



This is the accepted manuscript made available via CHORUS. The article has been published as:

Randomized resolvent analysis

Jean Helder Marques Ribeiro, Chi-An Yeh, and Kunihiro Taira

Phys. Rev. Fluids **5**, 033902 — Published 18 March 2020

DOI: [10.1103/PhysRevFluids.5.033902](https://doi.org/10.1103/PhysRevFluids.5.033902)

Randomized resolvent analysis

Jean Helder Marques Ribeiro,^{*} Chi-An Yeh,[†] and Kunihiko Taira[‡]

*Department of Mechanical and Aerospace Engineering,
University of California, Los Angeles, CA 90095, USA*

(Dated: February 28, 2020)

Performing global resolvent analysis for high-Reynolds-number turbulent flow calls for the handling of a large discrete operator. Even though such large operator is required in the analysis, most applications of resolvent analysis extracts only a few dominant resolvent response and forcing modes. Here, we consider the use of randomized numerical linear algebra to reduce the dimension of the resolvent operator for achieving computational speed up and memory saving compared to the standard resolvent analysis. To accomplish this goal, we utilize sketching of the linear operator with random test matrices with a Gaussian distribution and with insights from the base flow incorporated to perform singular value decomposition on a low-rank matrix holding dominant characteristics of the full resolvent operator. The strength of the randomized resolvent analysis is demonstrated on a turbulent separated flow over an airfoil. This randomized approach clears the path towards tackling resolvent analysis for higher-Reynolds number bi- and tri-global base flows.

Keywords: resolvent analysis, instabilities, randomized methods

I. INTRODUCTION

One of the central questions in turbulence and flow control is concerned with the evolution of perturbations. Gaining detailed understanding of the perturbation dynamics in turbulent flows is a daunting task, due to the complex nonlinear dynamics that takes place over a broad range of spatial and temporal scales. To modify the global flow characteristics with control, it is necessary that the nonlinear interactions involving the actuation input (perturbation) become appropriately large to alter the base flow. For flow control, we need not track all possible ways in which the actuation input can modify the flow, but instead we can focus on the dominant directions in which the perturbation can be amplified. This notion has led to the modal analysis based approaches [1–4], including the global stability analysis [5] and the resolvent analysis [6].

In the presence of sustained perturbations or forcing inputs, the linear system response can be described by the transfer function from control theory. This linear analysis is greatly simplified when the input to the system is sinusoidal and leads to the well-known Bode plots that reveal the gain and phase response of the system over a range of forcing frequencies. The transfer function that relates the system input to the output is called the resolvent and its analysis has been extended to fluid flows by Trefethen *et al.* [6]. The resolvent analysis is based on the pseudospectral analysis and has been used to study the transient energy growth [6] as well as the harmonic response of the system [7]. These initial studies of resolvent analysis were performed about stable laminar flows.

An extension to resolvent analysis to study turbulent flows was presented by McKeon and Sharma [8]. They considered the nonlinear advection term to be the self-sustained input within the natural feedback loop of the fluid flow. This viewpoint has enabled the use of time-averaged base flows to reveal the input-output dynamics of turbulent flows. Moreover, discounting or finite-time horizon based extension of the resolvent analysis has enabled resolvent analysis to study flows with unstable base states [9, 10]. Because resolvent analysis can determine the most amplified forcing and response directions, it serves as a powerful analytical tool to find effective active and passive flow control techniques [10, 11].

The resolvent analysis needs two key ingredients: (i) the base flow and (ii) the linearized Navier–Stokes operator. It is known that the accuracy of the base flow and the spatial discretization of the linear operators is critical for extracting response characteristics correctly [10]. The need for accurate discretization of the linearized Navier–Stokes operators calls for sufficient grid resolution and appropriate computational domain size. As such, the discrete resolvent operator becomes large with size $m \times m$. Here, m is essentially the number of variables times the size of the grid, which can easily be upward of $\mathcal{O}(10^6)$ for turbulent flows [12]. For resolvent analysis, the singular value decomposition (SVD) needs to be performed on the large resolvent operator with a taxing operation count of $\mathcal{O}(m^3)$. To enable resolvent analysis at high Reynolds numbers, we must find a relief to perform SVD of the resolvent operator.

^{*} jeanmarques@ucla.edu

[†] cayeh@seas.ucla.edu

[‡] ktaira@seas.ucla.edu

Although the resolvent analysis is performed on a very large matrix, only the leading forcing and response modes are generally sought. Based on the amount of necessary matrix data used to perform the analysis, the desired output is only a very small fraction of the input data size. For this reason, it would be natural to consider that all elements of the resolvent matrix are not necessary to determine the leading resolvent modes. Within the resolvent framework, the core of the computations lies with the SVD. In order to handle a large operator for SVD to find the leading modes, we can consider subsampling the matrix of interest and perform the SVD on the low-order representation of this large matrix. Randomized numerical linear algebra has recently emerged as an effective technique to reduce a large matrix to its low-order representation [13–15], for applications including big data compression and data transfer. The key idea is to pass a randomly generated low-rank test matrix through the large matrix to obtain the so-called sketch of the full matrix. This sketch is low-rank but holds key information about the full matrix and can be used to derive appropriate bases to represent the full matrix in a low-dimensional manner [13, 15–17]. Randomized techniques can be incorporated into SVD to achieve tremendous computational and memory savings [13, 16, 18].

In recent years, modal analyses are tackling flows over complex geometries and high-Reynolds number flows with increasingly large degrees of freedom [3, 4, 19–23]. To further aid this endeavor, randomized SVD have been incorporated into data-based modal analysis techniques, including the proper orthogonal decomposition [18] and dynamic mode decomposition [24]. For global operator-based analyses of high-Reynolds-number flows [19, 22, 23], the leading singular values and modes can be determined with significant reduction in computational costs with the aid of randomized techniques. The randomized technique presented in this study can greatly expand the applicability of the resolvent analysis to high-Reynolds number flows with multiple inhomogeneous spatial directions.

In fact, randomized SVD has been adopted in the one-dimensional resolvent analysis of turbulent channel flow by Moarref *et al.* [25]. Since the details on the use of the randomized techniques for large-scale resolvent analysis is not available, one of our objectives is to provide detailed guidance on the use of randomized resolvent analysis for multi-dimensional turbulent base flows. Moreover, the error bound of randomized SVD is generally not guaranteed and needs to be further characterized. We also note that the sketching strategy can be further improved with insights from the base flow. In the present study, we aim to address these issues on the use of randomized SVD in the context of resolvent analysis for fluid flows.

In what follows, we present the randomized resolvent analysis and demonstrate its use for turbulent flow over a canonical airfoil. In section II, we introduce the randomized resolvent analyses and discuss how randomized SVD can be implemented without the explicit generation of inverse matrices. We also propose an alternative route of recovering the response modes (left singular vectors) with significantly enhanced accuracy. In section III, the present randomized technique is applied to the resolvent analysis of turbulent flow over a NACA 0012 airfoil. We assess the level of error and the convergence behavior with respect to resolvent modes and gains obtained from the use of randomized SVD. To further enhance the accuracy of randomized resolvent analysis, we introduce a physics-informed sampling technique that leverages the insights from the base flow. At last, we provide concluding remarks in section IV.

II. APPROACH

A. Full resolvent analysis

Let us consider the flow state $\mathbf{q} \in \mathbb{R}^m$ as a sum of the time-invariant base state $\bar{\mathbf{q}}$ and the statistically stationary fluctuating component $\mathbf{q}'$. With this Reynolds decomposed flow variable and appropriate discretization, we can express the discrete Navier–Stokes equations as

$$\frac{\partial \mathbf{q}'}{\partial t} = \mathbf{L}_{\bar{\mathbf{q}}} \mathbf{q}' + \mathbf{f}', \quad (1)$$

where $\mathbf{L}_{\bar{\mathbf{q}}} \in \mathbb{R}^{m \times m}$ is the linearized Navier–Stokes operator about the base state $\bar{\mathbf{q}}$ and $\mathbf{f}'$ collects the nonlinear terms and the external forcing inputs. We gather the nonlinear terms as external forcing in the turbulent mean flow following the perspective of McKeon and Sharma [8], Farrell and Ioannou [26], and Schmid [27]. For the traditional resolvent analysis, $\bar{\mathbf{q}}$ is chosen to be the stable laminar equilibrium state such that $\mathbf{f}'$ can be considered as the forcing input to the system with the nonlinear term neglected [7]. More recently, turbulent mean flows has been used for $\bar{\mathbf{q}}$ with $\mathbf{f}'$ representing the nonlinear terms as sustained forcing input within the natural feedback system [8].

We can consider the Fourier transform $[\mathbf{q}'(\mathbf{x}, t), \mathbf{f}'(\mathbf{x}, t)] = \int_{-\infty}^{\infty} [\hat{\mathbf{q}}_{\omega}(\mathbf{x}), \hat{\mathbf{f}}_{\omega}(\mathbf{x})] e^{-i\omega t} d\omega$ and express the relationship between $\mathbf{q}'$ and $\mathbf{f}'$ in frequency space as

$$-i\omega \hat{\mathbf{q}}_{\omega} = \mathbf{L}_{\bar{\mathbf{q}}} \hat{\mathbf{q}}_{\omega} + \hat{\mathbf{f}}_{\omega}, \quad (2)$$

where ω is the frequency. Note that spatial Fourier transform can also be incorporated if directional homogeneity is present. For stable base flows, ω can be chosen to be real. To extend resolvent analysis to unstable base flows, we

can consider the use of finite-time/discounted analysis [9, 10] by choosing a complex frequency $\omega = \omega_r + i\beta$, where both ω_r and β are real and β discounts the modal growth rate of $\mathbf{L}_{\bar{\mathbf{q}}}$. The input-output relationship between $\hat{\mathbf{f}}$ and $\hat{\mathbf{q}}$ can be found from (2) as

$$\hat{\mathbf{q}}_\omega = \mathbf{A}\hat{\mathbf{f}}_\omega, \quad (3)$$

where

$$\mathbf{A} = [-i\omega\mathbf{I} - \mathbf{L}_{\bar{\mathbf{q}}}]^{-1} \in \mathbb{C}^{m \times m} \quad (4)$$

is referred to as the *resolvent operator* [6–8]. It serves as a transfer function that amplifies (or attenuates) the harmonic forcing input $\hat{\mathbf{f}}_\omega$ and maps it to the response $\hat{\mathbf{q}}_\omega$. The goal of resolvent analysis is to identify the dominant directions along which $\hat{\mathbf{f}}_\omega$ can be most amplified through $\mathbf{A}$ to form the corresponding responses in $\hat{\mathbf{q}}_\omega$. This question is addressed by the SVD of

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*, \quad (5)$$

where $\mathbf{V}^*$ denotes the Hermitian of $\mathbf{V}$. Resolvent analysis interprets left and right singular vectors $\mathbf{U} = [\hat{\mathbf{u}}_1, \hat{\mathbf{u}}_2, \dots, \hat{\mathbf{u}}_m] \in \mathbb{C}^{m \times m}$ and $\mathbf{V} = [\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \dots, \hat{\mathbf{v}}_m] \in \mathbb{C}^{m \times m}$ respectively as response modes and forcing modes, with the magnitude-ranked singular values $\mathbf{\Sigma} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_m) \in \mathbb{R}^{m \times m}$ being the amplification (gain) for the corresponding forcing-response pair. For unstable base flows, it is important that a finite-time window is chosen with β larger than the highest growth rate such that the resolvent analysis reveals the input-output relationship on a shorter time scale than that of the base flow instability [9, 10].

Performing the SVD of $\mathbf{A} \in \mathbb{C}^{m \times m}$ requires an theoretical operation count of $\mathcal{O}(m^3)$. In practice, some algorithms can reduce this operation count when only a few singular values are to be recovered, while still being computationally taxing for large m [28]. Such cases are encountered in high-Reynolds number flows and bi/tri-global analysis settings. However, we note that many applications of resolvent analysis call only for the dominant forcing and response modes $[\hat{\mathbf{v}}_1, \hat{\mathbf{u}}_1]$ associated with the highest gain σ_1 . This is appropriate when the first gain σ_1 is much larger than the rest of the gains $\sigma_{j>1}$ and shows a quick roll off. Such condition is related to non-normality of linear operator $\mathbf{L}_{\bar{\mathbf{q}}}$, which is encountered for flows with strong shear and separation [29].

When the contributions from the higher-order modes are neglected, it is referred to as the rank-1 assumption for which the flow response from forcing $\hat{\mathbf{f}}$ is approximated as $\hat{\mathbf{q}} \approx \hat{\mathbf{u}}_1 \sigma_1 \langle \hat{\mathbf{v}}_1, \hat{\mathbf{f}} \rangle$, provided that $\sigma_1 \gg \sigma_2$ and $\hat{\mathbf{f}}$ has reasonable magnitude along $\hat{\mathbf{v}}_1$. For seeking only the dominant modal insights from resolvent analysis, we discuss a remedy for performing large-scale resolvent analysis [19, 22] in a computationally tractable manner below.

B. Randomized resolvent analysis

For a flow that have dominant structures, we consider a low-rank representation of the resolvent operator $\mathbf{A}$. That is, instead of directly performing SVD of $\mathbf{A}$ and obtaining the leading-mode representation, we seek a low-rank representation of $\mathbf{A}$ and perform the SVD of the low-rank version of $\mathbf{A}$. We can consider finding an appropriate low-dimensional basis to project the large resolvent operator on a suitable subspace to derive the low-rank resolvent approximation.

The action of a full matrix on a vector should reveal some insights on which components are modified in the dominant directions. In the case of flow that can be described with the rank-1 approximation, there should be a low number of dominant directions. This very point can be taken advantage of through what is known as *sketching* in numerical linear algebra. Sketching refers to a procedure in which a tall and skinny test matrix $\mathbf{\Omega} \in \mathbb{R}^{m \times k}$ (or $\mathbb{C}^{m \times k}$), where $k \ll m$, is passed through $\mathbf{A}$

$$\mathbf{Y} = \mathbf{A}\mathbf{\Omega}. \quad (6)$$

Here, matrix $\mathbf{Y} \in \mathbb{C}^{m \times k}$ is called the *sketch* of the input matrix $\mathbf{A}$ [13, 15, 16]. The test matrix $\mathbf{\Omega}$ can be constructed using random values with Gaussian distribution [30] and be weighted by any input matrix insight, as will be discussed in section IIID for a physics-inspired random test matrix. As the sketch holds the dominant influence of $\mathbf{A}$, we can consider orthonormalizing $\mathbf{Y}$ using a QR decomposition to form the orthonormal basis with $\mathbf{Q} \in \mathbb{C}^{m \times k}$ upon which we can project the full matrix $\mathbf{A}$ to derive its low-rank approximation. In this way, it is possible to approximate $\mathbf{A}$ for a rank $k \ll m$ as long as this approximation preserves the features of the leading modes.

Given this $\mathbf{Q}$, a low-rank approximation of $\mathbf{A}$ can be found as $\mathbf{A} \approx \mathbf{Q}\mathbf{Q}^*\mathbf{A}$ [13]. We can view this as a low-rank decomposition of $\mathbf{A} \approx \mathbf{Q}\mathbf{B}$, where $\mathbf{B} = \mathbf{Q}^*\mathbf{A} \in \mathbb{C}^{k \times m}$. It is this reduced matrix $\mathbf{B}$ upon which we can perform the SVD

$$\mathbf{B} = \tilde{\mathbf{U}}\Sigma\mathbf{V}^* \quad (7)$$

Hence, as a low-rank approximation, we now have

$$\mathbf{A} \approx \mathbf{Q}\tilde{\mathbf{U}}\Sigma\mathbf{V}^* \quad (8)$$

where we can consider $\mathbf{U} \approx \mathbf{Q}\tilde{\mathbf{U}}$. This process is the *randomized SVD* [13], where the sketch $\mathbf{Y}$ was used to derive $\mathbf{Q}$. This approximation almost always satisfies $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| \leq (1 + 9\sqrt{k+p}\sqrt{m})\sigma_{k+1}$, where p is the oversampling parameter, which is applied to build an approximation of rank k while projecting the matrix $\mathbf{A}$ to the low-dimensional subspace with $(k+p)$ vectors. With this overall approach, the computational cost for SVD is reduced to $\mathcal{O}(mk^2)$ instead of $\mathcal{O}(m^3)$ for the full SVD. In our implementation, we use $\mathbf{v}_1$ from (8) and retrieve the leading singular value and left singular vector through

$$\mathbf{A}\hat{\mathbf{v}}_1 = \hat{\mathbf{u}}_1\sigma_1. \quad (9)$$

The singular value and vector can be separated by noticing that $\|\hat{\mathbf{u}}_1\| = 1$. Applied to resolvent analysis, the last equation provides more accurate leading singular value σ_1 and left singular vector $\hat{\mathbf{u}}_1$ compared to the original randomized SVD algorithm [13]. The same operation can be used to recover the higher-order modes, with better accuracy than using the original algorithm [13]. For applications where high-order modes and orthogonality are desired, we can solve for $\mathbf{U}\Sigma$ and compute its SVD. In the present randomized resolvent analysis, we emphasize that only the discrete linear operator $\mathbf{L}_{\bar{\mathbf{q}}}$ is needed for sketching $\mathbf{Y}$ and to find the reduced matrix $\mathbf{B}$. Unlike the original resolvent method, matrix linear solvers are used to avoid calling for the inverse within the resolvent operator. The resulting algorithm constitutes the *randomized resolvent analysis* summarized in Algorithm 1.

To utilize the randomized SVD for resolvent analysis, we must be aware that the resolvent operator $\mathbf{A} = [-i\omega\mathbf{I} - \mathbf{L}_{\bar{\mathbf{q}}}]^{-1}$ contains an inverse operation in its definition, which need not be numerically performed. We do not intend to perform an inverse operation within $\mathbf{A}$ in the present work. In the full resolvent analysis, when the matrices become too large and the inverse can not be performed (which is likely the case for 2D and 3D problems), one can focus on modes corresponding to the smallest singular values of $\mathbf{A}^{-1}$ to find those for the largest singular values of $\mathbf{A}$. Similar approaches have avoided the inverse computation, including the work by Jeun *et al.* [19].

Both full resolvent and randomized resolvent analyses are shown schematically in figure 1. Notice that we are not interested in all singular values and vectors of $[-i\omega\mathbf{I} - \mathbf{L}_{\bar{\mathbf{q}}}]^{-1}$, but only in a few subset of the largest σ_j and their corresponding $\hat{\mathbf{u}}_j$ and $\hat{\mathbf{v}}_j$. In the randomized resolvent analysis, we can approximate a low-rank representation of it using $[-i\omega\mathbf{I} - \mathbf{L}_{\bar{\mathbf{q}}}]$. Figure 1 shows an adaptation of the procedure from Halko *et al.* [13] in order to compute the largest singular values of the resolvent without performing its inverse. In the randomized resolvent analysis, we solve a linear system with $[-i\omega\mathbf{I} - \mathbf{L}_{\bar{\mathbf{q}}}]$, the columns of the random matrix $\mathbf{\Omega}$ form the right-hand side and the sketch columns of $\mathbf{Y}$ are the unknowns. By doing so, we sketch $[-i\omega - \mathbf{L}_{\bar{\mathbf{q}}}]^{-1}$ without finding the actual inverse matrix. The same procedure is performed to project the matrix to the low-dimensional subspace. The matrices are re-arranged in a way that the projection is performed using $[-i\omega\mathbf{I} - \mathbf{L}_{\bar{\mathbf{q}}}]$, but results in the low-dimensional projection of $[-i\omega - \mathbf{L}_{\bar{\mathbf{q}}}]^{-1}$ instead. Figure 1 also illustrates the procedures for recovering the left singular vectors and the singular values from the original algorithm [13] and from the present implementation.


1. Oversampling and power iteration schemes

Randomized algorithms can incorporate two additional procedures to improve performance and accuracy. Namely, they are oversampling [31] and power or subspace iterations [13, 18, 32]. Oversampling sketches the input matrix using $(k+p)$ vectors (with p extra vectors) and increases the low-dimensional subspace to accurately recover a smaller quantity of singular values k . For the randomized resolvent analysis, oversampling has the same outcome, in practice, of selecting a larger k and the influence of k will be discussed at the end of section III B. When k becomes large, it should be noticed that the memory consumption increases. Even for large sparse matrices, the sketch matrix and the subsequently reduced matrices that are formed are generally dense, which adds a computational burden.

The second procedure is the power or subspace iterations. These methods are a powerful tool when the singular values of the matrix decay slowly. For example, this type of spectral behavior appeared in the input-output analysis performed by Jeun *et al.* [19] for jet flows. The method consists of performing additional iterations after the sketch $\mathbf{Y}$ is evaluated. It should however be realized that such procedure calls for additional linear solvers, which is the most

(a) Full Resolvent Analysis

resolvent operator



(b) Randomized Resolvent Analysis

Solve Linear System for Y




FIG. 1. Schematics of the (a) full resolvent and (b) randomized resolvent analyses. For the full resolvent analysis, a direct SVD is applied. For the randomized resolvent analysis, the orthogonal basis Q is computed to project the operator into the low-dimensional subspace, where the SVD is performed on the reduced matrix B . Two options to recover the left singular vectors and singular values are provided, using the original randomized approach [13] and the present implementation. Orthogonal vectors are represented by long bars.

time consuming operation in Algorithm 1. For power iterations, one must compute the adjoint AA^* , where A^* is the Hermitian of A , q times and solve the linear system k times. For subspace iterations, additional QR decompositions are necessary, and the number of additional linear systems to be solved will be q times k . In practice, small values of q improves the accuracy of the results substantially. More information on the general applicability of the subspace and power iterations are discussed by Halko *et al.* [13] and Erichson *et al.* [31]. In section III D, we present another option for improving accuracy of the overall technique in a computationally inexpensive manner by constructing a

Algorithm 1: Randomized Resolvent Analysis

Require: Discrete linear operator $L_{\bar{q}} \in \mathbb{C}^{m \times m}$

Function randomized_resolvent(ω, k):

```

1   $\Omega \leftarrow \text{randn}(m, k)$  // Random normal matrix generator. For scaling, see Section IIID
2   $Y \leftarrow [-i\omega I - L_{\bar{q}}] \backslash \Omega$  // Solve linear system for  $Y$ ,  $\mathcal{O}(m^2 k)$ 
3   $(Q, \sim) \leftarrow \text{qr}(Y, 0)$  // Economy-sized QR decomposition,  $\mathcal{O}(mk^2)$ 
4   $B \leftarrow Q^* / [-i\omega I - L_{\bar{q}}]$  // Solve linear system for  $B$ ,  $\mathcal{O}(m^2 k)$ 
5   $(\sim, \sim, V) \leftarrow \text{svd}(B, \text{'econ'})$  // Reduced SVD decomposition,  $\mathcal{O}(mk^2)$ 
6   $U^\Sigma \leftarrow [-i\omega I - L_{\bar{q}}] \backslash V$  // solve linear system to recover  $U^\Sigma$ ,  $\mathcal{O}(m^2 k)$ 
7  for  $j \leftarrow 1$  to  $k$  do
8     $\Sigma_{j,j} \leftarrow \text{norm}(U_{1:m,j}^\Sigma, 2)$  // Recover singular values  $\Sigma$ 
9     $U_{1:m,j} \leftarrow U_{1:m,j}^\Sigma / \Sigma_{j,j}$  // Normalize  $U$ 
10 end
11  $(U, \Sigma, \tilde{V}) \leftarrow \text{svd}(U^\Sigma, \text{'econ'})$  // (Optional) Recover  $U$  and  $\Sigma$ ,  $\mathcal{O}(mk^2)$ 
12  $V \leftarrow V \tilde{V}$  // (Optional) Recover improved  $V$ ,  $\mathcal{O}(mk)$ 
13 return  $(U, \Sigma, V)$ 

```

physics-informed random test matrix Ω .

2. Test matrix Ω

The standard choice for the test matrix Ω is the random matrix generated with a normal distributions. Such matrix is known to present excellent performance and accuracy [15]. For some cases, especially for very large matrices or when the singular values present slow decay, larger values of k may be necessary to better approximate the matrix in the low-dimensional subspace. When large values of k are used, orthonormalization of the columns of the test matrix can be considered to improve numerical stability [13, 33]. The test matrix can also be generated using a Rademacher distribution [17]. It is also possible to build an ultrasparse matrix with Rademacher distribution in the non-sparse entries which allows for the control of cost, stability and reliability in the operations [15, 17].

When randomized SVD is applied, there is no a priori knowledge of the structure of the matrix. However, in the present application, we know how the resolvent operator is constructed. This theoretical insight can be used to build a random test matrix that outperforms the standard normal distribution matrix. We later propose a physics-informed test matrix Ω that can focus our sketching operation for regions of physical importance. In our application, the dominant directions are related to regions with the presence of high shear. The results from this approach will be discussed in Section IIID.

III. RANDOMIZED RESOLVENT ANALYSIS OF TURBULENT POST-STALL FLOW

We demonstrate the use of randomized resolvent analysis on turbulent flow over a NACA 0012 airfoil. In this example, the randomized resolvent analysis will be applied a resolvent operator of size $m \times m$, where $m \simeq 7 \times 10^5$, to reveal the dominant gain and modal structures with a thin sketching matrix having as little as $k = 10$ columns. The convergence of the gain and resolvent modes will also be reported with respect to the size of the sketching matrix. Influence of the ratio between the first and the second singular values of the resolvent operator will also be examined.

A. Problem setup

We consider the spanwise-periodic turbulent flow over a NACA 0012 airfoil at an angle of attack of 9° , a chord-based Reynolds number of $Re_{L_c} \equiv v_\infty L_c / \nu_\infty = 23,000$ and a free stream Mach number of $M_\infty \equiv v_\infty / a_\infty = 0.3$. Here, v_∞ is the free-stream velocity, L_c is the chord length, a_∞ is the free-stream sonic speed, and ν_∞ is the kinematic viscosity. The time- and spanwise-averaged turbulent flow is considered as the base flow for the full and randomized resolvent analyses. For this 2D base flow, we adopt the bi-global setting that decomposes $\mathbf{q}'$ into spanwise Fourier modes with the wavenumber k_z .

To obtain the base flow, large-eddy simulation (LES) is performed using a finite-volume compressible flow solver *CharLES* [34, 35], which is second-order accurate in space and third-order accurate in time. Vremen's sub-grid




FIG. 2. The instantaneous (left) and time/spanwise-averaged (right) flows over a NACA 0012 airfoil at $Re_{L_c} = 23,000$. The instantaneous flow visualization shows the isosurface of Q -Criterion ($QL_c^2/v_\infty^2 = 50$) colored by the instantaneous streamwise velocity.

scale model [36] is utilized in the LES. The LES is performed on a C-shaped mesh with the domain extent of $x/L_c \in [-19, 26]$, $y/L_c \in [-20, 20]$ and $z/L_c \in [-0.1, 0.1]$ in the streamwise, transverse and spanwise direction, respectively, with the airfoil leading edge at $x/L_c = y/L_c = 0$. Dirichlet boundary condition is specified at the far-field boundary as $(\rho, v_x, v_y, v_z, T) = (\rho_\infty, v_\infty, 0, 0, T_\infty)$, where ρ is the density, v_x , v_y and v_z are respectively the streamwise, transverse and spanwise velocity, and T is the temperature. Over the airfoil, the no-slip adiabatic boundary condition is prescribed. Along the outlet boundary, a sponge layer [37] is applied with a running-averaged state being the target state. The simulation has been validated with respect to the time-averaged pressure, lift and drag over the airfoil. The turbulent separated flow over the airfoil is visualized in figure 2. The visualization of the instantaneous flow shows the laminar separation from the leading edge. We have found that the shear layer physics dominates the pseudospectral behavior of the linearized Navier–Stokes operator, as shear is the main source of nonnormality in the operator. Further details regarding the computational setup, flow physics, and resolvent analysis based flow control of this setup are reported in Yeh and Taira [10].

The full and randomized resolvent analyses are performed on a separate mesh from that used in the LES. This mesh has a 2D rectangular domain with the extent of $x/L_c \in [-15, 16]$ and $y/L_c \in [-12, 12]$, comprising approximately 0.15 million cells. Compared to the LES mesh, the mesh for resolvent analysis is coarser over the airfoil and in the wake, but is much finer in the upstream of the airfoil in order to resolve the forcing mode structures. The time- and spanwise-averaged flow $\bar{\mathbf{q}}$ obtained from LES is interpolated onto this mesh. At the far-field boundary and over the airfoil, Dirichlet conditions are set for density and velocities and Neumann condition is prescribed for pressure in $\mathbf{q}'$. At the outlet boundary, Neumann condition is set for all flow variables. With these boundary conditions for $\mathbf{q}'$ and the base flow $\bar{\mathbf{q}}$, we construct the linearized Navier–Stokes operator $\mathbf{L}_{\bar{\mathbf{q}}}(k_z)$ for a chosen k_z . The size of $\mathbf{L}_{\bar{\mathbf{q}}}$ and the resolvent operator is approximately $0.75 \text{ million} \times 0.75 \text{ million}$.

For this large operator, we summarize in Table I the computational costs of performing resolvent analysis using the Krylov-based Arnoldi-iteration method with a range of parameter setups (i.e., number of singular values (n_{ev}), Krylov subspace dimension ($\dim(\mathcal{S})$), and tolerance) and compare them with those for the present randomized algorithm. The former was conducted by simply calling the `svds` command in MATLAB. It requires almost 80 gigabytes of memory and takes approximately 30 to 70 minutes (single-core) for each SVD. The high-memory demand necessitates the use of high performance computing resource to conduct the full resolvent analysis. In contrast, the randomized resolvent approach (Algorithm 1) achieves significant reductions in computational time and memory consumption. The present method only requires a third of the memory usage of the Arnoldi-iteration and cuts down the computational time by an order of magnitude. We also note that, in Algorithm 1, the linear systems solvers are the operations with higher computational cost. Since all the three linear systems solvers are conducted for the same operator, the LU decomposition of $[-i\omega\mathbf{I} - \mathbf{L}_{\bar{\mathbf{q}}}]$ is performed in the beginning of the algorithm and is passed through the three solvers. This decomposition becomes the main source of the memory consumption.

B. Results

We perform the full and randomized resolvent analyses for spanwise wavenumbers of $k_z L_c = 0$ and 20π . Since the base flow is found to be unstable [10], the finite-time approach is adopted with $v_\infty/\beta L_c = 3$ to ensure that the resolvent analysis is performed on a shorter time scale than that associated with the leading growth rate of the instability. Initially, for the randomized analysis, we consider $k = 10$ for the width of the test matrix $\mathbf{\Omega}$ with

Iteratively restarted Arnoldi method (MATLAB svds)				
n_{ev}	$\dim(\mathcal{S})$	tolerance	time (sec)	memory
10	30	10^{-14}	4185	78.6 GB
5	15	10^{-14}	2764	78.6 GB
2	6	10^{-14}	1486	78.6 GB
10	30	10^{-5}	2245	78.6 GB
10	15	10^{-14}	4194	78.6 GB

Randomized resolvent (present)			
n_{ev}	k	time (sec)	memory
2	2	354	28.2 GB
5	5	462	28.4 GB
10	10	615	28.8 GB

TABLE I. Comparison of the computational time and memory consumption for the implicitly restarted Arnoldi iteration (svds in MATLAB) with the present randomized approach for different parameter setups with the number of singular values (n_{ev}), Krylov subspace dimension ($\dim(\mathcal{S})$), and tolerance.

random Gaussian distribution. Later in Section IIID, we show that this value of k can be further reduced without compromising accuracy.

The leading response and forcing modes obtained from both full and randomized analyses are compared in figure 3 for representative frequencies St and spanwise wavelength $k_z L_c$. Although $k/m = 1.3 \times 10^{-5}$, we observe excellent agreement between the modes from the full resolvent analysis and the randomized algorithm. We observe that randomized forcing and response modes are very similar to full resolvent ones. Only at $St = 1$ and $k_z L_c = 20\pi$ we observe the appearance of spatially distributed errors in the background, which we refer to as background noise.

The forcing modes are recovered directly from the SVD of the low-dimensional subspace projection. As they are used to recover the response modes using the linear operator, the accuracy of the forcing modes affect the results of the response ones. In the particular case of $St = 1$ and $k_z L_c = 20\pi$, when forcing mode is affected by noise, the randomized approach returns some structures emanating from the trailing edge in the response mode $\hat{\mathbf{u}}_1^{\text{rand}}$, which was not present from the full resolvent analysis. This behavior is related to leakage from high-order modes, as σ_1 and σ_2 are close in the energy spectrum. This remarkable level of agreement over all frequencies and spanwise wavenumbers ensures that the randomized approach presented in Algorithm 1 can help extract insights into the spatial structures to identify regions of sensitivity and guide flow control efforts.

These results were obtained using the present implementation that extends the original randomized SVD algorithm. In figure 4, using the original randomized algorithm [13] within the resolvent analysis to recover the left singular vectors and singular values, the response modes contain background noise. Here, we use Algorithm 1 up to line 5, then SVD is performed as $\mathbf{B} = \tilde{\mathbf{U}}\Sigma\mathbf{V}^*$ and $\mathbf{U}$ is recovered a posteriori using the original procedure for the randomized SVD with the resolvent analysis [13, 25], by $\mathbf{U} = \mathbf{Q}\tilde{\mathbf{U}}$. Yeh and Taira [10] showed that the present problem setup presents a peak in the singular values near $St = 6$ for both spanwise wavenumbers, influenced by the eigenmodes associated with the shear-layer structure over the separation bubble. These eigenmodes are highly nonnormal and induce high-energy amplification through pseudo-resonance [6]. For the frequencies in a narrow region near $St = 6$, both implementations present similar results. Far from this band, the response modes obtained by the original procedure [13] are contaminated by random background noise or leakage from higher-order modes, as shown in figure 4 for $St = 0.5$ and 15 and $k_z L_c = 0$. In these critical cases, the original randomized approach may not provide meaningful insights into flow physics. The present implementation shown in Algorithm 1 improves solving for the response modes. Algorithm 1 does not enhance the forcing modes, as they are already accurate. When utilizing this technique to generate reduced-order models, one may perform steps 11 and 12 in Algorithm 1 to orthogonalize the left singular vectors. Considering the results obtained by our implementation, the modes are found very accurately for almost all frequencies and wavenumbers. To provide a concrete assessment, we quantitatively assess the accuracy of the randomized resolvent analysis.

The agreement between the full and randomized analyses with respect to the gain (leading singular value) and modes over a range of frequencies is presented in figure 5. When Algorithm 1 is applied to recover left singular vectors and singular values, the randomized analysis accurately captures the trend of gain distribution over $1 \lesssim St \lesssim 15$ in figure 5(a,b) for both wavenumbers. At the low and high-frequency ends, the gain shows deviations. The resemblance of the modal structure is quantified in figure 5(c,d) with the cosine similarities, i.e., the inner products, $\langle \hat{\mathbf{u}}_1^{\text{full}}, \hat{\mathbf{u}}_1^{\text{rand}} \rangle$ and $\langle \hat{\mathbf{v}}_1^{\text{full}}, \hat{\mathbf{v}}_1^{\text{rand}} \rangle$. As singular vectors are normalized, the cosine similarity of 1 suggests that perfect match is attained between the modes from full and randomized resolvent analyses. Since these modes are complex, the cosine similarity




FIG. 3. The leading response ($\hat{\mathbf{u}}_1$) and forcing ($\hat{\mathbf{v}}_1$) modes from the full and randomized resolvent analyses, using Algorithm 1, for $k_z L_c = 0$ and 20π at representative frequencies St . Modes are visualized with the streamwise velocity component with contour levels of $[-0.6, 0.6]$.

removes dependence on the phase difference. For almost the entire range of frequencies the cosine similarities are near unity, which means the agreement between full and randomized modes is excellent. When this value is reduced, the modes may be affected by noise, as seen for $\hat{\mathbf{v}}_1^{\text{rand}}$ at $St = 1$ and $k_z L_c = 0$ in figure 4. For the original approach, using the randomized SVD algorithm [13], the modes have good agreement for a narrow band of frequencies only. By comparing the results from figures 3 and 4 to the values in figure 5(c,d), we observe that the noise affects the modes when cosine similarity is below 0.5. For frequencies and wavenumbers with cosine similarity up to 0.8 or higher, there is no noise and the results for full and randomized resolvent agree well. For this reason, it is desirable to search for solutions that provide a reliable agreement up to this scale to a broad range of frequencies and both wavenumbers. With the high-gain frequency range well captured, randomized resolvent analysis has demonstrated its capability of predicting the dominant pathway for energy amplification over the spectral space with reduced computational cost.

As stated in section II B, the use of low-rank approximation in the randomized approach is built upon the assumption of the low-rank nature of the resolvent operator. The randomized resolvent analysis shows its strength when the




FIG. 4. Comparison of the leading response modes ($\hat{\mathbf{u}}_1$) recovered from the randomized resolvent analyses using $\mathbf{U} = \mathbf{Q}\tilde{\mathbf{U}}$ (the original approach in Halko *et al.* [13]) and $\mathbf{U} = \mathbf{A}\mathbf{V}\Sigma^{-1}$ (present, see equation (9)). The response modes from the full analysis are also shown for reference. Results are shown for $k_z L_c = 0$ at representative frequencies St . Modes are visualized with the streamwise velocity component with contour levels of $[-0.6, 0.6]$.

singular values exhibit fast decay, as evident from figure 6. The accuracy of the modal structure captured by randomized analysis is examined with respect to the ratio of the leading and second singular values, $(\sigma_1/\sigma_2)^{\text{full}}$ from the full resolvent analysis. The error in the modal structures exhibits a decreasing trend as this ratio increases. When this ratio is close to unity, the randomized technique may not accurately separate the first and second modes. In fact, the aforementioned trailing edge structure that appeared in the randomized response mode for $k_z L_c = 20\pi$ and $St = 1$ is caused by the leakage of the structures from the second response mode (see figure 3). When the ratio $(\sigma_1/\sigma_2)^{\text{full}}$ is above 30, the error decreases to $\lesssim 10^{-5}$ for forcing modes and $\lesssim 10^{-8}$ for response modes.

Next, we study the influence of the width of the test matrix k on the error in the leading singular values and modes, as presented in figure 7. When the value of k is varied from 2 to 500, the error from the use of randomized analysis decreases. For three representative frequencies, we observe the same rate of convergence $\approx \mathcal{O}(k)$ for both the gain and cosine similarity. As stated in section II B, increasing k has the same practical effect of oversampling, in the present application. For this flow, we observe that $k = 10$ is sufficient to achieve sufficient accuracy with $\lesssim 1\%$ error, which is remarkably low when compared to the high dimensionality of the resolvent operator.

The computational cost of the randomized resolvent technique can be further reduced. For instance, the biconjugate gradient stabilized (BiCGStab) or the generalized minimum residual (GMRES) methods can be utilized to solve linear systems with appropriate preconditioners (e.g., incomplete LU and Jacobian).

C. Higher-order modes

Let us discuss the performance of the randomized technique with respect to the high-order modes. For some cases, the second largest singular value may also be spaced apart from the higher-order singular values and be determined accurately. In figure 8, we show for $k_z L_c = 0$ and $St = 4$ a case where both the leading and second singular values are spread from the rest of the singular values. In this case, the randomized algorithms accurately capture the second modes. The flow structures at the trailing edge are perceived in the response modes. The forcing modes appear over the pressure side near the trailing edge. For the results obtained from Algorithm 1, the modes are the same for randomized resolvent and for the full resolvent. However, for the resolvent analysis using the randomized SVD algorithm [13] within the resolvent analysis, the secondary singular values are poorly captured and the modes are polluted by background noise and leakage from other modes (not shown here). For this reason, when applying randomized resolvent, the present implementation shown in Algorithm 1 must be considered as they can approximate the detached high-order modes accurately.

D. Choice of the test matrix Ω

For the randomized resolvent analysis with a test matrix size of $k = 10$, Gaussian, orthonormal, Rademacher and ultrasparse Rademacher test matrices provide similar results and no observable difference in computational savings. By using the implementation shown in Algorithm 1, all test matrices present similar accuracy as shown in figures 5




FIG. 5. (a,b) The leading amplification for the full and randomized resolvent analyses. (c-d) Cosine similarities for the leading response $\langle \hat{\mathbf{u}}_1^{\text{full}}, \hat{\mathbf{u}}_1^{\text{rand}} \rangle$ and forcing $\langle \hat{\mathbf{v}}_1^{\text{full}}, \hat{\mathbf{v}}_1^{\text{rand}} \rangle$ modes. Improvements of accuracy in the present randomized analysis can be observed by comparing the σ_1^{rand} and $\hat{\mathbf{u}}_1^{\text{rand}}$ recovered from equation (9) to those from the original approach of Halko *et al.* [13] (note that forcing modes obtained from both approaches are identical). Results for spanwise wavenumber of $k_z L_c = 0$ and 20π are shown in the left and right columns, respectively.

and 6. For very low and very high St numbers in the range of frequencies analyzed in this work, where the singular values decay slowly, one can increase the size k or apply subspace or power iterations. However, it is possible to obtain more accurate results by constructing a random test matrix Ω that incorporates physical insights from the base flow.

While the random test matrix is effective in yielding accurate results, we can consider constructing a test matrix that can generate the entries in a smart manner by incorporating the knowledge of the base flow. We know that regions of strong shear are important in amplifying forcing inputs. Moreover, regions with minimal velocity gradients are not that important. For these reasons, the velocity gradient at each grid point can be used to scale the test matrix. Here, we propose a physics-informed test matrix, Ω_p , scaled by the 2-norm of the velocity gradient, $\|\nabla \mathbf{v}\|_2$, where $\mathbf{v}$ is the velocity vector. We construct a scaling factor $\Phi_j = \|\nabla \mathbf{v}_j\|_2$ at each grid point j . The scaling vector Φ has to be stacked according to the number of variables to reach the size m of the linear operator. The physics-informed test matrix then becomes

$$\Omega_p = \text{diag}(\Phi)\Omega. \quad (10)$$

The results based on the physics-informed test matrix are shown in figure 9. While the results obtained from the use of a normal distribution test matrix have shown excellent accuracy, the use of physics-informed test matrix further improves the accuracy by a few orders of magnitude for the considered frequencies and spanwise wavenumbers. More importantly, we achieve results with comparable or higher level of accuracy using a extremely low width of the test matrix of $k = 2$. This results in a considerable reduction in computation time (see Table I). Using $k = 10$, linear systems are solved at least 30 times. Now, using $k = 2$, only 6 linear solvers are needed to obtain the same accuracy, which is achieved only with a physics-informed scaling of the test matrix. By combining randomized numerical linear algebra and some physical insights, we are now empowered to perform the input-output analysis for ever more complex 2D and 3D turbulent base flows on a standard computer, or on a high-performance computing cluster to expand the envelope of resolvent analysis.




FIG. 6. Error based on cosine similarity for the leading (a) forcing modes ($1 - \langle \mathbf{v}_1^{\text{full}}, \mathbf{v}_1^{\text{rand}} \rangle$) and (b) response modes ($1 - \langle \mathbf{u}_1^{\text{full}}, \mathbf{u}_1^{\text{rand}} \rangle$) over the leading gap from the full resolvent analysis. The symbols $\circ$ and $\triangle$ represent results from $k_z L_c = 0$ and 20π , respectively, colored by St .




FIG. 7. Influence of test matrix size k on the accuracy of (a) leading singular value, (b) forcing mode, and (c) response mode at $k_z L_c = 20\pi$. All exhibit $\mathcal{O}(k)$ convergence.

IV. CONCLUSION

Resolvent analysis has proven to be a powerful technique to reveal the input-output characteristics of fluid flows. However, the computational cost and memory allocation of the resolvent analysis can be taxing for high-Reynolds number flows, making it prohibitive to be applied to complex turbulent base flows. The major computational cost of the analysis is associated with the SVD of the resolvent operator. To remove this bottleneck, the randomized approach has been adopted to reduce the computational cost of SVD by considering the low-rank approximation of the resolvent operator. This was achieved by constructing the low-rank basis based on the insights gained from the sketch of the resolvent, which is obtained from a linear system solver. For flows with fast singular value decays, e.g., flows with strong shear and separation, the randomized resolvent analysis reveals its power to accurately capture the response and forcing modes as well as the gain. Moreover, we consider the use of the velocity gradient to scale the random test matrix. Such scaling enhanced the accuracy of the randomized resolvent analysis. The necessary computation time was significantly reduced as the number of linear systems to be solved is considerably smaller. To demonstrate the capability of the randomized resolvent analysis, we analyzed the turbulent separated flow over a NACA 0012 airfoil. Excellent agreement of the leading forcing and response modes and the gains were shown between the full and randomized resolvent methods. By incorporating the knowledge of the base flow in terms of its velocity gradient into the randomized test matrix, additional speed up and accuracy enhancement were achieved. With the computational cost and memory allocation being relieved with the randomized approach, the application of the resolvent analysis




FIG. 8. Recovery of the higher-order resolvent gains ($\sigma_{i=1-6}^{\text{rand}}$) and modes ($\hat{\mathbf{v}}_2$ and $\hat{\mathbf{u}}_2$) for $St = 4$ and $k_z L_c = 0$. The resolvent gains obtained from the present randomized analysis ($\diamond$, recovered by equation (9)) and those from the original approach of Halko *et al.* [13] ($\triangle$) are compared to those from the full resolvent analysis ($\circ$).




FIG. 9. Influence of test matrices on the relative error for (a) gain distribution and cosine similarities for (b) forcing and (c) response modes, $\hat{\mathbf{v}}_1$ and $\hat{\mathbf{u}}_1$, at $k_z L_c = 20\pi$. Results are shown for Gaussian random normal distribution test matrix Ω and physics-informed random test matrices $\Omega_p = \text{diag}(\Phi)\Omega$ with sizes $k = 10$ and $k = 2$.

can be significantly extended to higher-Reynolds number 2D and 3D base flows.

ACKNOWLEDGMENTS

We gratefully acknowledge the support from the Office of Naval Research (N00014-16-1-2443), Air Force Office of Scientific Research (FA9550-18-1-0040), and Army Research Office (W911NF-17-1-0118). We thank L. Mathelin, S. L. Brunton, and N. B. Erichson for the enlightening discussions. Some of the computations presented here were supported by the Department of Defense High Performance Computing Modernization Program.

-
- [1] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, *Turbulence, coherent structures, dynamical systems and symmetry*, Cambridge Univ. Press 2nd ed. (2012).
 - [2] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5–28 (2010).
 - [3] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordoyev, V. Theofilis, and L. S. Ukeiley, Modal analysis of fluid flows: An overview, *AIAA J.* **55**(12) 4013–4041 (2017).
 - [4] K. Taira, M. S. Hemati, S. L. Brunton, Y. Sun, K. Duraisamy, S. Bagheri, S. T. M. Dawson, and C.-A. Yeh, Modal Analysis of Fluid Flows: Applications and Outlook, *AIAA J.* accepted, available online (2019).

- [5] V. Theofilis, Global linear instability, *Annu. Rev. Fluid Mech.* **43**, 319–352 (2011).
- [6] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, Hydrodynamic stability without eigenvalues, *Science* **261** 578–584 (1993).
- [7] M. R. Jovanović and B. Bamieh, Componentwise energy amplification in channel flows, *J. Fluid Mech.* **534** 578–584 (2005).
- [8] B. J. McKeon and A. S. Sharma, A critical-layer framework for turbulent pipe flow, *J. Fluid Mech.* **658** 336–382 (2010).
- [9] M. R. Jovanović, Modeling, analysis, and control of spatially distributed systems, PhD Thesis, University of California at Santa Barbara, Dept. of Mechanical Engineering (2004).
- [10] C.-A. Yeh and K. Taira, Resolvent-analysis-based design of airfoil separation control, *J. Fluid Mech.* **867** 572–610 (2019).
- [11] S. Nakashima, K. Fukagata, and M. Luhar, Assessment of suboptimal control for turbulent skin friction reduction via resolvent analysis, *J. Fluid Mech.* **828** 496–526 (2017).
- [12] T. Kajishima and K. Taira, *Computational fluid dynamics: incompressible turbulent flows*, Springer (2017).
- [13] N. Halko, P.-G. Martinsson, and J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Rev.* **53** 217–288 (2011).
- [14] P. Drineas and M. W. Mahoney, RandNLA: randomized numerical linear algebra, *Comm. ACM* **59** 80–90 (2016).
- [15] J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher, Practical sketching algorithms for low-rank matrix approximation, *SIAM J. Matrix Anal. Appl.* **38** 1454–1485 (2017).
- [16] F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert, A fast randomized algorithm for the approximation of matrices, *Appl. Comput. Harmon. Anal.* **25** 335–366 (2008).
- [17] K. L. Clarkson and D. P. Woodruff, Numerical linear algebra in the streaming model, *Proceedings of the forty-first annual ACM symposium on Theory of computing ACM* 205–214 (2007).
- [18] V. Rokhlin, A. Szlam, and M. Tygert, A randomized algorithm for principal component analysis, *SIAM J. Matrix Anal. Appl.* **31** 1100–1124 (2009).
- [19] J. Jeun, J. W. Nichols, and M. R. Jovanović, Input-output analysis of high-speed axisymmetric isothermal jet noise, *Phys. Fluids* **28** 047101 (2016).
- [20] J. H. M. Ribeiro and W. R. Wolf, Identification of coherent structures in the flow past a NACA0012 airfoil via proper orthogonal decomposition, *Phys. Fluids* **29** 085104 (2017).
- [21] T. R. Ricciardi, W. R. Wolf, and R. Speth, Acoustic prediction of lagoon landing gear: Cavity noise and coherent structures, *AIAA J.* **56** 4379–4399 (2018).
- [22] O. T. Schmidt, A. Towne, G. Rigas, T. Colonius, and G. A. Brès, Spectral analysis of jet turbulence, *J. Fluid Mech.* **855** 953–982 (2018).
- [23] A. Dwivedi, G. S. Sidharth, J. W. Nichols, G. V. Candler, and M. R. Jovanović, Reattachment streaks in hypersonic compression ramp flow: an input–output analysis, *J. Fluid Mech.* **880** 113–135 (2019).
- [24] N. B. Erichson, L. Mathelin, J. N. Kutz, and S. L. Brunton, Randomized Dynamic Mode Decomposition, *SIAM J. Appl. Dyn. Syst.* **18**(4) 1867–1891 (2019).
- [25] R. Moarref, A. S. Sharma, J. A. Tropp, and B. J. McKeon, Model-based scaling of the streamwise energy density in high-Reynolds-number turbulent channels, *J. Fluid Mech.* **734** 275–316 (2013).
- [26] B. F. Farrell and P. J. Ioannou, Variance maintained by stochastic forcing of non-normal dynamical systems associated with linearly stable shear flows, *Phys. Rev. Lett.* **72** 1188–1191 (1994).
- [27] P. J. Schmid, Nonmodal stability theory, *Annu. Rev. Fluid Mech.* **39**(1) 129–162 (2007).
- [28] R. B. Morgan and M. Zeng, A harmonic restarted Arnoldi algorithm for calculating eigenvalues and determining multiplicity, *Lin. Alg. Appl.* **415** 96–113 (2006).
- [29] P. J. Schmid and D. S. Henningson, *Stability and Transition in Shear Flows*, Springer (2001).
- [30] P.-G. Martinsson, V. Rokhlin, and M. Tygert, A randomized algorithm for the decomposition of matrices, *Appl. Comput. Harmon. Anal.* **30** 47–68 (2011).
- [31] N. B. Erichson, S. Voronin, S. L. Brunton, and J. N. Kutz, Randomized matrix decompositions using R, *J. Stat. Soft.* **89**(11) 1–48 (2019).
- [32] M. Gu, Subspace iteration randomization and singular value problems, *SIAM J. Sci. Comput.* **37** A1139–A1173 (2015).
- [33] J. Demmel, I. Dumitriu, and O. Holtz, Fast linear algebra is stable, *Num. Math.* **108** 59–91 (2007).
- [34] 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, *AIAA Paper* 2011-2886 (2011).
- [35] G. A. Brès, F. E. Ham, J. W. Nichols, and S. K. Lele, Unstructured Large-Eddy Simulations of Supersonic Jets, *AIAA J.* **55** 1164–1184 (2017).
- [36] A. W. Vreman, An eddy-viscosity subgrid-scale model for turbulent shear flow: algebraic theory and applications, *Phys. Fluids* **16** 3670–3681 (2004).
- [37] J. B. Freund, Proposed inflow/outflow boundary condition for direct computation of aerodynamic sound, *AIAA J.* **35** 740–742 (1997).