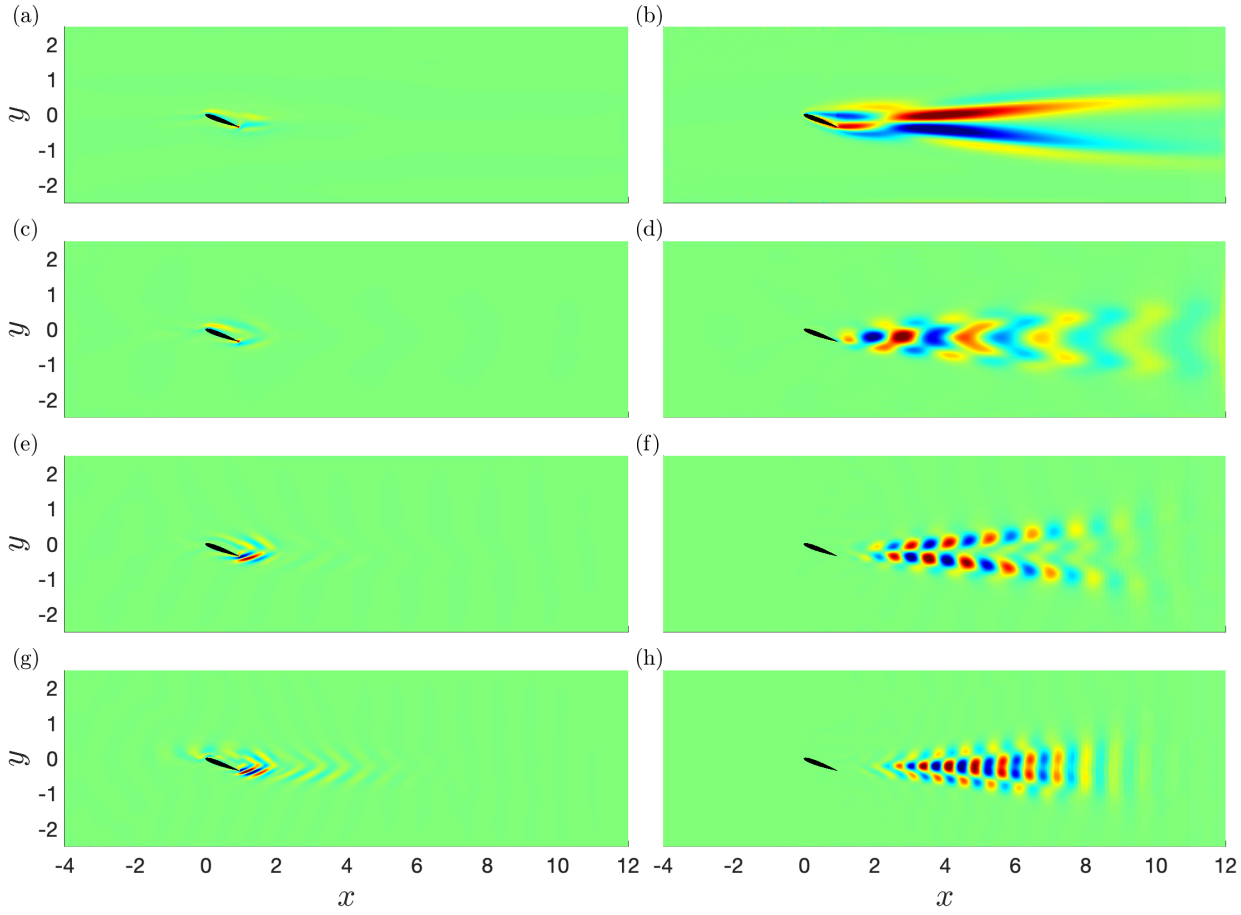

Figure 7: Harmonic resolvent results for the airfoil test case: real part of the vorticity field for the optimal input (a, c, e, g) and output (b, d, f, h) modes at (a, b) $St = 0$, (c, d) $St = St_f$, (e, f) $St = 2St_f$, and (g, h) $St = 3St_f$, where St_f is the fundamental frequency. Color bar ranges are adjusted for visualization.

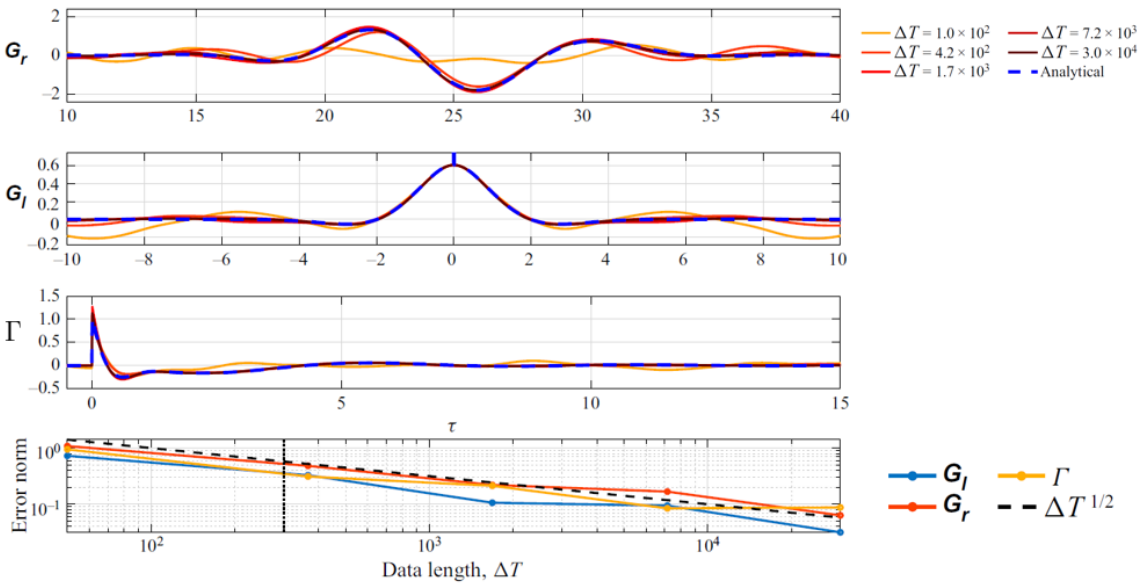

Figure 8: Results for the data-driven formulation for the backward-facing step problem [17]. Convergence to the operator-based kernels is achieved as the amount of data increases.

problem includes a second set of coefficients, and these can be computed from impulse response measurements. That is, we simply ping the flow with our actuators and measure the response, and the optimal control kernels can be constructed using this data along with the sensor-target correlations.

This data-driven approach is significant because it broadens the applicability of the resolvent-based estimation and control methods. In simulations, we can compute the control kernels without the need for a linearized code or adjoint capabilities. Even more significant, we can compute the control kernels entirely from experimental measurements by measuring correlations in the uncontrolled flow and impulse responses.

We initially tested this data-driven approach using the backward-facing-step problem. Figure 8 shows two of the Wiener-Hopf coefficients and the final control kernel computed using different amounts of data. With enough data, the data-driven results match the analytical ones. The main downside of the data-driven approach is that the convergence is slow: like most sample averages, it converges like the square root of the number of samples. Thus, it is advantageous to use the time-marching method when possible, but the data-driven approach extends the applicability of the theory to experiments and simulations without the needed capabilities, making it easier for others to adopt our methods.

Accomplishment 8: Enabled efficient exploration of sensor and actuator placement (Objective 2)

Selecting appropriate locations and types of sensors (and actuators) is important for achieving effective estimation (and control). We have developed an approach that allows us to efficiently explore sensor and ac-

tuator placements within our resolvent-based framework. This requires the ability to (i) efficiently construct many control and estimation kernels for different configurations and (ii) efficiently evaluate the performance of each configuration.

First, while building the resolvent-based estimation and control kernels is already inexpensive, we have devised a way to bring the extra cost of building kernels for many different sensor and actuator configurations to essentially zero. The key observation is that virtually all of the cost of our procedure is incurred in the time-integration step. Once this step is complete, we can build the needed resolvent operators for other choices of sensor and actuator placements at negligible additional cost by simply multiplying the time-marching results by different input and output matrices.

Second, we can evaluate the performance of the controller for a given choice of sensors and actuators without actually applying the control. This is enabled by an expression for the cross-spectral density (CSD) of the target for the controlled flow that depends only on the estimation and control kernels rather than simulation data [17]. The diagonal of the CSD is the power spectral density (PSD), and the trace is the quantity the controller is trying to minimize. This means that the performance of the controller can be estimated (under the assumption of a linear response to the controller) without ever actually applying the controller to the flow. Figure 9 contains an example of this capability. The back line shows the PSD of the target for the backward-facing-step problem. The red solid and dashed lines show the predicted PSD using two possible sensor/actuator setups, while the corresponding blue lines show the actual PSD of the controlled flow for these same two setups. While the predictions are not exact (due to the assumption of linearity), they successfully identify which of the setups is more effective at reducing the target.

Putting together these two capabilities – inexpensive variation of the kernels and a priori performance evaluation – we can perform inexpensive exploration of sensor and actuator placement, even for large systems.

Accomplishment 9: Implemented the resolvent-based tools within CharLES (Objective 3)

We have implemented the resolvent-based estimation and control tools into the high-fidelity CharLES simulation code. CharLES solves the spatially-filtered compressible Navier-Stokes equations on unstructured grids using a control-volume-based finite-volume method and includes advanced sub-grid and wall modeling capabilities to capture unresolved scales.

Implementing the resolvent-based estimation and control tools required four main tasks. First, we linearized the code via a non-intrusive, optimal sampling of the CharLES DNS/LES operator. Specifically, the

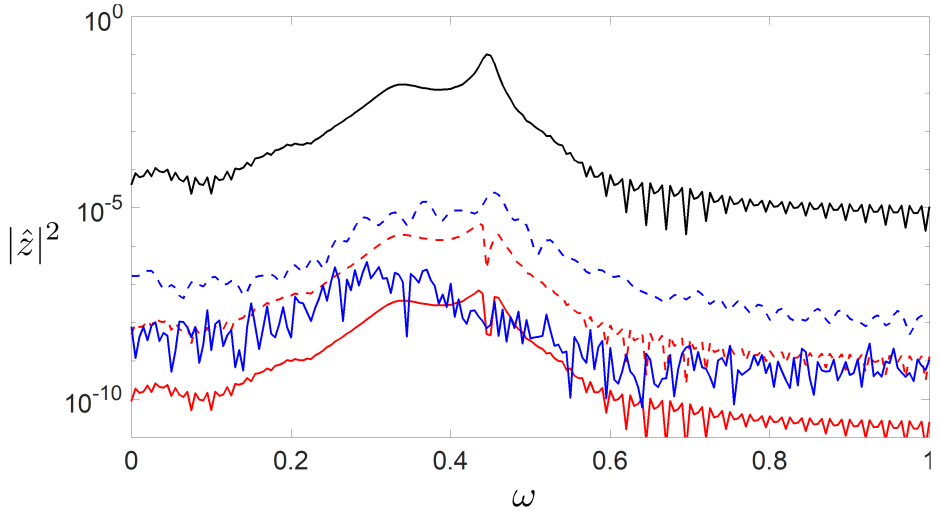

Figure 9: Example of a priori approximation of control efficacy for the backward-facing step: (black) uncontrolled flow; (red) predicted PSD of the controlled flow for two different control setups; (blue) actual PSD of the controlled flow for the same two setups.

linear operator A is constructed using Frechet derivatives,

$$A(:,j) \leftarrow \frac{\mathcal{F}(\bar{\boldsymbol{u}} + \varepsilon \hat{\boldsymbol{e}}_j) - \mathcal{F}(\bar{\boldsymbol{u}})}{\varepsilon}, \tag{16}$$

where $\mathcal{F}$ is the DNS/LES approximation of the Navier-Stokes operator, ε is a small number, $\hat{\boldsymbol{e}}_j$ is the j-th unit vector of size N, and N is the total number of degrees of freedom of the problem (i.e., the number grid points times the number of flow variables). Naively building the operator in this way would require $\mathcal{F}$ to be evaluated independently for all N degrees of freedom of the problem, which is prohibitively expensive for large problems. Instead, we leverage the finite size of the numerical stencil (e.g., only nearby control volumes contribute to the flux) to perturb many degrees of freedom simultaneously [28]. This makes the number of evaluations of $\mathcal{F}$ independent of the problem size, allowing the linear operator to be obtained for large systems at low cost.

Second, we implemented the time-marching algorithm that enables the estimation and control kernels to be generated efficiently. This process involves integrating the direct and adjoint linear equations, often in succession, i.e., the time-varying solution of one run serves as the forcing term for the next run. For large problems, storing these time-varying solutions in memory becomes problematic and reading and writing to disk is slow. To minimize this memory overhead, our implementation includes an optimized check-pointing algorithm, which eliminates the need to store time-series data, with a penalty of roughly 50% higher CPU cost (which is of no concern given the low cost of our method).

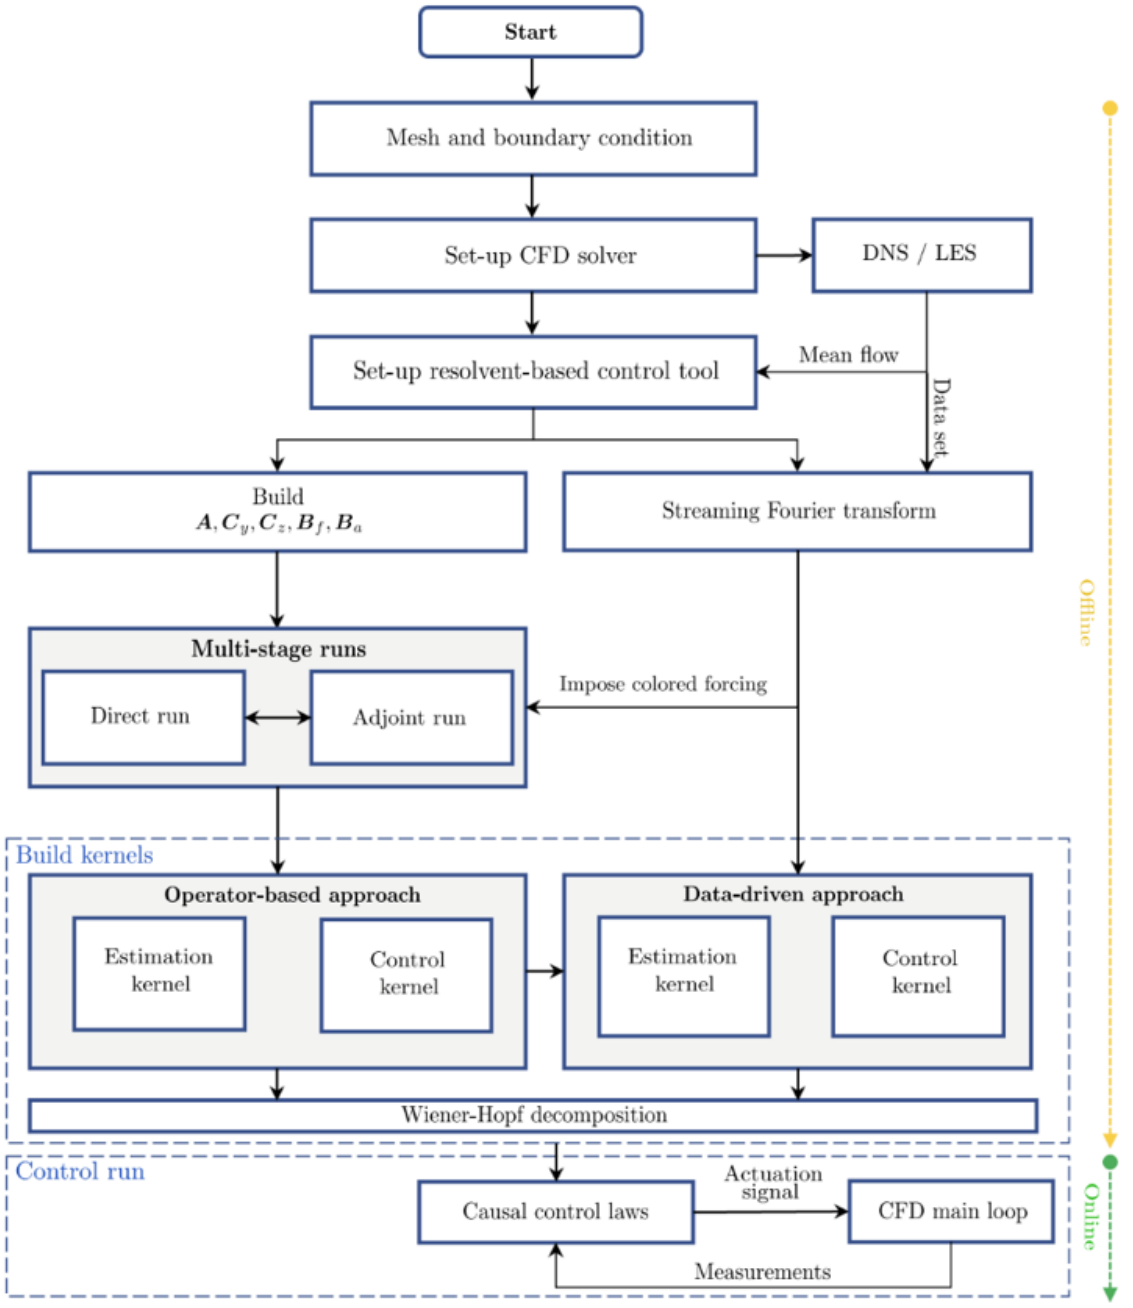

Figure 10: Implementation of the resolvent-based tools within CharLES.

Third, we added the option to use data from preliminary simulations to compute the data-driven estimation and control kernels. Fourth, we implemented the Wiener-Hopf factorization within the CharLES code. Either the operator-based or data-driven noncausal kernels are given as input and are factorized to yield the corresponding causal kernels.

In aggregate, these three new capabilities allow us to complete the entire process of creating and employ-

ing the resolvent-based estimation and control kernels without leaving the CharLES environment. Critically, the entire code is efficiently parallelized using a combination of the native CharLES architecture and PETSc, a suite of data structures and routines for the scalable solution of scientific applications developed and maintained by Argonne National Laboratory. This allows our estimation and control tools to be applied to very large problems and, in principle, to any flow that can be simulated in CharLES. The overall implementation of the resolvent-based tools into CharLES is shown schematically in Figure 10.

Accomplishment 10: Performed estimation for laminar flow over an airfoil (Objective 3)

As a precursor to addressing turbulent aerodynamic flows, we applied our new tools to laminar flow over an airfoil, starting with the estimation problem [18]. Specifically, we considered a two-dimensional laminar airfoil flow at Reynolds number $Re = 5000$, Mach number $M = 0.3$, and angle of attack $AoA = 6.5^\circ$. We considered two different freestream conditions: clean and noisy, the latter of which was included to make the problem more challenging. The noisy environment is generated by randomly forcing the flow well upstream of the airfoil. This prevents the vortex shedding from falling onto its natural quasi-periodic limit cycle, leading to a chaotic flow that is more challenging to estimate and qualitatively similar to turbulence. In both cases, a small number of sensors are placed on the surface of the airfoil, and the goal is to estimate fluctuations downstream in the wake.

Figure 11 shows sample estimation results for the laminar airfoil [18]. While we have considered many different sensor and target configurations, here, as an example, we use the three shear-stress sensors shown in Figure 11(a) to estimate the streamwise velocity at the indicated downstream target positions. The causal resolvent-based estimator (blue line) outperforms a truncated noncausal estimator (red dashed line) in matching the DNS data (black line) at both targets, both when the airfoil is immersed in the clean (b-c) and noisy (d-e) freestream. For the clean inflow case, both methods accurately estimate the further downstream z_2 target, but the optimal causal estimator performs slightly better. For the closer z_1 target, the optimal causal estimator provides a large improvement over the truncated noncausal estimator. This is in line with our general observation that optimally enforcing causality is increasingly important when the sensors and targets are close together. The same observations hold for the noisy freestream case but with somewhat larger errors. Nevertheless, the agreement is quite good, considering that this is a chaotic flow due to the noisy freestream.

Figure 12 offers another visualization of the estimation results. Namely, we show a snapshot of the streamwise velocity fluctuation in the spatial plane at one instant in time. The estimated flow fields (bottom row) closely match the fields from the DNS (tow row) for both the clean (left) and noisy (right) freestream conditions. This exemplifies the ability of the optimal causal resolvent-based estimator to provide real-time

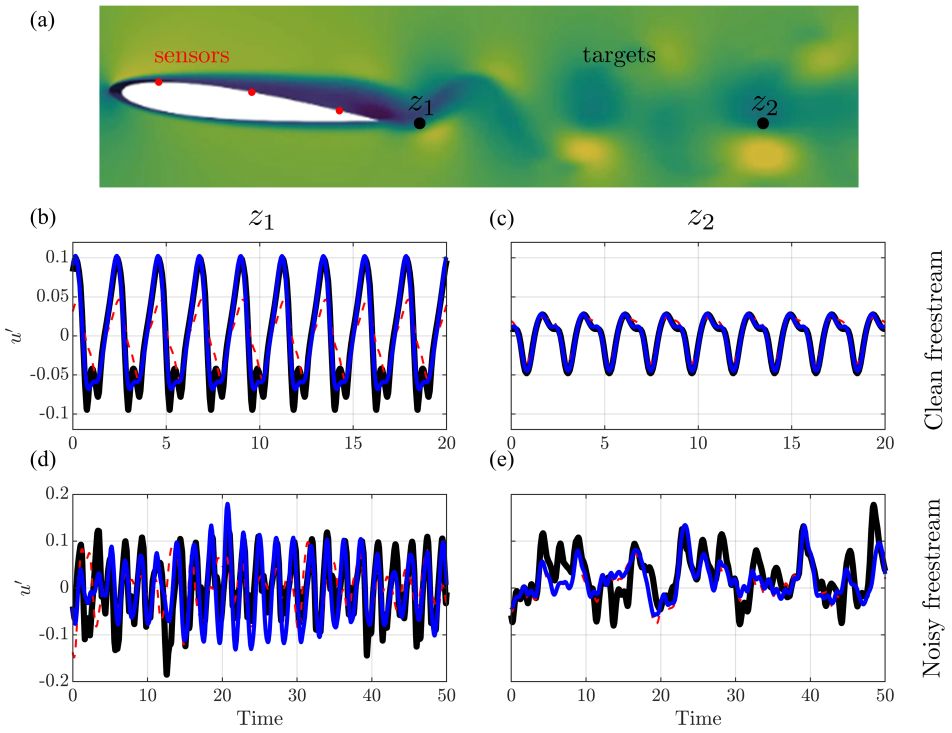

Figure 11: Estimation results for the laminar airfoil: (a) sensor and target placement; (b,c) estimation in noiseless environment; (d,e) estimation in noisy environment. Colors: (black) DNS data; (pink) noncausal method; (blue) optimal causal resolvent-based method.

information about the flow field surrounding an aerodynamic body using limited surface-mounded sensors.

Accomplishment 11: Performed control for laminar flow over an airfoil (Bonus!)

While control is beyond the original scope of the project, we have demonstrated the capabilities we developed in Accomplishment 3 by controlling the laminar airfoil flow with clean freestream conditions discussed in Accomplishment 10. As shown in Figure 13(a), here, we use one shear-stress sensor and one actuator that injects momentum into the flow, and the goal is to minimize the streamwise velocity fluctuation at the target. After a brief transient when the controller is turned on, the largest negative velocity fluctuations at the target are reduced by over 50%, as shown in Figure 13(b). We expect improved performance when multiple sensors and actuators are employed; we are currently performing these calculations and will report results in an upcoming publication.

Accomplishment 12: Performed estimation for turbulent flow over an airfoil (Objective 3)

Finally, we apply our resolvent-based controller to the turbulent flow over an airfoil [27]. This final task is the confluence and culmination of the theoretical development of the resolvent-based methods, their implemen-

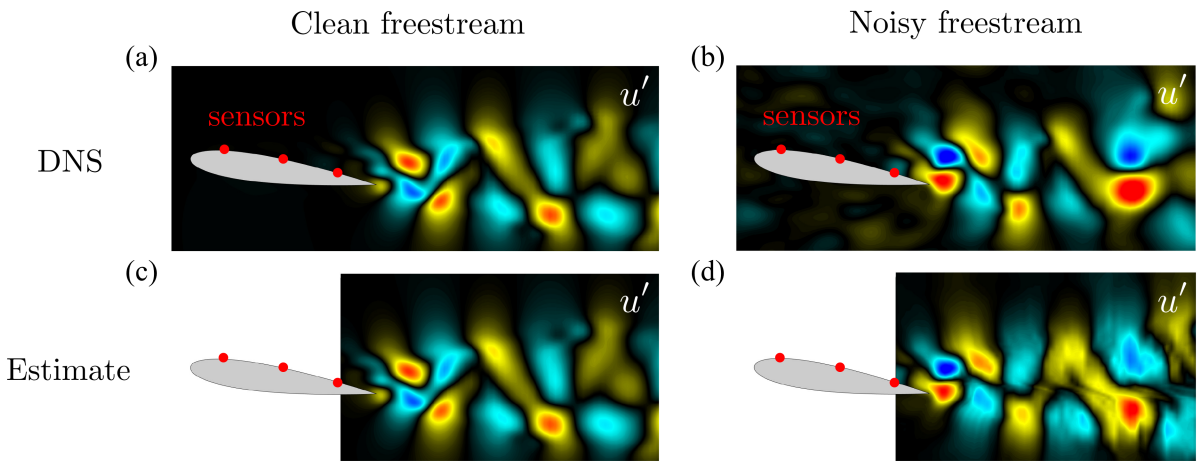

Figure 12: Snaphots of streamwise velocity for the laminar flow over an airfoil immersed in a (a,c) clean and (b,d) noisy freestream: (a,b) DNS, (c,d) resolvent-based estimates.

tation, and their application to aerodynamic flows. Following Yeh & Taira [13], we consider a spanwise-periodic NACA0012 airfoil at chord-based Reynolds number $Re = 23,000$, Mach number $M = 0.3$, and angle of attack $AoA = 6^\circ$. We first perform a large-eddy simulation (LES) of the flow using CharLES, and data from the simulation are used to construct the data-driven estimation kernels, in addition to providing sensor data and target data to evaluate the accuracy of the estimates. By construction, the kernels include the impact of the colored statistics of the nonlinear terms from the Navier-Stokes equations that act as a forcing on the linear dynamics. Moreover, using the data-driven kernels allows us to circumvent the global instability of the linear operator for this problem [13].

Figure 14 shows sample estimation results for the turbulent airfoil [27]. This time, we use the six shear-stress sensors shown in Figure 14(a) to estimate the streamwise velocity at the indicated downstream target positions. Since the turbulent flow is three-dimensional, we consider multiple options for handling the spanwise coordinate. In Figure 14(b-c), we exclusively estimate the spanwise-averaged fluctuations, while in Figure 14(d-e) we estimate the fluctuations in the mid-span plane. The latter task is more challenging since the three-dimensional turbulent motions lead to fluctuations that are observed at the targets in the mid-span plane but cannot be sensed in that same plane. We are currently exploring whether additional out-of-plane sensors can improve the estimation accuracy. Nevertheless, it is clear in all cases that the causal resolvent-based estimator outperforms the truncated noncausal estimator and provides reasonable accuracy.

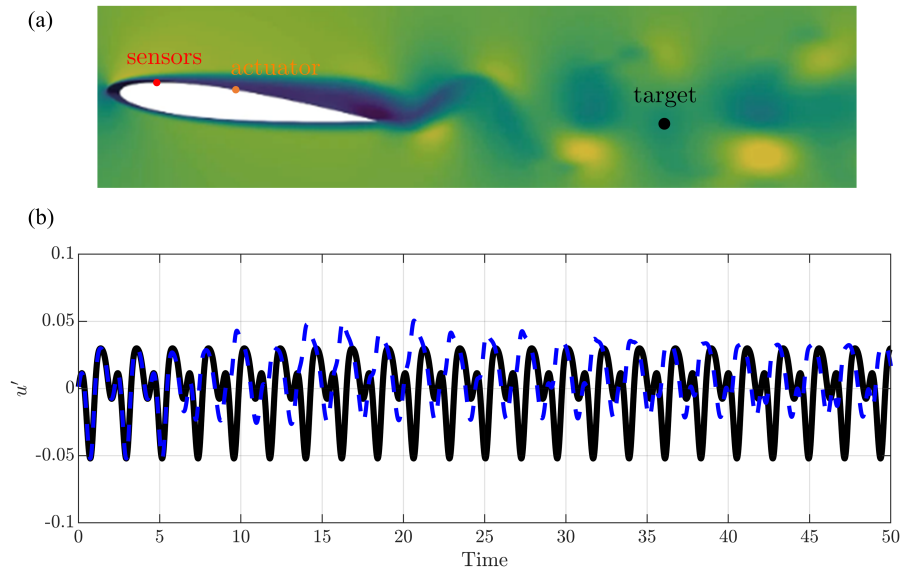

Figure 13: Example control results for the laminar airfoil. Colors: (black) uncontrolled; (blue) controlled.

3 Assessment of objectives and changes

The accomplishments described in Section 2 satisfy the three objectives defined in the project proposal and summarized in Section 1.3. Specifically, Objective 1 (develop an optimal resolvent-based estimation method) was achieved by the optimal noncausal and causal estimators derived in Accomplishments 1 and 2 and by their efficient and flexible time-stepping and data-driven implementations developed in Accomplishments 4 and 7. Objective 2 (determine the maximum accuracy of resolvent-based models for aerodynamic flows) is automatically satisfied by the formal optimality of the resolvent-based estimators in Accomplishments 1 and 2 and by the ability to efficiently explore sensor placements made possible by Accomplishment 8. Objective 3 (Assess the performance of resolvent-based estimation for aerodynamic flows) is satisfied by the application of our resolvent-based estimator to the laminar and turbulent flows over airfoils reported in Accomplishments 10 and 12 and made possible by the embedding of those tools within the CharLES framework described in Accomplishment 9.

The final execution of the project involved some minor departures from the original work plan. First, the originally proposed approach to Objective 2 was to use the theoretical relationship between SPOD and resolvent analysis to assess the performance/optimality of the resolvent-based estimation method [7]. However, this aim was made obsolete by the provable optimality of our estimation framework and the a priori performance estimates, i.e., the combination of Accomplishments 1, 2, and 8 satisfies Objective 2 without the need to explicitly appeal to SPOD modes. Second, we introduced the laminar airfoil test case, which

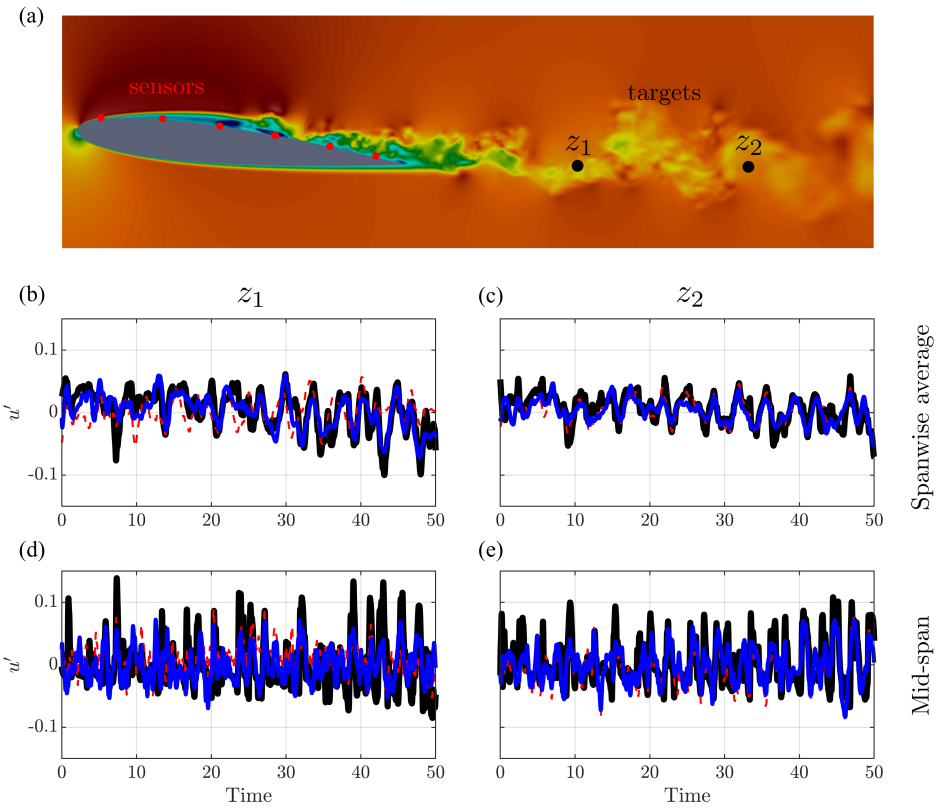

Figure 14: Estimation results for the turbulent airfoil: (a) sensor and target placement; (b,c) target z_1; (c,e) target z_2; using (b,c) spanwise-averaged data; (d,e) mid-span data. Colors: (black) LES data; (pink) noncausal method; (blue) optimal causal resolvent-based method.

was not described in the proposal, before taking on the proposed turbulent case. The laminar case provided a useful stepping stone that allows us to test our methods for aerodynamic flows in a less computationally demanding environment. We did ultimately return to the turbulent case in Accomplishment 12.

Additionally, we made two important contributions beyond the scope of the original project proposal. First, we derived an optimal resolvent-based controller and applied it to the laminar flow over an airfoil in Accomplishments 3 and 11, respectively. Second, we developed a novel algorithm to enable resolvent and harmonic resolvent of large-scale problems that could not be tackled with previous state-of-the-art algorithms, drastically expanding the reach of resolvent analysis and a slew of resolvent-based tools, as described in Accomplishments 5 and 6.

4 Dissemination

The results of this project have been disseminated through journal and conference publications, invited lectures, and conference presentations delivered by the students supported by the project.

4.1 Publications

1. Martini, E., Cavalieri, A. V., Jordan, P., Towne, A., & Lesshafft, L. (2020). Resolvent-based optimal estimation of transitional and turbulent flows. *Journal of Fluid Mechanics*, 900:A2.

2. Amaral, F. R., Cavalieri, A. V., Martini, E., Jordan, P., & Towne, A. (2021). Resolvent-based estimation of turbulent channel flow using wall measurements. *Journal of Fluid Mechanics*, 927:A17.

3. Farghadan, A., Towne, A., Martini, E. & Cavalieri, A. (2021). A randomized time-domain algorithm for efficiently computing resolvent modes. *AIAA paper #2021-2896*.

4. Martini, E., Jung, J., , Jordan, P., & Towne, A. (2022). Resolvent-based tools for optimal estimation and control via the Wiener-Hopf formalism. *Journal of Fluid Mechanics*, 937:A19.

5. Jung, J., Bhagwat, R., & Towne, A. (2023). Resolvent-based estimation of laminar flow around an airfoil. *AIAA paper #2023-0077*.

6. Jung, J. & Towne, A. (2024). Toward turbulent wake estimation using a resolvent-based approach. *AIAA paper #2024-0057*.

7. Farghadan, A., Martini, E. & Towne, A. (2024). Scalable resolvent analysis for three-dimensional flows. *Submitted to Journal of Computational Physics, arXiv:2309.04617*.

8. Farghadan, A., Jung, J., Bhagwat, R. & Towne, A. (2024). Efficient harmonic resolvent analysis via time-stepping. *Theoretical and Computational Fluid Dynamics (accepted), arXiv:2312.05766*.

9. Jung, J., Bhagwat, R., & Towne, A. (2024). Resolvent-based estimation and control of a laminar airfoil flow. *To be submitted to Journal of Fluid Mechanics*.

10. Jung, J. & Towne, A. (2024). Resolvent-based estimation of turbulent wakes. *To be submitted to Journal of Fluid Mechanics*.

4.2 Invited seminars

1. Towne, A. Resolvent-based optimal estimation and control of turbulent flows. *Aerospace Seminar, Ohio State University*, February 2021.

2. Towne, A. Advances in resolvent-based modeling of turbulent flows. *Aerospace Seminar, University of Illinois, Urbana-Champaign*, September 2021.

3. Towne, A. Resolvent-based estimation and control. *Workshop of the USP Aerospace Center*, November 2021.

4. Towne, A. Resolvent-based estimation and control of turbulent shear flows. *Isaac Newton Institute for Mathematical Sciences Workshop; Wall-bounded turbulence: beyond current boundaries*, March 2022.

5. Towne, A. Resolvent-based estimation and control of aerodynamic flows. *AIAA Aviation Forum; Flow Control Open Forum*, June 2022.

6. Towne, A. Aeroacoustic Noise Control. *Applied Physics Seminar, University of Michigan*, 21 September 2022.

7. Towne, A. Resolvent-based optimal estimation and control of turbulent flows. *Illinois Institute of Technology, Mechanical, Materials, and Aerospace Engineering Seminar*, 2 November 2022.

8. Towne, A. Resolvent-based estimation and control of turbulent flows. *Aerospace Chairs Distinguished Seminar Series, University of Michigan*, 19 January 2023.

9. Towne, A. Advances in non-modal hydrodynamic stability theory. *GALCIT Colloquium, California Institute of Technology*, 24 February 2023.

10. Towne, A. Three advances in non-modal stability theory for transitional and turbulent flows. *Department of Computational Mathematics, Science, and Engineering Colloquium, Michigan State University*, 3 April 2023.

11. Towne, A. Advances in resolvent-based modeling of turbulent flows. *Advanced Modeling & Simulation (AMS) Seminar, NASA Ames Research Center*, 5 October 2023.

4.3 Student conference presentations

1. Jung, J. Optimal resolvent-based estimation for flow control. *APS Division of Fluid Dynamics Meeting*, November 2020.

2. Farghadan, A. A randomized time-domain algorithm for efficiently computing resolvent modes, *AIAA Aviation Forum*, August 2021.

3. Farghadan, A. A new algorithm for computing global resolvent modes in a CPU and memory efficient manner, *APS Division of Fluid Dynamics Meeting*, November 2021.

4. Jung, J., Real-time estimation and control of flow over a NACA0012 airfoil using resolvent-based tools, *APS Division of Fluid Dynamics Meeting*, November 2022.

5. Farghadan, A. Global resolvent analysis of three-dimensional jets using randomized linear algebra and time stepping, *APS Division of Fluid Dynamics Meeting*, November 2022.

6. Jung, J., Resolvent-based estimation of laminar flow around an airfoil, *AIAA SciTech Forum*, January 2023.

7. Jung, J., Resolvent-based estimation for turbulent wakes, *APS Division of Fluid Dynamics Meeting*, November 2023.

8. Farghadan, A. Harnessing time-stepping for cost-effective harmonic resolvent analysis, *APS Division of Fluid Dynamics Meeting*, November 2023.

9. Jung, J., Toward turbulent wake estimation using a resolvent-based approach, *AIAA SciTech Forum*, January 2024.

10. Farghadan, A. An efficient algorithm for resolvent and harmonic resolvent analyses, *13th International Symposium on Turbulence and Shear Flow Phenomena*, June 2024 (scheduled).

5 Impacts

Finally, we describe the technical and educational impacts of the project and suggest some opportunities for further investigation.

5.1 Technical impact

This project has enabled improved estimation and control of aerodynamic flows. Namely, our resolvent-based framework offers two key improvements over classical control methods like Kalman filters and LQG controllers: it can account for the colored nonlinearity of the flow in a statistical sense, improving accuracy, and it can be applied to large systems at low cost, making it applicable to the large problems typical in fluid dynamics. Our implementation within a high-performance computing code makes these tools available for a wide range of problems, and a data-driven implementation expands their applicability to experimental scenarios. The method is flexible with regard to both the location and type of measurements (e.g., pressure, temperature, shear stress) and actuators (e.g., micro-jets or plasma), and its real-time deployment requires only the evaluation of a small number of integrals.

These capabilities will enable physics-based prediction and control for flow objectives of AFOSR interest. Improved flow estimation capabilities enable aerodynamic performance improvements that contribute to the Air Force's mission to maintain and expand its technological superiority [29]. Flow estimation can enable improved flight control by allowing an aircraft to detect its surrounding flow environment using sensors on its surface. This could enable the prediction of future unsteady aerodynamic loads and lead to improved aerodynamic performance, maneuverability, and safety. This capability could also benefit missile guidance systems and path planning for future autonomous unmanned aerial vehicles (UAVs). The combination of flow estimation and control offers powerful capabilities. Time-dependent predictions of the flow in the boundary layer around an aircraft can be used as input to active control efforts to manipulate the flow via suitable actuation to achieve objectives such as reducing drag, increasing lift, delaying transition, preventing separation, and so on. These capabilities will be especially crucial for efforts to make hypersonic flight increasingly routine and cost-effective [29]. More generally, flow estimation is a necessary component for achieving the Air Force's fundamental goal of developing physically based predictive models and control strategies for turbulent, aerodynamic flows [30].

Additionally, the improved algorithm for computing resolvent modes, a popular model for coherent structures in turbulent flows, extends the applicability of this important tool beyond canonical flows to complex flows of practical engineering interest. This will unlock the latent potential of resolvent analysis for flows of relevance to the Air Force. Indeed, given the popularity and power of resolvent analysis, our new algorithm may, on its own, prove to be similarly impactful as the resolvent-based estimation and control framework itself.

5.2 Training

Two Ph.D. students received training and mentorship through this project. The primary student funded by the project presented at six conferences (plus one more upcoming), wrote two conference papers, and authored or coauthored four journal articles (two nearing submission). He passed his dissertation proposal examination in October 2023 and will defend his Ph.D. thesis in July 2024. A second student who was partially supported by this project presented at five conferences (plus one more upcoming), wrote one conference paper, and authored three journal articles (one accepted, one under review, and one nearing submission). He passed his dissertation proposal examination in March 2024 and is on track to defend his Ph.D. thesis in August 2024. Both students received ongoing mentorship through a formal mentoring plan, weekly one-on-one meetings with me, quarterly individualized reviews, and regular group meetings, including a paper-reading campaign to strengthen students' understanding of the field. Both students also participated in training and mentoring junior students in the group to ensure effective knowledge transfer.

5.3 Future opportunities

The tools developed during this project can be applied and extended in numerous directions. First, our methods could be applied to additional flows of DoD interest. The PI currently has an ONR project in which the resolvent-based estimator and controller are being used to mitigate noise-generating coherent structures in turbulent jets. Given the successful estimation for the turbulent airfoil problem, the resolvent-based controller could also be applied and evaluated for turbulent aerodynamic flows. Due to the efficiency of the methods for computing the kernels and/or resolvent modes, we could also take on more geometrically complex, three-dimensional aerodynamic problems. For example, our algorithms could be used to compute, for the first time, resolvent modes for a full aircraft geometry. Of particular interest would be studying how the realistic flow around the aircraft impacts the resonance properties of, and control methods to suppress, the open cavity flow created by an open weapons bay.

There are also several opportunities for methodological extensions of our tools. For both the laminar and turbulent flows we considered, the ability to inexpensively explore sensor and actuator locations could be leveraged to obtain optimal configurations. To date, we have addressed this via parametric exploration; a more formal approach could be pursued by wrapping an optimization routine around the entire sensor and actuator exploration process. The resolvent-based estimation and control framework could be extended to incorporate structural measurements such as elastic strain within an aircraft's structure, enabling flow estimates without the need for flow measurements. Another interesting direction would be to extend the

resolvent-based estimation and control framework to non-stationary problems, starting with the periodic harmonic-relevant problem.

In summary, the capabilities developed during this project have the potential to aid in the Air Force's mission to maintain its technological superiority by enabling physics-based prediction and closed-loop control, leading to aerodynamic performance improvements.

Bibliography and References Cited

[1] J. Walker, H. Helin, and D. Chou. Unsteady surface pressure measurements on a pitching airfoil. Technical report, United States Air Force Academy, 1985.

[2] Y. Nagaoka, H. G. Alexander, W. Liu, and C.-M. Ho. Shear stress measurements on an airfoil surface using micro-machined sensors. *JSME Int. J. B – Fluids Therm. Eng.*, 40(2):265–272, 1997.

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

[4] M.R. Jovanović. From bypass transition to flow control and data-driven turbulence modeling: an input-output viewpoint. *Annu. Rev. Fluid Mech.*, 53:311–345, 2021.

[5] B. J. McKeon. The engine behind (wall) turbulence: perspectives on scale interactions. *J. Fluid Mech.*, 817, 2017.

[6] K. Taira, S. L. Brunton, S. Dawson, C. W. Rowley, B. J. Colonius, T.and McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley. Modal analysis of fluid flows: An overview. *AIAA J.*, 2017.

[7] A. Towne, O. T. Schmidt, and T. Colonius. Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.*, 847:821–867, 2018.

[8] A. Chavarin and M. Luhar. Resolvent analysis for turbulent channel flow with riblets. *AIAA J.*, 58(2):589–599, 2020.

[9] S. Beneddine, D. Sipp, A. Arnault, J. Dandois, and L. Lesshafft. Conditions for validity of mean flow stability analysis. *J. Fluid Mech.*, 798:485–504, 2016.

[10] A. Towne, A. Lozano-Durán, and X. Yang. Resolvent-based estimation of space–time flow statistics. *J. Fluid Mech.*, 883:A17, 2020.

[11] O. Kamal, M.T. Lakebrink, and T. Colonius. Global receptivity analysis: physically realizable input–output analysis. *J. Fluid Mech.*, 956:R5, 2023.

[12] M. Luhar, A. S. Sharma, and B. J. McKeon. Opposition control within the resolvent analysis framework. *J. Fluid Mech.*, 749:597–626, 2014.

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

[14] E. Martini, A. V. G. Cavalieri, P. Jordan, A. Towne, and L. Lesshafft. Resolvent-based optimal estimation of transitional and turbulent flows. *J. Fluid Mech.*, 900:A2, 2020.

[15] F. R. Amaral, A. V. G. Cavalieri, E. Martini, P. Jordan, and A. Towne. Resolvent-based estimation of turbulent channel flow using wall measurements. *J. Fluid Mech.*, 927:A17, 2021.

[16] J. Jung, E. Martini, A. Cavalieri, P. Jordan, L. Lesshafft, and A. Towne. Optimal resolvent-based estimation for flow control. *Bulletin of the American Physical Society*, 2020.

[17] E. Martini, J. Jung, A. Cavalieri, P. Jordan, and A. Towne. Resolvent-based optimal estimation of transitional and turbulent flows. *J. Fluid Mech.*, 937:A19, 2022.

[18] J. Jung, R. Bhagwat, and A. Towne. Resolvent-based estimation of laminar flow around an airfoil. *AIAA Paper #2023-0077*, 2023.

[19] A. Towne, G. A. Brès, and S. K. Lele. A statistical jet-noise model based on the resolvent framework. *AIAA Paper #2017-3406*, 2017.

[20] P. Morra, O. Semeraro, D.S. Henningson, and C. Cossu. On the relevance of reynolds stresses in resolvent analyses of turbulent wall-bounded flows. *J. Fluid Mech.*, 867:969–984, 2019.

[21] P. Morra, P.A.S. Nogueira, A.V.G. Cavalieri, and D.S. Henningson. The colour of forcing statistics in resolvent analyses of turbulent channel flows. *J. Fluid Mech.*, 907:A24, 2021.

[22] A. Farghadan, A. Towne, E. Martini, and A. Cavalieri. A randomized time-domain algorithm for efficiently computing resolvent modes. *AIAA Paper #2021-2896*, 2021.

[23] A. Farghadan, E. Martini, and A. Towne. Scalable resolvent analysis for three-dimensional flows. *arXiv:2309.04617*, 2023.

[24] J. H. M. Ribeiro, C. Yeh, and K. Taira. Randomized resolvent analysis. *Phys. Rev. Fluids*, 5(3):033902, 2020.

[25] A. Farghadan, J. Jung, R. Bhagwat, and A. Towne. Efficient harmonic resolvent analysis via time-stepping. *arXiv:2312.05766*, 2023.

[26] A. Padovan, S. E. Otto, and C. W. Rowley. Analysis of amplification mechanisms and cross-frequency interactions in nonlinear flows via the harmonic resolvent. *J. Fluid Mech.*, 900:A14, 2020.

[27] J. Jung and A. Towne. Toward turbulent wake estimation using a resolvent-based approach. *AIAA Paper #2024-0057*, 2024.

[28] E. J. Nielsen and W. L. Kleb. Efficient construction of discrete adjoint operators on unstructured grids using complex variables. *AIAA J.*, 44(4):827–836, 2006.

[29] Science and technology strategy: strenghtening USAF science and technology for 2030 and beyond. Technical report, United States Air Force, 2019.

[30] Air Force Office of Scientific Research Broad Agency Announcement. Technical report, United States Air Force, 2019.