Variants of dynamic mode decomposition: boundary conditions, Koopman, and Fourier analyses*

Kevin K. Chen,[†‡] Jonathan H. Tu,[†] and Clarence W. Rowley[†]

Oct 24, 2011

Abstract

Dynamic mode decomposition (DMD) is an Arnoldi-like method based on the Koopman operator that analyzes empirical data, typically generated by nonlinear dynamics, and computes eigenvalues and eigenmodes of an approximate linear model. Without explicit knowledge of the dynamical operator, it extracts frequencies, growth rates, and spatial structures for each mode. We show that expansion in DMD modes is unique under certain conditions. When constructing mode-based reduced-order models of partial differential equations, subtracting a mean from the data set is typically necessary to satisfy boundary conditions. Subtracting the mean of the data exactly reduces DMD to the temporal discrete Fourier transform (DFT); this is restrictive and generally undesirable. On the other hand, subtracting an equilibrium point generally preserves the DMD spectrum and modes. Next, we introduce an "optimized" DMD that computes an arbitrary number of dynamical modes from a data set. Compared to DMD, optimized DMD is superior at calculating physically relevant frequencies, and is less numerically sensitive. We test these decomposition methods on data from a two-dimensional cylinder fluid flow at a Reynolds number of 60. Time-varying modes computed from the DMD variants yield low projection errors.

*Submitted to the Journal of Nonlinear Science

[†]Department of Mechanical & Aerospace Engineering, Princeton University, Princeton, NJ 08544

[‡]Email address for correspondence: kkchen@princeton.edu

1 Introduction

Data-based modal decomposition is frequently employed to investigate complex dynamical systems. When dynamical operators are too sophisticated to analyze directly—because of nonlinearity and high dimensionality, for instance—data-based techniques are often more practical. Such methods numerically analyze empirical data produced by experiments or simulations, yielding modes containing dynamically significant structures. In fluid flow control, for example, decompositions of an evolving flow field may identify coherent structures such as vortices and eddies that are fundamental to the underlying flow physics (Holmes et al. 1996). A subset of the computed modes can then form the basis of a reduced-order model—that is, a low-dimensional approximation of the original dynamical system—using Galerkin projection. This is often critical in applications such as controller or observer design in a control system, where matrix operations on a high-dimensional dynamical system could be computationally intractable.

One common decomposition technique, proper orthogonal decomposition (POD, otherwise known as principal component analysis or Karhunen-Loève decomposition), generates orthogonal modes that optimally capture the vector energy of a given data set (Holmes et al. 1996). This method, while popular, suffers from a number of known issues. For instance, reduced-order models generated by POD may be inaccurate, because the principal directions in a set of data may not necessarily correspond with the dynamically important ones. As a result, the selection of POD modes to retain in a reduced-order model is nontrivial and potentially difficult (Ilak and Rowley 2008). Improved methods are available, such as balanced truncation (Moore 1981) or balanced POD (an approximation of balanced truncation for high-dimensional systems; Rowley 2005), but both of these methods are applicable only to linear systems.

Dynamic mode decomposition (DMD) is a relatively recent development in the field of modal decomposition (Rowley et al. 2009; Schmid 2010). Dynamic mode decomposition approximates the modes of the Koopman operator, which is a linear, infinite-dimensional operator that represents nonlinear, finite-dimensional dynamics without linearization (Mezić and Banaszuk 2004; Mezić 2005), and is the adjoint of the Perron-Frobenius operator. The method can be viewed as computing, from empirical data, eigenvalues and eigenvectors of a linear model that approximates the underlying dynamics, even if those dynamics are nonlinear. Unlike POD and balanced POD, this decomposition yields growth rates and frequencies associated with each mode, which can be found from the magnitude and phase of each corresponding eigenvalue. If the data are generated by a linear dynamical operator, then the decomposition recovers the leading eigenvalues and eigenvectors of that operator; if the data are periodic, then the decomposition is equivalent to a temporal discrete Fourier transform (DFT) (Rowley et al. 2009). Applications of DMD to experimental and numerical data can be found in Rowley et al. (2009), Schmid (2010, 2011), and Schmid et al. (2011).

In this paper, we present some new properties and some modifications of DMD, highlighting the relationship with traditional Fourier analysis. Section 2 briefly reviews the Koopman operator, DMD, and two algorithms for computing DMD. In Section 3, we prove the uniqueness of the decomposition under specific conditions. Section 4 discusses applications to reduced-order models of partial differential equations, and in particular compares methods to ensure that models based on DMD satisfy boundary conditions correctly. We explore the subtraction of either the data mean or an equilibrium point (in fluid mechanics, this is commonly known as "subtracting a base flow"), and examine the consequences of each. In Section 5, we present advances toward an alternate, "optimized" DMD formulation, which tailors the decomposition to the user-specified number of modes, and is more accurate than truncated DMD models. We provide examples in Section 6 by considering numerically generated data of a cylinder fluid flow, evolving via the nonlinear Navier-Stokes equations from an unstable equilibrium to a limit cycle.

2 The Koopman operator and dynamic mode decomposition

In this section, we briefly review the results of Rowley et al. (2009). The *Koopman operator* U is defined for a dynamical system

$$\boldsymbol{\xi}_{k+1} = \mathbf{f}\left(\boldsymbol{\xi}_k\right) \tag{1}$$

evolving on a finite-dimensional manifold M. (The dynamics need not be discrete in time, but we retain this form since we presume that the data to be analyzed are such.) It acts on scalar functions $g : M \to \mathbb{R}$ or $\mathbb{C}$ according to

$$Ug\left(\boldsymbol{\xi}\right) \triangleq g\left(\mathbf{f}\left(\boldsymbol{\xi}\right)\right). \tag{2}$$

Although the underlying dynamics are nonlinear and finite-dimensional, the Koopman operator is a linear infinite-dimensional operator; the linearity property is a result of the definition in (2), not linearization. Given the eigendecomposition

$$U\phi_j\left(\boldsymbol{\xi}\right) = \lambda_j \phi_j\left(\boldsymbol{\xi}\right), \qquad j = 1, 2, \ldots \tag{3}$$

of U, we can generally express vector-valued "observables" $\mathbf{g} : M \to \mathbb{R}^n$ or $\mathbb{C}^n$ of our choice in terms of the Koopman eigenfunctions ϕ_j by

$$\mathbf{g}\left(\boldsymbol{\xi}\right) = \sum_{j=1}^{\infty} \phi_j\left(\boldsymbol{\xi}\right) \mathbf{v}_j, \tag{4}$$

where $\{\mathbf{v}_j\}_{j=1}^{\infty}$ is a set of vector coefficients called *Koopman modes*. (Here, we assume that each component of $\mathbf{g}$ lies within the span of the eigenfunctions.) Using (3) and the definition in (2), we can explicitly write

$$\mathbf{g}\left(\boldsymbol{\xi}_k\right) = \sum_{j=1}^{\infty} \lambda_j^k \phi_j\left(\boldsymbol{\xi}_0\right) \mathbf{v}_j. \tag{5}$$

The Koopman eigenvalues $\{\lambda_j\}_{j=1}^{\infty}$ therefore dictate the growth rate and frequency of each mode.[1] For simplicity, we will hereafter incorporate the constant $\phi_j\left(\boldsymbol{\xi}_0\right)$ into $\mathbf{v}_j$.

It is straightforward to show that, if the dynamics (1) are linear, with $\mathbf{f}(\boldsymbol{\xi}) = \mathbf{A}\boldsymbol{\xi}$, then the eigenvalues λ_j of $\mathbf{A}$ are also eigenvalues of the Koopman operator. Furthermore, if the observable is $\mathbf{g}(\boldsymbol{\xi}) = \boldsymbol{\xi}$, then the Koopman modes $\mathbf{v}_j$ are the corresponding eigenvectors of $\mathbf{A}$ (Rowley et al. 2009).

The practical idea behind the Koopman analysis is to collect a set of data $\{\boldsymbol{\xi}_k\}$, identify an observable $\mathbf{g}$ of interest from the data, and express the observable in terms of Koopman modes and eigenvalues. For example, an experimenter may wish to collect fluid flow data from a wind tunnel experiment, measuring three-component velocity fields in a spatial region of interest, and decompose the fields into eigenvalues and modes containing spatial structures.

The DMD algorithm approximates Koopman modes and eigenvalues from a finite set of data. The algorithm itself, described in Rowley et al. (2009) and Schmid (2010), is a variant of a standard Arnoldi method, and we summarize it here. Suppose that the vector observable is expressed by $\mathbf{g}\left(\boldsymbol{\xi}_k\right) = \mathbf{x}_k$. The algorithm operates on the real or complex data set $\{\mathbf{x}_k\}_{k=0}^{m}$ and

[1] In ergodic theory, the map $\mathbf{f}$ is assumed to be measure preserving, making the Koopman operator unitary (Petersen 1983). Here, we make no such assumption, allowing for the analysis of dynamics not lying on an attractor. As such, the growth rates given by $|\lambda_j|$ may differ from unity.

identifies complex Ritz values $\{\lambda_j\}_{j=1}^m$ and complex Ritz vectors $\{\mathbf{v}_j\}_{j=1}^m$ such that

$$\mathbf{x}_k = \sum_{j=1}^m \lambda_j^k \mathbf{v}_j, \qquad k = 0, \ldots, m-1 \tag{6a}$$

$$\mathbf{x}_m = \sum_{j=1}^m \lambda_j^m \mathbf{v}_j + \mathbf{r}, \qquad \mathbf{r} \perp \operatorname{span}(\mathbf{x}_0, \ldots, \mathbf{x}_{m-1}), \tag{6b}$$

provided that $\{\lambda_j\}_{j=1}^m$ are distinct. This algorithm approximates the data—potentially generated by a nonlinear process—as the trajectory of a linear system. The first m data vectors are exactly represented by Ritz values and vectors. Projecting the final data vector, however, results in a residual $\mathbf{r}$ that is orthogonal to all previous data vectors, and hence to all Ritz vectors as well.

The computation of DMD modes proceeds as follows. Let $\mathbf{K} \triangleq \begin{bmatrix}\mathbf{x}_0 & \cdots & \mathbf{x}_{m-1}\end{bmatrix}$ be a data matrix. Let $\mathbf{c} = \begin{bmatrix}c_0 & \cdots & c_{m-1}\end{bmatrix}^{\mathrm{T}}$ be a vector of coefficients that best constructs (in a least squares sense) the final data vector $\mathbf{x}_m$ as a linear combination of all previous data vectors. That is,

$$\mathbf{x}_m = \mathbf{K}\mathbf{c} + \mathbf{r}, \qquad \mathbf{r} \perp \operatorname{span}(\mathbf{x}_0, \ldots, \mathbf{x}_{m-1}). \tag{7}$$

One solution for $\mathbf{c}$ is given by

$$\mathbf{c} = \mathbf{K}^+ \mathbf{x}_m, \tag{8}$$

where $(\cdot)^+$ indicates the Moore-Penrose pseudoinverse. The solution for (7) is unique if and only if $\mathbf{K}$ has linearly independent columns; in this case, we can use the pseudoinverse expansion $\mathbf{K}^+ = (\mathbf{K}^*\mathbf{K})^{-1}\mathbf{K}^*$. Now construct the companion matrix

$$\mathbf{C} \triangleq \begin{bmatrix} 0 & 0 & \cdots & 0 & c_0 \\ 1 & 0 & \cdots & 0 & c_1 \\ 0 & 1 & & 0 & c_2 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & & 1 & c_{m-1} \end{bmatrix}. \tag{9}$$

One possible diagonalization of this matrix is

$$\mathbf{C} = \mathbf{T}^{-1}\mathbf{\Lambda}\mathbf{T}, \tag{10}$$

where $\mathbf{\Lambda}$ is a diagonal matrix with $\lambda_1, \ldots, \lambda_m$ on the diagonal, $\mathbf{T}$ is the Vandermonde matrix defined by $T_{ij} \triangleq \lambda_i^{j-1}$, and $\{\lambda_j\}_{j=1}^m$ are distinct. The matrix $\mathbf{\Lambda}$ yields the Ritz values. Finally, the Ritz vectors are given by $\mathbf{V} \triangleq \begin{bmatrix}\mathbf{v}_1 & \cdots & \mathbf{v}_m\end{bmatrix}$, where

$$\mathbf{V} = \mathbf{K}\mathbf{T}^{-1}. \tag{11}$$

To understand how this algorithm yields the decomposition in (6), consider that an index-shifted data matrix $\mathbf{K}_* \triangleq \begin{bmatrix}\mathbf{x}_1 & \cdots & \mathbf{x}_m\end{bmatrix}$ can be written

$$\mathbf{K}_* = \mathbf{K}\mathbf{C} + \mathbf{r}\mathbf{e}^{\mathrm{T}} \tag{12a}$$

$$= \mathbf{K}\mathbf{T}^{-1}\mathbf{\Lambda}\mathbf{T} + \mathbf{r}\mathbf{e}^{\mathrm{T}}, \tag{12b}$$

if $\mathbf{e} \triangleq \begin{bmatrix}0 & \cdots & 0 & 1\end{bmatrix}^{\mathrm{T}}$. Finally, using (11), we obtain

$$\mathbf{K} = \mathbf{V}\mathbf{T} \tag{13a}$$

$$\mathbf{K}_* = \mathbf{V}\mathbf{\Lambda}\mathbf{T} + \mathbf{r}\mathbf{e}^{\mathrm{T}}, \tag{13b}$$

which together are equivalent to (6).

The algorithm above is analytically correct, but may be ill-conditioned in practice. This may especially be the case when using noisy experimental data. Schmid (2010) recommends the following alternate algorithm, which is exact when $\mathbf{K}$ is full-rank. Let $\mathbf{K} = \mathbf{U}\mathbf{\Sigma}\mathbf{W}^*$ be the economy-sized singular value decomposition (SVD) of $\mathbf{K}$. Solve the diagonalization problem

$$\mathbf{U}^*\mathbf{K}_*\mathbf{W}\mathbf{\Sigma}^{-1} = \mathbf{Y}\mathbf{\Lambda}\mathbf{Y}^{-1} \tag{14}$$

and set $\mathbf{V} = \mathbf{U}\mathbf{Y}$; $\mathbf{\Lambda}$ is as above. Finally, scale each column of $\mathbf{V}$ by the appropriate complex scalar so that (6a) is satisfied.

For data sets with $\dim(\mathbf{x}_k) \gg m$, we recommend a method of snapshots approach. In fluid mechanics, for instance, $\mathbf{K}$ may (very roughly) have a row dimension $\mathcal{O}\left(10^5\right)$ to $\mathcal{O}\left(10^{10}\right)$ and a column dimension $\mathcal{O}\left(10^1\right)$ to $\mathcal{O}\left(10^3\right)$. In this case, $\mathbf{K}^*\mathbf{K}$ is much smaller than $\mathbf{K}$. The method of snapshots can be carried out with lower memory requirements, and may be faster to compute. Evaluate the matrix $\mathbf{K}^*\mathbf{K}$, and perform the diagonalization

$$\mathbf{K}^*\mathbf{K} = \mathbf{W}\mathbf{\Sigma}^2\mathbf{W}^*. \tag{15}$$

The diagonal matrix $\mathbf{\Sigma}$ with positive entries and the unitary matrix $\mathbf{W}$ are as above. Complete the algorithm of Schmid above, starting at (14), with the substitution $\mathbf{U} = \mathbf{K}\mathbf{W}\mathbf{\Sigma}^{-1}$.

A key advantage of the SVD and method of snapshots approaches comes to light when $\mathbf{K}$ is rank-deficient or nearly so. In this case, the parts of $\mathbf{U}$, $\mathbf{\Sigma}$, and $\mathbf{W}$ corresponding to small singular values can be truncated to approximate (6) with a smaller number of modes (Schmid 2010).

3 Uniqueness of dynamic mode decomposition

Rowley et al. (2009) proved by construction that for a set of real or complex data $\{\mathbf{x}_k\}_{k=0}^m$, there exist complex $\{\lambda_j\}_{j=1}^m$ and $\{\mathbf{v}_j\}_{j=1}^m$ such that (6) is satisfied, provided that $\{\lambda_j\}_{j=1}^m$ are distinct. Here, we extend this result by showing that the decomposition is unique.

Theorem 1. *The choice of $\{\lambda_j\}_{j=1}^m$ and $\{\mathbf{v}_j\}_{j=1}^m$ in (6) is unique up to a reordering in j, if and only if $\{\mathbf{x}_k\}_{k=0}^{m-1}$ are linearly independent and $\{\lambda_j\}_{j=1}^m$ are distinct.*

Proof. We begin with the backward proof. Since λ_j are distinct, the Vandermonde matrix $\mathbf{T}$ is full-rank. Recall that (6a) is equivalent to $\mathbf{K} = \mathbf{V}\mathbf{T}$, and that we assume $\mathbf{K}$ is full-rank. It follows that $\mathbf{V}$ is full-rank, and that $\{\mathbf{x}_k\}_{k=0}^{m-1}$ and $\{\mathbf{v}_j\}_{j=1}^m$ must each be linearly independent bases of the same space. Equations (6b, 7) therefore imply that $\mathbf{x}_m$ can be decomposed into two additive parts,

$$\sum_{j=1}^m \lambda_j^m \mathbf{v}_j = \mathbf{K}\mathbf{c} \in \operatorname{span}(\mathbf{x}_0, \ldots, \mathbf{x}_{m-1}) \tag{16}$$

and $\mathbf{r} \perp \operatorname{span}(\mathbf{x}_0, \ldots, \mathbf{x}_{m-1})$.

Let us rewrite (16) as

$$\mathbf{V}\begin{bmatrix}\lambda_1^m \\ \vdots \\ \lambda_m^m\end{bmatrix} = \mathbf{K}\mathbf{c} \tag{17a}$$

$$= \mathbf{V}\mathbf{T}\mathbf{c}. \tag{17b}$$

Given that the columns of $\mathbf{V}$ are linearly independent, this then reduces to

$$\begin{bmatrix} \lambda_1^m \\ \vdots \\ \lambda_m^m \end{bmatrix} = \mathbf{T}\mathbf{c}, \tag{18}$$

from which we extract m scalar equations

$$\begin{aligned} \lambda_1^m &= c_0 + c_1\lambda_1 + c_2\lambda_1^2 + \cdots + c_{m-1}\lambda_1^{m-1} \\ &\vdots \\ \lambda_m^m &= c_0 + c_1\lambda_m + c_2\lambda_m^2 + \cdots + c_{m-1}\lambda_m^{m-1}. \end{aligned} \tag{19}$$

Therefore, the m distinct values of λ_j are precisely the m roots of the polynomial

$$p(\lambda; \mathbf{c}) \triangleq \lambda^m - \sum_{k=0}^{m-1} c_k \lambda^k. \tag{20}$$

(By extension, we may also write $p(\lambda; \mathbf{c}) = \prod_{j=1}^m (\lambda - \lambda_j)$; this will be used later.) Note from (8) that for linearly independent $\{\mathbf{x}_k\}_{k=0}^{m-1}$, $\{c_k\}_{k=0}^{m-1}$ exists and is unique. Thus, the m roots $\{\lambda_j\}_{j=1}^m$ of $p(\lambda; \mathbf{c})$ must be unique as well, up to a reordering in j. Finally, the vectors $\{\mathbf{v}_j\}_{j=1}^m$ must be unique, as a result of (11).

Next, we show by contradiction that if $\{\mathbf{x}_k\}_{k=0}^{m-1}$ are linearly dependent or $\{\lambda_j\}_{j=1}^m$ are not distinct, then the choice of $\{\lambda_j\}_{j=1}^m$ and $\{\mathbf{v}_j\}_{j=1}^m$ is not unique, even beyond a reordering in j. The latter condition is straightforward to prove—suppose that $\lambda_q = \lambda_r$ for some $q \neq r$. Then for any arbitrary $\mathbf{v}_* \in \mathbb{C}^n$, $\mathbf{v}_q$ may be replaced with $\mathbf{v}_q + \mathbf{v}_*$, and $\mathbf{v}_r$ with $\mathbf{v}_r - \mathbf{v}_*$, and (6) will still hold.

Suppose now that $\{\mathbf{x}_k\}_{k=0}^{m-1}$ are linearly dependent, so $\operatorname{rank}(\mathbf{K}) < m$. We maintain the condition that the λ_j be distinct; otherwise we refer back to the statement just proven. Thus, the Vandermonde matrix $\mathbf{T}$ is full-rank, so $\operatorname{rank}(\mathbf{V}) = \operatorname{rank}(\mathbf{K}) < m$; furthermore, $\operatorname{span}(\mathbf{x}_0, \ldots, \mathbf{x}_{m-1}) = \operatorname{span}(\mathbf{v}_1, \ldots, \mathbf{v}_m)$.

The remainder of this proof follows similarly to the backward proof. One difference, however, is that $\mathbf{c}$ must be *non-unique* because $\mathbf{K}$ does not have full column rank. To be precise, $\dim(\operatorname{Ker}(\mathbf{K})) > 0$; if some $\mathbf{c}$ satisfies (7), then $\mathbf{c} + \mathbf{c}_*$ does as well, for any $\mathbf{c}_* \in \operatorname{Ker}(\mathbf{K})$. In addition, (18, 19) are now sufficient but not necessary for (17). In other words, the m solutions to the polynomial equation $p(\lambda; \mathbf{c}) = 0$, along with $\mathbf{V} = \mathbf{K}\mathbf{T}^{-1}$, are sufficient for satisfying (6).

Since $\mathbf{c}$ is not unique, there must exist different coefficient sets $\{c_k\}_{k=0}^{m-1}$ and $\{d_k\}_{k=0}^{m-1}$ such that

$$\lambda^m - \sum_{k=0}^{m-1} c_k \lambda^k = \prod_{j=1}^m (\lambda - \lambda_j) = 0 \tag{21a}$$

$$\mu^m - \sum_{k=0}^{m-1} d_k \mu^k = \prod_{j=1}^m (\mu - \mu_j) = 0 \tag{21b}$$

produce roots $\{\lambda_j\}_{j=1}^m$ and $\{\mu_j\}_{j=1}^m$ each satisfying (6). If $\{\lambda_j\}_{j=1}^m$ and $\{\mu_j\}_{j=1}^m$ are identical, then according to the above relations, $\{c_k\}_{k=0}^{m-1}$ and $\{d_k\}_{k=0}^{m-1}$ must also be identical. Since the coefficient sets are different, however, there must exist distinct sets $\{\lambda_j\}_{j=1}^m$ and $\{\mu_j\}_{j=1}^m$ satisfying (6). □

Remark. The polynomial $p(\lambda; \mathbf{c})$, defined in (20), is precisely the characteristic polynomial of the companion matrix in (9) used to construct the DMD. Thus, as in (10), the scalars $\{\lambda_j\}_{j=1}^m$ are the eigenvalues of $\mathbf{C}$.

4 Boundary condition handling

4.1 Overview

An important application of modal decompositions, such as POD or DMD, is to obtain low-dimensional approximations of partial differential equations. When used in this setting, the governing dynamics (1) become infinite dimensional, although in practice, one often uses approximate numerical solutions (e.g., using finite-difference or spectral methods), and can treat these as systems with large, but finite, dimension. Modal decompositions such as POD and DMD may then be used to obtain further approximations of surprisingly low dimension, by expanding the solution in terms of a linear combination of modes (Holmes et al. 1996).

Of course, partial differential equations differ fundamentally from initial value ordinary differential equations and maps in that they require boundary conditions. Subtleties of imposing boundary conditions on such low-dimensional models are frequently overlooked, for instance as pointed out in Noack et al. (2005), but in many cases, the situation is straightforward. For instance, in a fluid flow, often the boundaries of the domain are at a physical boundary (a wall), where velocity is zero, and a simple Dirichlet boundary condition is appropriate. Other cases such as inflow or outflow conditions can be more difficult, but in practice, often a Dirichlet or Neumann boundary condition suffices for constructing low-dimensional models. When POD or DMD modes are used in these situations, careful attention must be given to ensure that the boundary conditions are satisfied, at least approximately, by the low-dimensional models.

In this section, we consider the case where the dynamics (1) represent an approximation of a partial differential equation, discretized in both time and space, and the observable $\mathbf{x} = \mathbf{g}(\boldsymbol{\xi})$ consists of the dependent variables in some spatial region Ω. This is a common situation for fluid flows, since flow information is typically known only in a limited region of space, whether data is obtained from experiments or from numerical simulations. We consider boundary conditions of the form

$$B\left(\mathbf{x}\right) = \gamma, \tag{22}$$

where B maps the observed variables $\mathbf{x}$ to certain values on the boundary $\partial\Omega$. Suppose the boundary condition also satisfies the linearity condition

$$B\left(\alpha_1 \mathbf{x}_1 + \alpha_2 \mathbf{x}_2\right) = \alpha_1 B\left(\mathbf{x}_1\right) + \alpha_2 B\left(\mathbf{x}_2\right). \tag{23}$$

Given a set of modes $\{\mathbf{v}_j\}_{j=1}^m$, one may construct a low-dimensional approximation of the solution by expanding $\mathbf{x}$ as a linear combination of modes. Using $\mathbf{x}(k)$ to denote $\mathbf{x}_k$ from the previous section, we have the expansion

$$\mathbf{x}(k) = \mathbf{x}_b + \sum_{j=1}^{m} a_j(k)\mathbf{v}_j, \tag{24}$$

where the constant offset $\mathbf{x}_b$ is commonly referred to as a "base flow" in the context of fluids. It is then convenient to require that the base flow itself satisfy the boundary conditions:

$$B(\mathbf{x}_b) = \gamma. \tag{25}$$

Equations (22–25) imply that for a set $\{\mathbf{x}(k)\}_{k=0}^m$ of observations each satisfying (22),

$$B(\mathbf{x}(k) - \mathbf{x}_b) = 0, \qquad k = 0, \ldots, m. \tag{26}$$

That is, if each solution $\mathbf{x}(k)$ satisfies the boundary conditions, then the subtraction of the base flow from each solution attains homogeneous boundary conditions. As a result, modes decomposed from the base-flow-subtracted solutions also acquire homogeneous boundary conditions.

Therefore, given a set of modes $\{\mathbf{v}_j\}_{j=1}^m$ constructed from base-flow-subtracted data, then as long as the base flow $\mathbf{x}_b$ satisfies (25), the expansion (24) will also satisfy the boundary conditions (22), for any choice of $a_j(k)$. When employing DMD modes computed *without* first subtracting a base flow, it is generally not possible to satisfy boundary conditions. The modes will typically have nonzero boundary conditions, and any nonzero values along those mode boundaries will be scaled by $a_j(k)$.

In the application of POD, it is common to use the mean of a data set as the base flow (e.g., Noack et al. 2003). We show in the following sections, however, that the use of the data mean as the base flow exactly reduces DMD to the temporal DFT, and that this is typically undesirable. On the other hand, using an equilibrium point as the base flow leads to better-behaved decompositions.

4.2 Mean subtraction

The data mean we describe here is defined by

$$\bar{\mathbf{x}} \triangleq \frac{1}{m+1}\sum_{k=0}^{m}\mathbf{x}_k, \tag{27a}$$

and the mean-subtracted data by

$$\mathbf{x}'_k \triangleq \mathbf{x}_k - \bar{\mathbf{x}}, \qquad k = 0,\ldots,m. \tag{27b}$$

We note briefly that the mean must include the vector $\mathbf{x}_m$. Otherwise, the mean-subtracted data would be given by the rank-deficient matrix $\begin{bmatrix}\mathbf{x}_0 - (1/m)\sum_{k=0}^{m-1}\mathbf{x}_k & \cdots & \mathbf{x}_{m-1} - (1/m)\sum_{k=0}^{m-1}\mathbf{x}_k\end{bmatrix}$; according to Section 3, DMD would not be unique. Therefore, let $\mathbf{K}' \triangleq \begin{bmatrix}\mathbf{x}'_0 & \cdots & \mathbf{x}'_{m-1}\end{bmatrix}$ be the mean-subtracted data matrix, with the mean given by (27a). Then

$$\mathbf{K}' = \begin{bmatrix}\mathbf{x}_0 & \cdots & \mathbf{x}_m\end{bmatrix}\begin{bmatrix}1-\frac{1}{m+1} & & -\frac{1}{m+1}\\ & \ddots & \\ -\frac{1}{m+1} & & 1-\frac{1}{m+1}\\ -\frac{1}{m+1} & \cdots & -\frac{1}{m+1}\end{bmatrix}. \tag{28}$$

The $(m+1)\times m$ rightmost matrix has a full column rank m. If $\begin{bmatrix}\mathbf{x}_0 & \cdots & \mathbf{x}_m\end{bmatrix}$ also has a full column rank, then $\mathbf{K}'$ does as well.

Subtracting $\bar{\mathbf{x}}$ from data vectors has the unexpected result of determining the Ritz values in a way that is completely independent of the data vectors' content. If $\mathbf{c} = \begin{bmatrix}-1 & \cdots & -1\end{bmatrix}^{\mathrm{T}}$, then $\mathbf{K}'\mathbf{c}$ reduces precisely to $\mathbf{x}'_m$. Therefore, the final snapshot $\mathbf{x}'_m$ is perfectly represented as a linear combination of the previous mean-subtracted snapshots in $\mathbf{K}'$. With this value of $\mathbf{c}$, an eigendecomposition of the corresponding companion matrix $\mathbf{C}$ yields Ritz values that satisfy

$$0 = \sum_{k=0}^{m}\lambda_j^k \tag{29a}$$

$$= \frac{1-\lambda_j^{m+1}}{1-\lambda_j}. \tag{29b}$$

(This is also evident from (20)). From here, it is easy to see that the Ritz values are the $(m+1)^{\text{th}}$ roots of unity, excluding unity itself. We write this explicitly as

$$\lambda_j = \exp\left(\frac{2\pi i j}{m+1}\right), \qquad j = 1,\ldots,m. \tag{30}$$

Thus, the Ritz values are completely independent of the data; they can be found with only one parameter, m. Inserting this result into (6), we find

$$\mathbf{x}'_k = \sum_{j=1}^{m} \exp\left(\frac{2\pi ijk}{m+1}\right) \mathbf{v}_j, \qquad k = 0, \ldots, m. \tag{31}$$

Note that the residual typically found on $k = m$ is absent because of the perfect reconstruction of that snapshot.

4.3 Equivalence with temporal DFT and harmonic average

A temporal DFT of the mean-subtracted data $\{\mathbf{x}'_k\}_{k=0}^m$ is given by

$$\{\hat{\mathbf{x}}_j\}_{j=0}^m = \mathcal{F}\left[\{\mathbf{x}'_k\}_{k=0}^m\right] \triangleq \left\{\frac{1}{m+1}\sum_{k=0}^{m} \exp\left(-\frac{2\pi ijk}{m+1}\right) \mathbf{x}'_k\right\}_{j=0}^m, \tag{32a}$$

and its inverse transform by

$$\{\mathbf{x}'_k\}_{k=0}^m = \mathcal{F}^{-1}\left[\{\hat{\mathbf{x}}_j\}_{j=0}^m\right] \triangleq \left\{\sum_{j=0}^{m} \exp\left(\frac{2\pi ijk}{m+1}\right) \hat{\mathbf{x}}_j\right\}_{k=0}^m. \tag{32b}$$

(In this representation, we move the scaling term $1/(m+1)$ typically found on the inverse transform to the forward transform.) Note that (32a) implies $\hat{\mathbf{x}}_0 = \mathbf{0}$, since the set $\{\mathbf{x}'_k\}_{k=0}^m$ has a zero mean. Thus, (32b) reduces to

$$\mathbf{x}'_k = \sum_{j=1}^{m} \exp\left(\frac{2\pi ijk}{m+1}\right) \hat{\mathbf{x}}_j \tag{33}$$

for $k = 0, \ldots, m$. The decompositions in (31, 33) are identical if we let $\mathbf{v}_j = \hat{\mathbf{x}}_j$ for $j = 1, \ldots, m$. Appealing to the uniqueness of DMD shown in Section 3, we conclude that if a linearly independent data set has a zero mean, then DMD is exactly equivalent to a temporal DFT. This is perhaps surprising, since the DMD algorithm and its derivation are largely unrelated to those of the temporal DFT (see Section 2, Rowley et al. 2009, and Schmid 2010). Whereas the temporal DFT model is based on a superposition of oscillating modes, the DMD model is based on the linear combination of eigenvectors that grow or decay according to their eigenvalues.

Mean-subtracted DMD and the temporal DFT are also equivalent to harmonic averaging, for the correct choices of averaging frequencies. Harmonic averaging is a concept that was applied to Koopman analysis (Mezić and Banaszuk 2004; Mezić 2005) before the DMD algorithms of Rowley et al. (2009) and Schmid (2010) were developed. Suppose we pick an oscillation rate $\omega_j = 2\pi j/((m+1)\Delta t)$ for $j \in \{1, \ldots, m\}$, along with time values $t_k = k\Delta t$, where Δt is the time interval between snapshots and $k = 0, \ldots, m$. Multiplying $\mathbf{x}'_k$ by this discrete-time oscillation $\exp(-i\omega_j t_k)$ and computing its average, we recover (32a). Hence, this choice of averaging frequencies exactly reproduces the temporal DFT. Thus, it is also equivalent to DMD with mean-subtracted data.

We also note briefly that if a data set has a nonzero mean, then a temporal DFT or a harmonic average yields $\hat{\mathbf{x}}_0 = \bar{\mathbf{x}}$. This is akin to a DMD analysis on mean-subtracted data, in which we separately account for the data mean.

4.4 Implications of mean subtraction

To make a clear distinction between DMD and mean-subtracted DMD, we will hereafter refer to the latter as a temporal DFT.

A key limitation of the temporal DFT is that its Ritz values depend only on the number of data points and not the data content; see (30). Although a Ritz vector energy would still dictate the degree to which that particular DFT mode is present in the data, the decomposition is only capable of outputting a predetermined set of frequencies. To ensure that a particular frequency is properly captured, the data must cover an integer number of corresponding periods. If multiple frequencies are of interest, then this constraint may be prohibitive, especially if the frequencies are unknown or are not related by a simple rational number.

In addition, this constraint implies that the longest period the temporal DFT can capture is the time span of the data set. We will see in the discussion of the cylinder flow example (Section 6.5) that DMD and a variant we will introduce in Section 5 are not subject to this limitation. Although DMD and its variant are still subject to the Nyquist frequency constraint, there is no theoretical lower bound on the frequencies they can compute.

Secondly, the temporal DFT is wholly incapable of determining modal growth rates. Equation (30) implies that all Ritz values have unit magnitude. This—like the predetermined mode frequencies—is an artifact of the mean subtraction, and is independent of the dynamics underlying the data. Dynamic mode decomposition is predicated on the assumption that meaningful information about a trajectory can be inferred from related linear dynamics. When the mean is removed, important dynamical information is stripped away, and the modal decomposition necessarily yields a periodic trajectory.

Finally, we stress the well-known fact that a temporal DFT on non-periodic data produces a slow decay in the modal energy. We also demonstrate this on the cylinder flow example in Section 6.5. In this example, many modes need to be retained to reproduce non-periodic data correctly.

4.5 Subtracting an equilibrium

As an alternative to subtracting the mean from a given dataset, one can consider subtracting the observable corresponding to an *equilibrium point* of the dynamics. For instance, if the dynamics are given by (1), with linear boundary conditions (22), suppose there is an equilibrium solution $\boldsymbol{\xi}_e$ that satisfies $\mathbf{f}(\boldsymbol{\xi}_e) = \boldsymbol{\xi}_e$ along with the boundary conditions $B(\mathbf{g}(\boldsymbol{\xi}_e)) = \gamma$. This is often the case for fluid flows, for instance, where the equilibrium solution is a steady (laminar) flow, while the observed data consists of snapshots of an unsteady (possibly turbulent) flow that satisfies the same inflow and wall boundary conditions.

If we choose the "base flow" $\mathbf{x}_b$ to be the observable corresponding to this equilibrium (i.e., $\mathbf{x}_b = \mathbf{x}_e = \mathbf{g}(\boldsymbol{\xi}_e)$), then the base flow satisfies the boundary conditions (25). Thus, the modes computed from base-flow-subtracted data will satisfy homogeneous boundary conditions, just as in the case where the base flow is the mean of the data. However, using the equilibrium ($\mathbf{x}_b = \mathbf{x}_e$) appears to have significant advantages over using the mean ($\mathbf{x}_b = \bar{\mathbf{x}}$, as in the previous section), mainly because then the DMD algorithm no longer reduces to a traditional DFT, and we can once again obtain frequency and growth rate information from the data. Subtracting the equilibrium can be viewed as merely translating the coordinate system so that the equilibrium is at the origin. As we will see in Section 6.6, this shift does not appear to change the DMD spectrum and modes in a significant way.

5 Optimized dynamic mode decomposition

The DMD formulation in (6) can be thought of as a way to "curve-fit" $m+1$ nonlinearly-generated data vectors $\{\mathbf{x}_k\}_{k=0}^{m}$ to a linear model of m modes. The model achieves perfect accuracy for $k = 0, \ldots, m-1$, as in (6a), and minimizes the error at $k = m$, as in (6b). This model, however, has potential pitfalls. First, if the user provides $m+1$ data points but requires fewer than m modes, then it is not entirely clear how to truncate the mode set—especially when the governing dynamics are not simple. This scenario is common in the construction of reduced-order models from data-based decompositions, where the desired number of modes is typically much less than the number of data vectors, and the dynamics are often quite complex.

Second, the DMD algorithm places a residual only on the final data vector $\mathbf{x}_m$, since the algorithm is built upon the reconstruction of $\mathbf{x}_m$ from $\{\mathbf{x}_k\}_{k=0}^{m-1}$. Numerical experiments show that DMD results are more sensitive to variations $\mathbf{x}_m$ than to variations in other data vectors. This is because the Ritz values are the eigenvalues of the companion matrix $\mathbf{C}$ (9), which dictates this reconstruction. The presence of noise in $\mathbf{x}_m$ could drastically change the contents and hence the eigenvalues of the companion matrix.

To address these issues, we propose a new, optimized DMD, in which we fit p points on a trajectory to a linear model of m modes. We must satisfy the inequality $m < p$, but m and p are otherwise arbitrary. We allow the linear model to contain a residual at each of the p points, but the overall residual is minimized for the values of m and p that we choose.

More formally, suppose that $\{\mathbf{x}_k\}_{k=0}^{p-1}$ is a set of real or complex vectors. Given $m < p$, we seek complex scalars $\{\lambda_j\}_{j=1}^{m}$ and vectors $\{\mathbf{v}_j\}_{j=1}^{m}$ such that

$$\mathbf{x}_k = \sum_{j=1}^{m} \lambda_j^k \mathbf{v}_j + \mathbf{r}_k, \qquad k = 0, \ldots, p-1, \tag{34a}$$

and

$$\Gamma \triangleq \sum_{k=0}^{p-1} \|\mathbf{r}_k\|_2^2 \tag{34b}$$

is minimized.

To construct the optimized DMD, define

$$\mathbf{K} \triangleq \begin{bmatrix} \mathbf{x}_0 & \cdots & \mathbf{x}_{p-1} \end{bmatrix} \tag{35a}$$

$$\mathbf{V} \triangleq \begin{bmatrix} \mathbf{v}_1 & \cdots & \mathbf{v}_m \end{bmatrix} \tag{35b}$$

$$\mathbf{T} \triangleq \begin{bmatrix} 1 & \lambda_1 & \lambda_1^2 & \cdots & \lambda_1^{p-1} \\ 1 & \lambda_2 & \lambda_2^2 & \cdots & \lambda_2^{p-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & \lambda_m & \lambda_m^2 & \cdots & \lambda_m^{p-1} \end{bmatrix} \tag{35c}$$

$$\mathbf{r} \triangleq \begin{bmatrix} \mathbf{r}_1 & \cdots & \mathbf{r}_p \end{bmatrix}. \tag{35d}$$

We seek $\mathbf{V}$ and the Vandermonde matrix $\mathbf{T}$ such that $\mathbf{K} = \mathbf{V}\mathbf{T} + \mathbf{R}$, and the squared Frobenius norm $\Gamma = \|\mathbf{R}\|_F^2$ is minimized. This can be viewed as a least-squares problem for $\mathbf{V}$. Of all the choices of $\mathbf{V}$ that minimize $\|\mathbf{R}\|_F^2$, the one with the smallest Frobenius norm is

$$\mathbf{V} = \mathbf{K}\mathbf{T}^{+}. \tag{36}$$

(If $\{\lambda_j\}_{j=1}^m$ are distinct and therefore $\mathbf{T}$ is full-rank, then $\mathbf{T}^+ = \mathbf{T}^* (\mathbf{T}\mathbf{T}^*)^{-1}$.) The residual is therefore given by $\mathbf{R} = \mathbf{K}(\mathbf{I} - \mathbf{T}^+\mathbf{T})$. Using the trace expansion of the Frobenius norm, we then find that

$$\Gamma = \operatorname{tr}(\mathbf{K}^*\mathbf{K}) - \operatorname{tr}\left(\mathbf{T}^+\mathbf{T}\mathbf{K}^*\mathbf{K}\right) \tag{37a}$$

$$= \|\mathbf{K}\|_F^2 - \left\|\mathbf{K}\mathbf{T}^+\mathbf{T}\right\|_F^2 . \tag{37b}$$

Although the second form is more intuitive, the first form is faster to compute, since the small matrix $\mathbf{K}^*\mathbf{K}$ can be precomputed and stored. The objective is to find a set of complex numbers $\{\lambda_j\}_{j=1}^m$ that minimizes (37); the modes are then given by (36).

At present, we do not have an analytic algorithm for computing optimized DMD. To find the choice of $\{\lambda_j\}_{j=1}^m$ that minimizes Γ, we employ a global optimization technique that combines simulated annealing and the Nelder-Mead simplex method (Press et al. 2007). At each iteration, the technique adds thermal fluctuations to the function evaluations at the points of the simplex. When a new point is tested, a thermal fluctuation is subtracted from the function evaluation; therefore, there is a calculable probability of accepting a worse point. As the annealing completes, the fluctuations vanish, and the technique reduces exactly to the Nelder-Mead method of local optimization. We also use the Broyden-Fletcher-Goldfarb-Shanno quasi-Newton iterator for purely local minimization. In this case, we apply the SVD-based DMD algorithm of Schmid (2010) and use select Ritz values as the initial condition for the iterator. The two iterators produce consistent results for $m = 3$, which we report in Section 6.4.

This minimization problem is similar to that of POD, where $\|\mathbf{K}\|_F^2 - \left\|\mathbf{V}(\mathbf{V}^*\mathbf{V})^{-1}\mathbf{V}^*\mathbf{K}\right\|_F^2$ is minimized with the requirement that $\mathbf{V}$ be unitary. In optimized DMD, $\mathbf{T}$ is not unitary, but rather obeys the structure given in (35c).

Finally, we present a property of optimized DMD as a theorem below.

Theorem 2. *Suppose the data $\{\mathbf{x}_k\}_{k=0}^{p-1}$ are linearly independent. Then the Ritz values $\{\lambda_j\}_{j=1}^m$ resulting from optimized DMD are distinct.*

Proof. This proof is motivated by the notion that an orthogonal projection onto the rows of a Vandermonde matrix attains the maximum norm only when the rows are distinct.

Suppose that $\{\lambda_j\}_{j=1}^m$ are not distinct. In particular, suppose this set has $q < m$ distinct elements $\{\mu_j\}_{j=1}^q$. Let $\{\mu_j\}_{j=q+1}^m$ be any set of distinct complex scalars *not* in $\{\lambda_j\}_{j=1}^m$, so that $\{\mu_j\}_{j=1}^m$ is a distinct set of complex scalars. Construct the complex conjugate Vandermonde matrices

$$\mathbf{T}^* = \begin{bmatrix} 1 & \cdots & 1 \\ \bar{\lambda}_1 & \cdots & \bar{\lambda}_m \\ \vdots & & \vdots \\ \bar{\lambda}_1^{p-1} & \cdots & \bar{\lambda}_m^{p-1} \end{bmatrix} \qquad \mathbf{S}^* \triangleq \begin{bmatrix} 1 & \cdots & 1 \\ \bar{\mu}_1 & \cdots & \bar{\mu}_m \\ \vdots & & \vdots \\ \bar{\mu}_1^{p-1} & \cdots & \bar{\mu}_m^{p-1} \end{bmatrix} \tag{38}$$

and the corresponding orthogonal projection matrices $\mathbf{P}_{\mathbf{T}^*} \triangleq \mathbf{T}^*(\mathbf{T}^*)^+$ and $\mathbf{P}_{\mathbf{S}^*} \triangleq \mathbf{S}^*(\mathbf{S}^*)^+$. Also, let $\{\mathbf{y}_k\}_{k=1}^n$ be the columns of $\mathbf{K}^*$, so that

$$\mathbf{K}^* = \begin{bmatrix} \mathbf{x}_0^* \\ \vdots \\ \mathbf{x}_{p-1}^* \end{bmatrix} = \begin{bmatrix} \mathbf{y}_1 & \cdots & \mathbf{y}_n \end{bmatrix} . \tag{39}$$

Thus, each vector $\mathbf{y}_k$ contains the time history of the data at a particular index in $\mathbf{x}$.

First, for $k = 1, \ldots, n$, consider the separation of $\mathbf{y}_k$ into its projection onto the range of $\mathbf{S}^*$ and the orthogonal complement,

$$\mathbf{y}_k = \mathbf{P}_{\mathbf{S}^*}\mathbf{y}_k + \boldsymbol{\zeta}_k, \qquad \boldsymbol{\zeta}_k \notin \operatorname{range}(\mathbf{S}^*). \tag{40}$$

Since $\operatorname{range}(\mathbf{T}*) \subset \operatorname{range}(\mathbf{S}^*)$, we deduce that $\mathbf{P}_{\mathbf{T}^*}\boldsymbol{\zeta}_k = \mathbf{0}$. Thus, applying to the above equation a projection onto the range of $\mathbf{T}^*$, we obtain $\mathbf{P}_{\mathbf{T}^*}\mathbf{y}_k = \mathbf{P}_{\mathbf{T}^*}\mathbf{P}_{\mathbf{S}^*}\mathbf{y}_k$. Using this fact to separate $\mathbf{P}_{\mathbf{S}^*}\mathbf{y}_k$ into its projection onto $\mathbf{T}^*$ and the orthogonal complement,

$$\mathbf{P}_{\mathbf{S}^*}\mathbf{y}_k = \mathbf{P}_{\mathbf{T}^*}\mathbf{y}_k + \boldsymbol{\eta}_k, \qquad \mathbf{P}_{\mathbf{T}^*}\mathbf{y}_k \perp \boldsymbol{\eta}_k. \tag{41}$$

Furthermore, the Pythagorean theorem states that

$$\|\mathbf{P}_{\mathbf{S}^*}\mathbf{y}_k\|_2^2 = \|\mathbf{P}_{\mathbf{T}^*}\mathbf{y}_k\|_2^2 + \|\boldsymbol{\eta}_k\|_2^2. \tag{42}$$

Define $\mathbf{H} = \begin{bmatrix}\boldsymbol{\eta}_1 & \cdots & \boldsymbol{\eta}_n\end{bmatrix}$; stack (41) and sum (42) to form

$$\mathbf{P}_{\mathbf{S}^*}\mathbf{K}^* = \mathbf{P}_{\mathbf{T}^*}\mathbf{K}^* + \mathbf{H} \tag{43a}$$

$$\|\mathbf{P}_{\mathbf{S}^*}\mathbf{K}^*\|_F^2 = \|\mathbf{P}_{\mathbf{T}^*}\mathbf{K}^*\|_F^2 + \|\mathbf{H}\|_F^2. \tag{43b}$$

The cost function in (37) can be written $\Gamma = \|\mathbf{K}^*\|_F^2 - \|\mathbf{P}_{\mathbf{T}^*}\mathbf{K}^*\|_F^2$. By the construction of optimized DMD, $\mathbf{T}^*$ is the Vandermonde matrix that maximizes $\|\mathbf{P}_{\mathbf{T}^*}\mathbf{K}^*\|_F^2$. That is, $\|\mathbf{P}_{\mathbf{T}^*}\mathbf{K}^*\|_F^2 \geq \|\mathbf{P}_{\mathbf{S}^*}\mathbf{K}^*\|_F^2$ since $\mathbf{S}^*$ is also a Vandermonde matrix, and therefore $\mathbf{H} = \mathbf{0}$. Equation (43a) reduces to $\mathbf{P}_{\mathbf{S}^*}\mathbf{K}^* = \mathbf{P}_{\mathbf{T}^*}\mathbf{K}^*$.

The rank property of Vandermonde matrices then yields

$$q = \operatorname{rank}(\mathbf{P}_{\mathbf{T}^*}) = \operatorname{rank}(\mathbf{P}_{\mathbf{T}^*}\mathbf{K}^*) \tag{44a}$$

$$m = \operatorname{rank}(\mathbf{P}_{\mathbf{S}^*}) = \operatorname{rank}(\mathbf{P}_{\mathbf{S}^*}\mathbf{K}^*); \tag{44b}$$

the right-most side comes from the fact that $\mathbf{K}^*$ has full row rank. We arrive at a contradiction: $\mathbf{P}_{\mathbf{S}^*}\mathbf{K}^*$ and $\mathbf{P}_{\mathbf{T}^*}\mathbf{K}^*$ are equal, yet they have different ranks. Therefore, $\{\lambda_j\}_{j=1}^m$ must be distinct. □

6 Application to a cylinder flow

6.1 Overview

The two-dimensional incompressible flow past a cylinder is a canonical problem in fluid mechanics. Its dynamics are relatively well-understood (see Williamson 1996), and real-life examples with similar dynamics can be found in everyday objects such as flagpoles, smokestacks, and even entire skyscrapers. A number of recent studies have focused specifically on dynamical modeling, modal analysis, and flow control of the cylinder wake; see Roussopoulos (1993), Gillies (1998), Noack et al. (2003), Giannetti and Luchini (2007), Marquet et al. (2008), and Tadmor et al. (2010) for a short selection.

The dynamics of the incompressible cylinder flow are governed by the continuity and Navier-Stokes equations,

$$\nabla \cdot \mathbf{u} = 0 \tag{45a}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla p}{\rho} + \nu \nabla^2 \mathbf{u}. \tag{45b}$$

We denote the horizontal (free-stream) direction by x and the vertical (normal) direction by y. The fluid velocity $\mathbf{u} = \begin{bmatrix} u_x(x,y,t) & u_y(x,y,t) \end{bmatrix}^{\mathrm{T}}$ and pressure $p(x,y,t)$ vary in space and time, and the fluid density ρ and kinematic viscosity ν are assumed constant. Denoting the horizontal unit vector by $\mathbf{e}_x$, the boundary conditions are

$$\mathbf{u} = U\mathbf{e}_x, \qquad x, y \to \pm\infty \tag{46a}$$

$$p = p_0, \qquad x, y \to \pm\infty \tag{46b}$$

$$\mathbf{u} = \mathbf{0}, \qquad \sqrt{x^2 + y^2} = \frac{D}{2}, \tag{46c}$$

where $U\mathbf{e}_x$ is the externally-imposed free-stream velocity, p_0 is the constant far-field pressure, and D is the cylinder diameter.

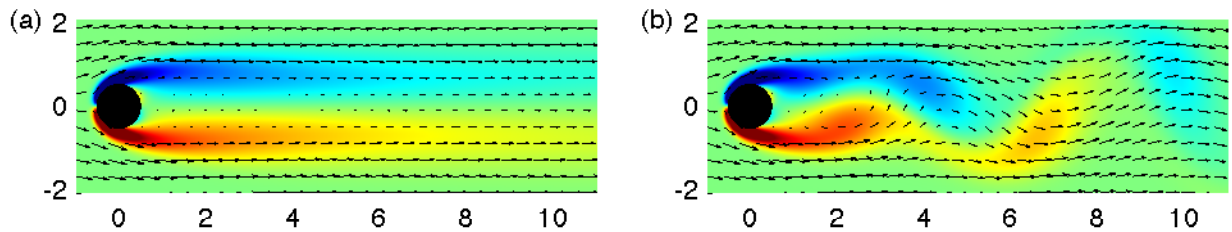

Figure 1: Color fields of vorticity (i.e., velocity rotation), with overlaid velocity vectors, for the cylinder flow at $Re = 60$. Clockwise vorticity is shown in blue, and counterclockwise in red. An exterior uniform flow from left to right is imposed. (a): The unstable equilibrium of the cylinder flow. (b): A snapshot from the limit cycle.

The qualities of the incompressible cylinder flow can be characterized by only one parameter, the Reynolds number $Re \triangleq UD/\nu$. At Reynolds numbers between 47 (Provansal et al. 1987) and roughly 194 (Williamson 1996), there exists one unstable equilibrium (Figure 1a), consisting of a long separation bubble with two counter-rotating vortices. There also exists one stable limit cycle (Figure 1b), known as a von Kármán vortex street, characterized by a periodic shedding of vortices from the top and bottom surfaces. To test our decomposition methods, we generate data from numerical simulations of a cylinder flow at $Re = 60$. We intentionally pick a Reynolds number close to 47—where the Hopf bifurcation occurs—so that dynamical parameters such as the shedding frequency do not change too drastically from the equilibrium to the limit cycle.

As the cylinder flow evolves from the unstable equilibrium to the limit cycle, three distinct regimes can be identified (see Figure 2). Near the equilibrium state, perturbations in the flow from the equilibrium are small, and the dynamics are approximately linear. The flow oscillates at a frequency f given by a nondimensional Strouhal number of $St \triangleq fD/U = 0.126$. The lift coefficient—defined by $C_L \triangleq 2F_L/\left(\rho U^2 D\right)$ for a lift force per depth F_L—increases in oscillation amplitude by $1.054^{\hat{t}}$, where $\hat{t} \triangleq tU/D$ is the nondimensional advection time. In the transient regime, the Strouhal number increases and the C_L peaks grow more slowly than exponentially. Finally, near the limit cycle, we find that $St = 0.141$ and that C_L peaks roughly reach their asymptotic values. This Strouhal number is consistent with findings by Roshko (1954).

6.2 Numerical method

The data are generated using a multi-domain immersed boundary method; see Taira and Colonius (2007) and Colonius and Taira (2008) for details. Each successively smaller domain covers half

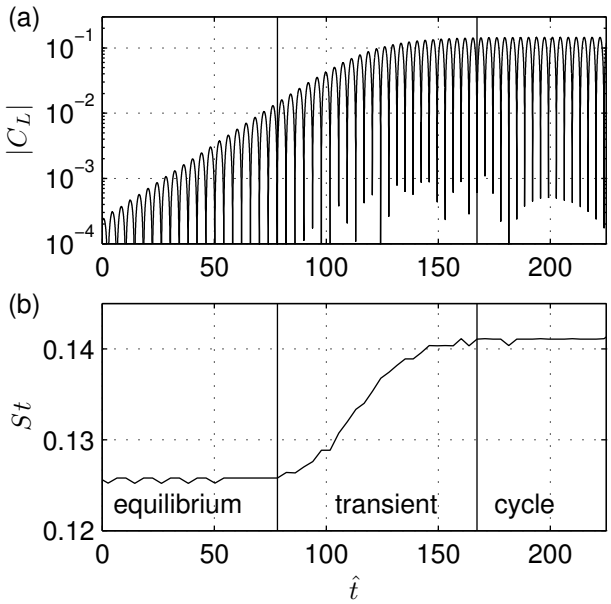

Figure 2: (a): The lift coefficient magnitude as the cylinder flow evolves from the unstable equilibrium to the limit cycle. (b): The corresponding instantaneous Strouhal number. The near-equilibrium, transient, and limit cycle regimes are indicated.

the streamwise and normal length, and samples the space twice as finely. The coarsest domain spans 48 diameters in the streamwise direction and 16 diameters in the normal direction. On the third and finest domain, shown in Figure 1, the computational domain spans 12 diameters in the streamwise direction and four diameters in the normal direction, with a resolution of 480 by 160 pixels. In this study, we only use velocity data from this innermost domain in DMD. We use a time step of $\Delta\hat{t} = 0.01875$, which satisfies the Courant-Friedrichs-Levy condition for the numerical scheme implemented. The two references above have verified the convergence of the scheme and agreement with analytic models.

We use selective frequency damping (SFD; see Åkervik et al. 2006) to solve for the unstable equilibrium. This solution is used as the initial condition of the immersed boundary solver. Because time derivatives are extremely small at the start of the fluid simulation, the computed equilibrium is likely close to the true equilibrium. The numerical error in the SFD solution causes the trajectory to depart from the equilibrium, and ultimately to approach the limit cycle. Only one simulation, the one shown in Figures 1 and 2, is used in this study.

6.3 Dynamic mode decomposition of the cylinder flow

From each of the three regimes shown in Figure 2, we compute DMD on x- and y-velocity data spanning approximately two periods of oscillation. Every 25th time step from the immersed boundary solver is retained, so the sampling time between successive data vectors is $\Delta\hat{t}_s = 25\Delta\hat{t} = 0.46875$. The data samples from the near-equilibrium, transient, and limit cycle regimes respectively cover $0 \leq \hat{t} < 15.94$, $95.63 \leq \hat{t} < 111.09$, and $210.94 \leq \hat{t} < 225.00$. To cover two periods, we respectively require 34, 33, and 30 snapshots. Although we feed an integer number of periods into DMD, our experience has been that DMD is generally insensitive to this condition. This is in contrast to other methods such as POD, which may yield nonsensical results when using non-integer periods of data.

The spectra of Ritz values are shown in Figure 3. The Strouhal number is computed by

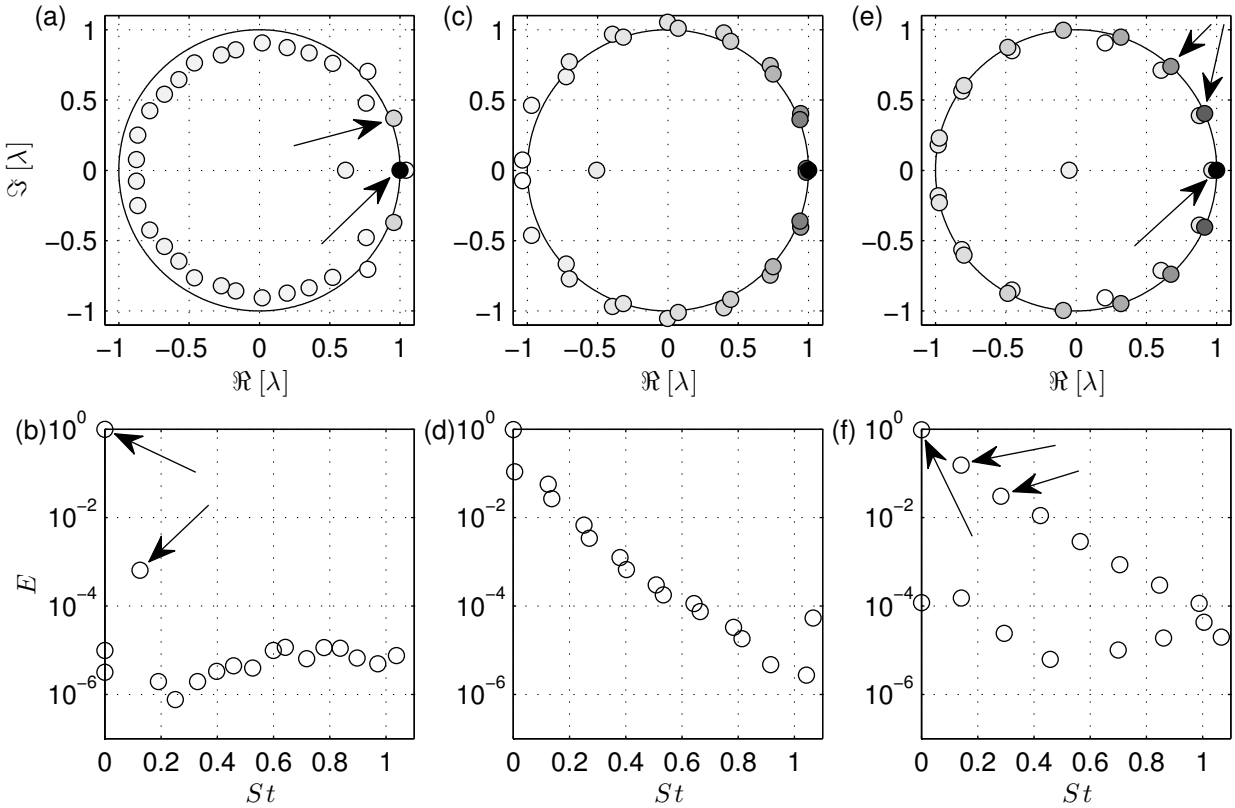

Figure 3: Ritz values λ computed from DMD on two periods of data. Top row: as the real and imaginary parts; the 2-norms of the corresponding Ritz vectors are indicated by the color of the circles, with darker colors indicating greater norms. Bottom row: as the normalized vector energy $E = \sqrt{m+1}\,\|\mathbf{v}\|_2 \,/\, \|[\mathbf{K}\ \mathbf{x}_m]\|_F$ versus the Strouhal number. Left column: near-equilibrium regime; middle column: transient regime; right column: limit cycle regime. The arrows indicate Ritz values for which the corresponding modes are plotted in Figures 4, 6, and 7.

$St = \arg(\lambda) \,/\, \left(2\pi\Delta\hat{t}_s\right)$, and vector energies $\|\mathbf{v}\|_2$ are normalized by the root-mean-square 2-norm of the data, $\left\|\left[\mathbf{K} \quad \mathbf{x}_m\right]\right\|_F / \sqrt{m+1}$. The most energetic mode typically has a Ritz value of $\lambda \approx 1$. This mode has the resemblance of (but is not exactly the same as) the mean of the data set; see Figures 4(a, b) and 6.

In the near-equilibrium regime (Figure 3a, b), where the dynamics are nearly linear, DMD easily recognizes the complex conjugate pair of eigenmodes primarily responsible for the oscillating exponential growth away from the unstable equilibrium. One mode from the pair is shown in Figure 4(c–e). The spatial oscillations are the result of instabilities that convect downstream. The magnitude of complex velocity mode, as shown in Figure 4(e), attains its maximum value far downstream of the cylinder. The qualitative nature of this unstable mode is consistent with previous results, such as the eigenmode analyses of Giannetti and Luchini (2007) and Marquet et al. (2008) of the linearized Navier-Stokes equations. An additional unstable pair of modes is visible in Figure 3(a, b); we posit that this is a result of numerical noise.

We momentarily skip the transient regime and discuss the limit cycle, as shown in Figure 3(e, f). The spectrum appears well-ordered; all but the least energetic modes lie very close to the unit circle. The mode corresponding to $\lambda = 1$ (Figure 6) exhibits a shorter tail than the equivalent mode from near the equilibrium. This is expected, and is the basis of "shift mode" analyses of

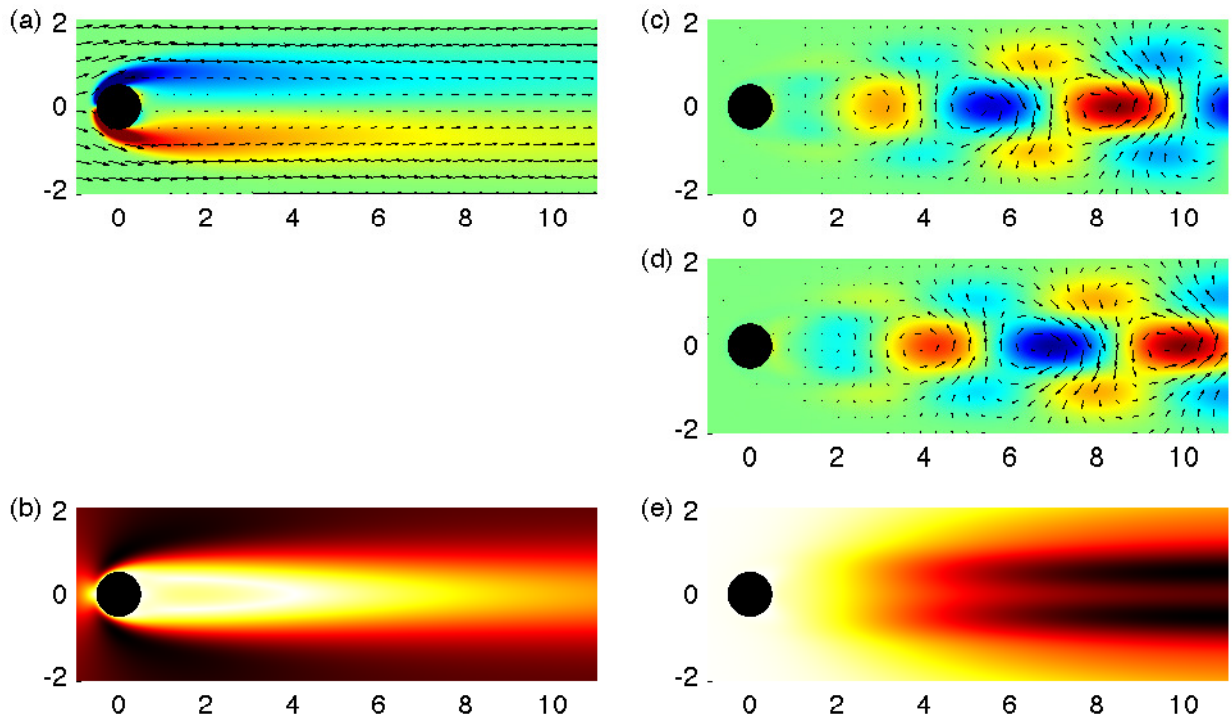

Figure 4: Modes $\mathbf{v}$ computed from DMD on two periods of data in the near-equilibrium regime. Left column: $\lambda = 1.000$; right column: $\lambda = 0.956 + 0.371i$ (see Figure 3a, b). Top row: the real part $\Re[\mathbf{v}]$; middle row: the imaginary part $\Im[\mathbf{v}]$. The modes are depicted as velocity vectors and vorticity color fields, with clockwise rotation in red and counterclockwise rotation in blue. Note that the mode corresponding to $\lambda = 1.000$ is real. Bottom row: the magnitude of the complex velocity, with zero shown in white and the highest magnitude shown in black.

bluff body wakes, as in Noack et al. (2003) and Tadmor et al. (2010). Furthermore, the oscillatory mode frequencies are nearly integral multiples of the base frequency $St = 0.141$. An aliasing effect can be seen in Figure 3(f) as the multiples exceed the Nyquist frequency of 1.067. The integral multiples suggest that the periodic orbit of the cylinder flow exhibits harmonic-like behavior. The first and second harmonics are respectively shown in Figure 7(a–c) and (d–f). We observe that the first harmonic is symmetric in vorticity about the horizontal axis, and the second harmonic is antisymmetric. Thus, the former corresponds to the large-scale convection of fluid structures in the wake, and the latter corresponds to the top-bottom oscillation visible in Figure 1(b).

We also observe that the DMD modes from the limit cycle match very closely with POD modes (see Holmes et al. 1996). The parallel is not exact, because oscillatory DMD modes are complex, whereas POD modes from real data are always real. Nevertheless, let us index DMD modes in order of decreasing vector energy, so that $\mathbf{v}_2$ is the DMD mode corresponding to $\lambda_2 = 0.915 + 0.404i$ (Figure 7a–c), and $\mathbf{v}_4$ is the mode corresponding to $\lambda_4 = 0.674 + 0.739i$ (Figure 7d–f). Similarly, let $\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_4$ be the first four POD modes from the same data set. We find that

$$\left| \frac{\langle \mathbf{v}_2, \boldsymbol{\phi}_1 + i\boldsymbol{\phi}_2 \rangle}{\|\mathbf{v}_2\|_2 \, \|\boldsymbol{\phi}_1 + i\boldsymbol{\phi}_2\|_2} \right|^2 = 0.9996 \tag{47a}$$

$$\left| \frac{\langle \mathbf{v}_4, \boldsymbol{\phi}_3 + i\boldsymbol{\phi}_4 \rangle}{\|\mathbf{v}_4\|_2 \, \|\boldsymbol{\phi}_3 + i\boldsymbol{\phi}_4\|_2} \right|^2 = 0.9999, \tag{47b}$$

where $\langle \mathbf{x}_j, \mathbf{x}_k \rangle \triangleq \mathbf{x}_k^* \mathbf{x}_j$ is an inner product. That is to say that in the limit cycle, POD modes are

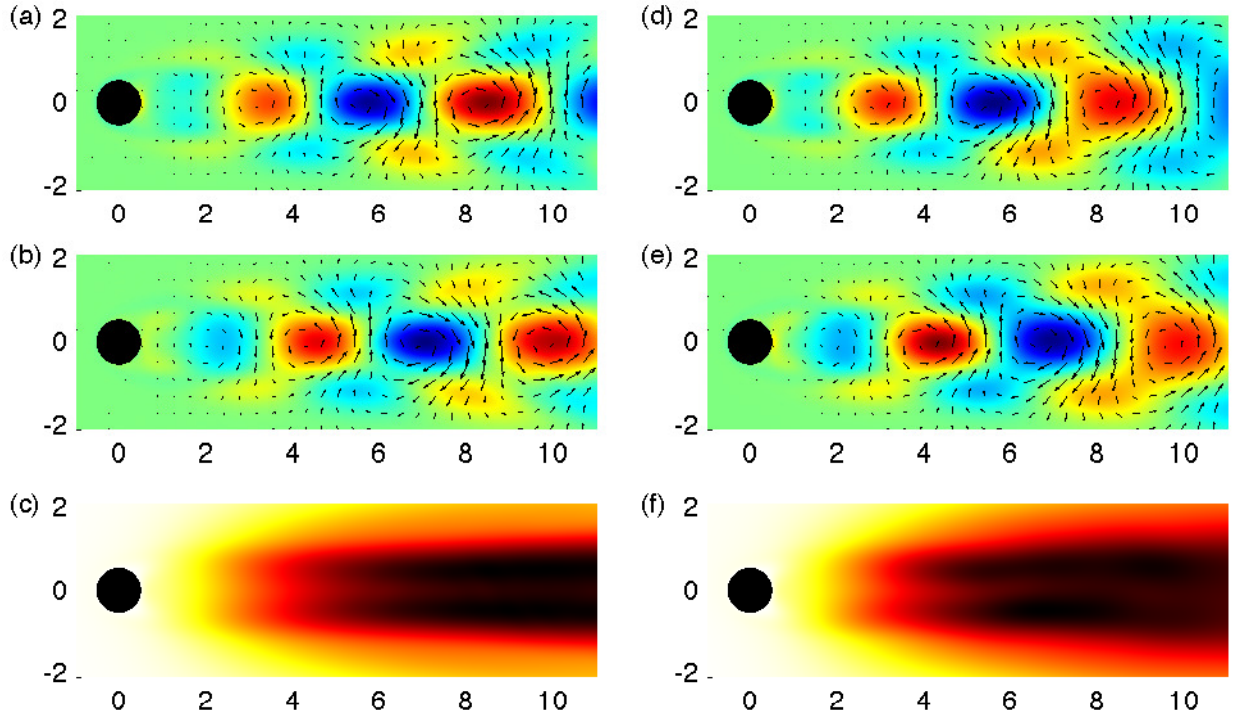

Figure 5: Modes $\mathbf{v}$ computed from two periods of data in the transient regime. Left column: DMD mode at $\lambda = 0.940 + 0.361i$; right column: optimized DMD mode ($m = 3$) at $\lambda = 0.940 + 0.378i$. (See Figure 10.) Top row: the real part $\Re[\mathbf{v}]$; middle row: the imaginary part $\Im[\mathbf{v}]$. Bottom row: the magnitude of the complex velocity.

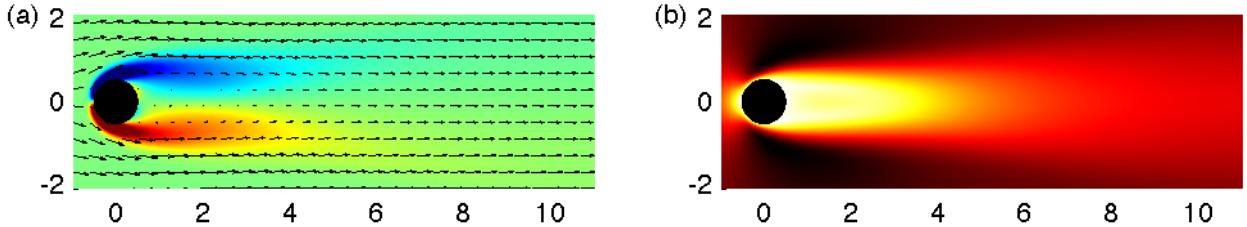

Figure 6: The DMD mode corresponding to $\lambda = 1.000$, from two periods of data in the limit cycle regime (see Figure 3e, f). (a): Velocity and vorticity color fields. (b): The magnitude of the velocity.

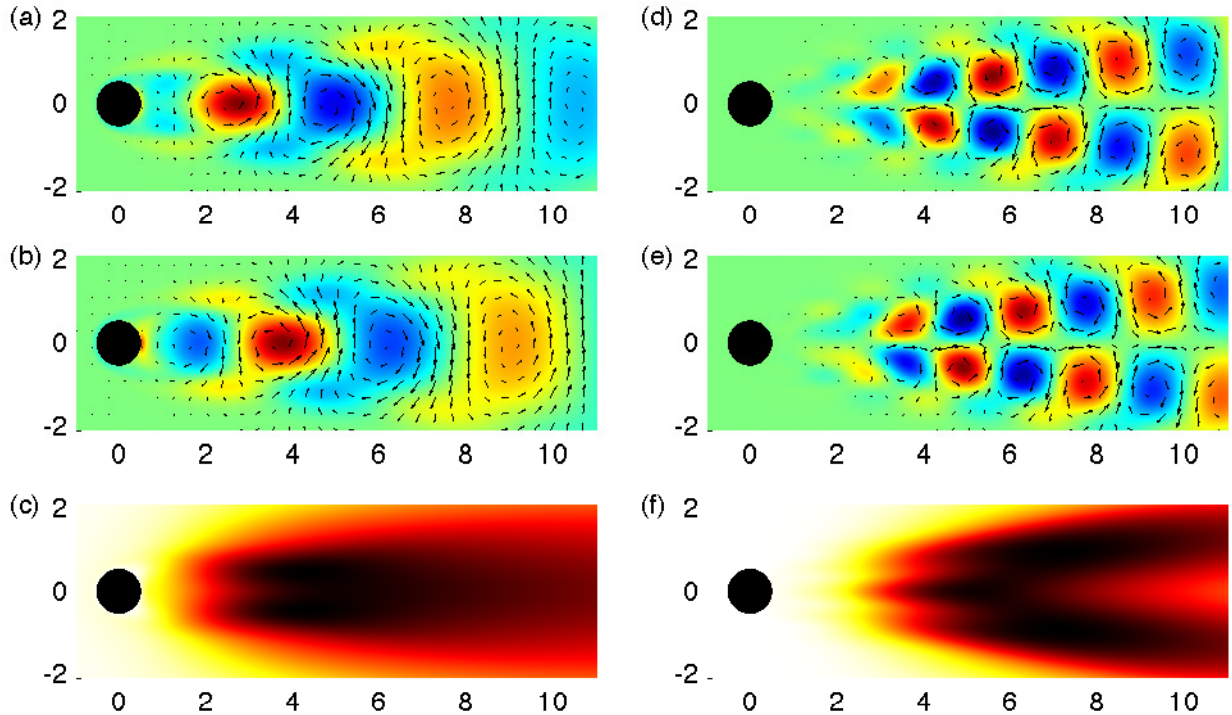

Figure 7: Modes $\mathbf{v}$ computed from DMD on two periods of data in the limit cycle regime. Left column: $\lambda = 0.915 + 0.404i$; right column: $\lambda = 0.674 + 0.739i$ (see Figure 3e, f). Top row: the real part $\Re[\mathbf{v}]$; middle row: the imaginary part $\Im[\mathbf{v}]$. Bottom row: the magnitude of the complex velocity.

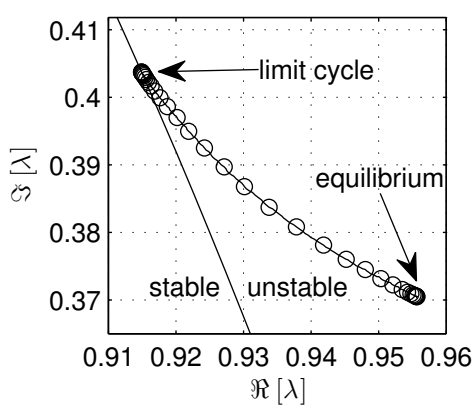

Figure 8: The movement of the primary unstable Ritz value near the equilibrium (see Figures 3a and 4c–e) to the primary oscillatory Ritz value near the limit cycle (see Figures 3e and 7a–c), as a window of 16 snapshots (i.e., roughly one period) traverses the cylinder data. The unit circle is shown as a solid curve, and the movement of the Ritz value is shown as a solid curve with circles.

extremely similar to the real and imaginary parts of DMD modes, up to a complex multiplicative factor.

Of the three regimes, the transient one (Figure 3c, d) is the most difficult to analyze. Rowley et al. (2009) showed that the analysis of the Koopman operator, on which DMD is based, yields elegant solutions for linear dynamics and periodic dynamics. The transient regime, however, is neither. Rather, it is the part of the highly nonlinear post-Hopf-bifurcation dynamics between the unstable equilibrium and the limit cycle. At $Re = 60$, a smooth "morphing" can be seen in the DMD results when a window of snapshots traverses the cylinder data from the equilibrium to the limit cycle (Figure 8). Therefore, the Ritz values and vectors are loosely between those of the near-equilibrium regime and those of the limit cycle. For instance, compare the most energetic oscillatory mode, shown in Figure 5(a–c), to Figure 4(c–e) and Figure 7(a–c).

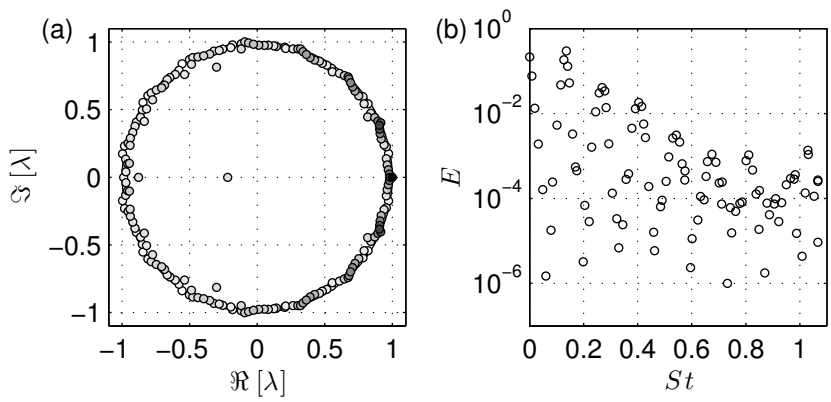

Figure 9: Ritz values λ computed from DMD on the entire transient regime. (a): As the real and imaginary parts. (b): As the normalized vector energy $E = \sqrt{m+1}\,\|\mathbf{v}\|_2 \,/\, \|[\mathbf{K}\ \mathbf{x}_m]\|_F$ versus the Strouhal number.

We obtain different results if the entire transient regime, given by $78.21 \leq \hat{t} < 167.33$, is processed by DMD instead of just two periods; see Figure 9. Interpreting this global analysis is especially challenging, because as Figure 2(b) implies, the frequency of large-scale oscillation is no longer approximately constant within this window. The most energetic mode exhibits a frequency of $St = 0.136$, which is between the near-equilibrium frequency of $St = 0.126$ and the limit cycle frequency of $St = 0.141$. In addition, each cluster of Ritz values in Figure 3(d) is now roughly an inverted parabola. That is, there is now a characterizable spread of modes around each major frequency. Schmid (2010) observed a similar pattern of inverted parabolae when analyzing the flow of a jet between two cylinders, at a Reynolds number of 3,000 based on the volume flux velocity and the cylinder diameter. Although Schmid attributes the parabolae to the spread of spatial and temporal scales in the modes of advective-diffusive flows (e.g., in Schmid and Henningson 2000), we believe that they are caused in this case by the increasing Strouhal number of the vortex shedding.

6.4 Comparison with optimized DMD

Given a set of cylinder flow data, DMD and the m-mode optimized DMD typically yield very similar results, when the former is truncated to retain only the m most energetic modes. For this reason, we do not show the optimized DMD Ritz values and vectors of the cylinder flow. The transient regime, however, is an exception. Figure 10 compares the Ritz values from this regime, using DMD and a three-mode optimized DMD. Whereas DMD clusters Ritz values with similar vector energies (as in Figure 3c, d), optimized DMD combines these clusters into single modes. This is advantageous,

since nearly as much dynamical information can be contained in three optimized DMD modes as in seven DMD modes. Furthermore, in Figure 14 and its surrounding discussion, we show that the DMD mode clusters may be problematic when truncating the mode set. Comparing Figure 5(a–c) and (d–f), we observe small but noticeable differences between the primary oscillatory DMD and optimized DMD modes.

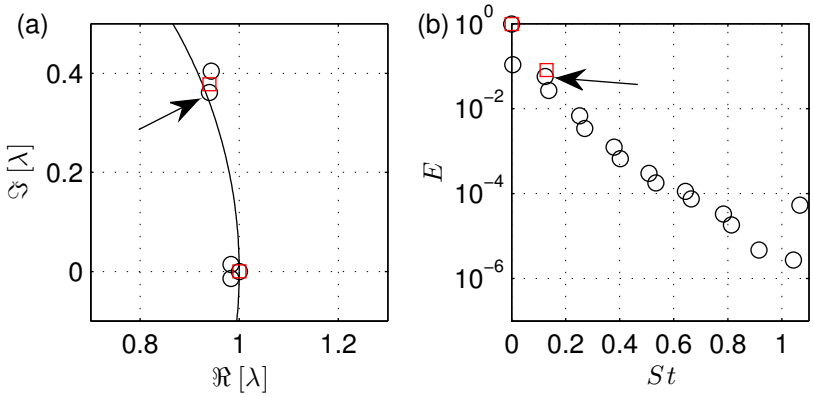

Figure 10: Comparison of the DMD (black circles) and optimized DMD ($m = 3$; red squares) spectra from two periods of data in the transient regime. (a): As the real and imaginary parts. (b): As the normalized vector energy ($E = \sqrt{m+1}\,\|\mathbf{v}\|_2 \,/\, \|[\mathbf{K}\ \mathbf{x}_m]\|_F$ for DMD, $E = \sqrt{p}\,\|\mathbf{v}\|_2 \,/\, \|\mathbf{K}\|_F$ for optimized DMD) versus the Strouhal number. The DMD data are reproduced from Figure 3(c, d). The arrows indicate the DMD Ritz value for which the corresponding mode is plotted in Figure 5(a–c). The nearby optimized DMD mode is plotted in Figure 5(d–f).

To a large extent, the DMD algorithm is successful at extracting frequencies and growth rates from empirical data. Rowley et al. (2009) demonstrates this by comparing the numerical simulation of a three-dimensional jet in crossflow to its DMD results. The DMD algorithm is similarly successful for the cylinder flow (whose dynamics are simpler than those of the jet in crossflow), except in the nonlinear transient regime. Table 1 compares frequencies and growth rates calculated by three techniques. We extract one set of values from C_L oscillation peaks. In the case of DMD, we consider the most energetic oscillatory mode pair, as in Section 6.3. Three-mode optimized DMD recovers a similar complex conjugate pair of modes. Denoting the growth rate by β such that growth is dictated by $\beta^{\hat{t}}$, we have $\beta = |\lambda|^{1/\Delta\hat{t}_s}$.

Table 1: Comparison of Strouhal numbers St and growth rates β over two periods of data from each of the three regimes. Results are obtained from lift force data, DMD, and three-mode optimized DMD.

	St			β		
Regime	C_L	DMD	opt. DMD	C_L	DMD	opt. DMD
Near-equilibrium	0.126	0.126	0.126	1.054	1.054	1.054
Transient	0.130	0.124	0.130	1.044	1.017	1.028
Limit cycle	0.141	0.141	0.141	1.000	1.000	1.000

The frequencies extracted from DMD and optimized DMD match lift force values to very good precision in the near-equilibrium and limit cycle regimes. In the transient regime, however, the three-mode optimized DMD captures the dominant Strouhal number better than DMD. This is because DMD imposes the constraint that the number of modes must be one fewer than the

number of data vectors. By considering only three modes, we truncate the majority of the modal decomposition. In nonlinear dynamics, as in the transient regime, these truncated modes may contain important information that should not be thrown out carelessly. On the other hand, the three-mode optimized DMD tailors the decomposition to find a single complex conjugate pair of modes. Thus, optimized DMD provides the "best" single-frequency representation of the transient flow.

Similarly, growth rates agree near the equilibrium, where dynamics are largely linear, and near the limit cycle, where $\beta = 1$ by definition. A larger variation is evident in the transient regime.

6.5 Comparison with the temporal DFT

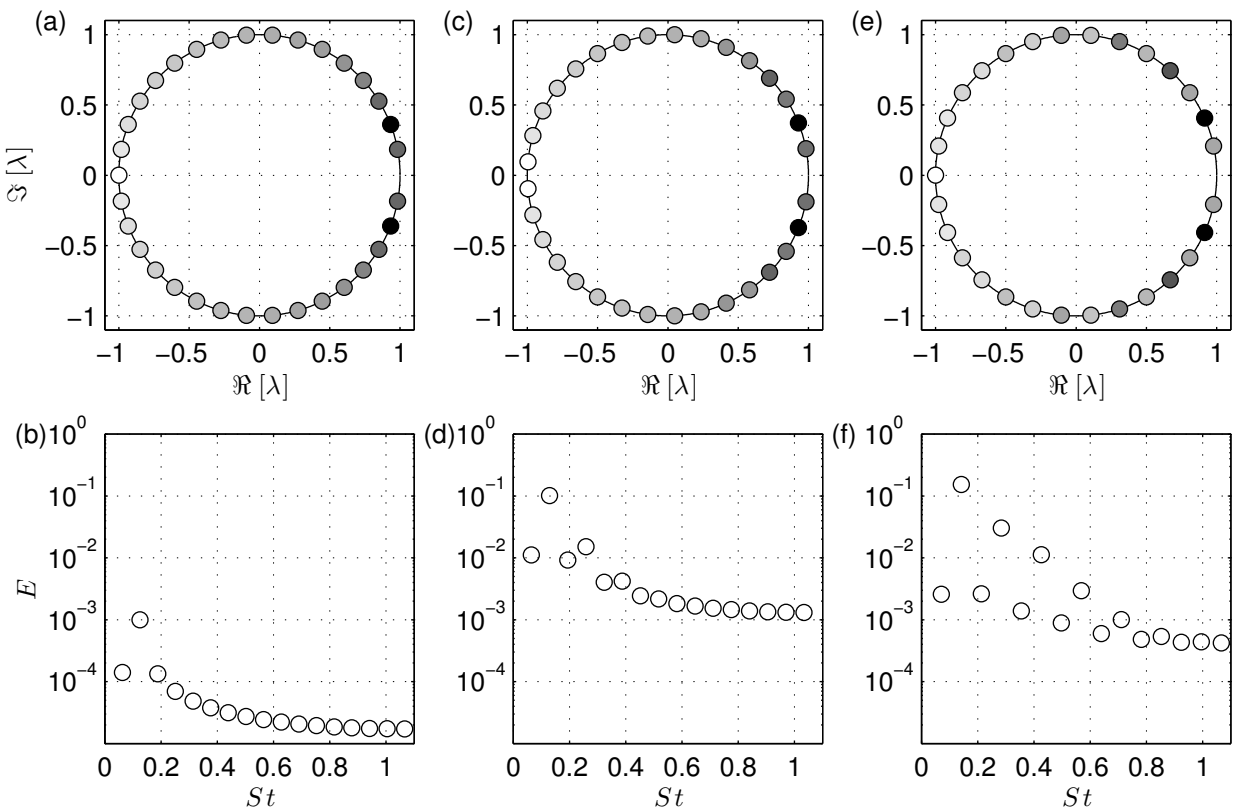

Figure 11: Ritz values λ computed from temporal DFT on two periods of data. Top row: as the real and imaginary parts. Bottom row: as the normalized vector energy $E = \sqrt{m+1}\,\|\mathbf{v}\|_2 \,/\, \|[\mathbf{K}\,\mathbf{x}_m]\|_F$ versus the Strouhal number. Left column: near-equilibrium regime; middle column: transient regime; right column: limit cycle regime.

The spectra produced by the temporal DFT in each of the three data sets is shown in Figure 11. As (30) predicts, the Ritz values are exactly the $m+1$ roots of unity, excluding unity itself. Therefore, to produce modes that correctly capture the large-scale features of the flow, it is imperative that the data set cover an integer number of oscillation periods. Comparing the temporal DFT vector energies to the DMD Ritz vector energies in Figure 3, we observe that the primary oscillation modes are well-captured. These DFT modes look the same as equivalent optimized DMD modes to the naked eye (Figures 4c–e, 5d–f, and 7).

The temporal DFT spectra, however, have a much slower decay rate. This is expected

in the near-equilibrium and transient regimes, since data sets there are non-periodic. The slow decay is generally undesirable, since the implication is that a large number of modes may need to be retained to construct an accurate solution. In addition, the low-energy modes in the transient regime are irregular and generally without meaningful physical interpretation. This is typically not the case with low-energy DMD and optimized DMD modes.

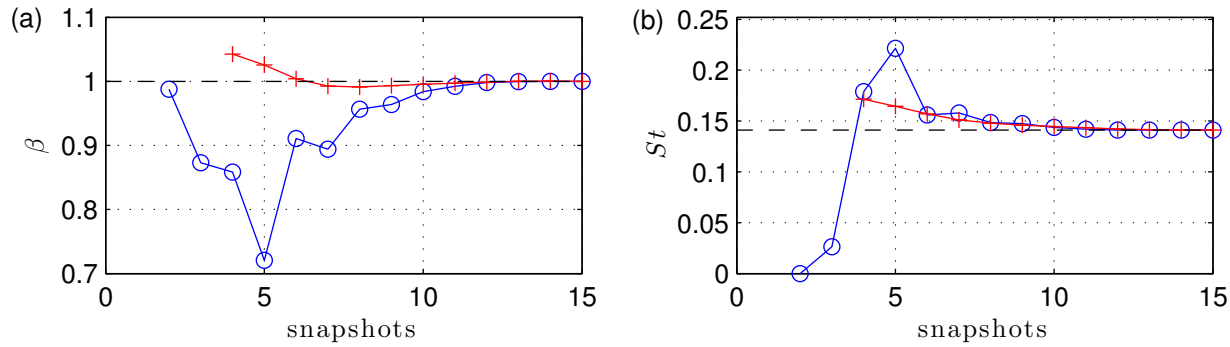

Figure 12: The Ritz value of the limit cycle's primary oscillatory mode, using a full period of data or less in DMD (blue line with circles) and three-mode optimized DMD (red line with pluses). A full period is 15.1 snapshots. The dashed line indicates the correct value. (a): As the growth rate β. (b): As the Strouhal number.

An additional advantage of DMD and optimized DMD over the temporal DFT is highlighted in Figure 12. As explained in Section 4.4, the largest period that the temporal DFT can compute is the time span of the data set. Therefore, it would have limited utility in analyzing small data sets that cover less than a full period of oscillation. On the other hand, DMD and especially optimized DMD are not constrained as severely. In the cylinder's limit cycle, where a full period is 15.1 snapshots, we compute DMD on two to 15 snapshots, and we compute a three-mode optimized DMD on four to 15 snapshots. The former computes the growth rate to within a 5% error of unity when eight or more snapshots are used, and the latter when any number of snapshots is used. Similarly, DMD computes the Strouhal number to within a 5% error of 0.141 when nine or more snapshots are used, and optimized DMD when eight or more are used. With only half a period of data, DMD can compute the modal decomposition modestly accurately; optimized DMD performs even better.

6.6 Comparison with equilibrium subtraction

The DMD spectra of the three data sets—with the equilibrium subtracted—are shown in Figure 13. Upon comparing these spectra with the ones from the original data sets (Figure 3), we notice that the results are nearly identical. One key difference is that the mode closest to $\lambda = 1$ is now significantly weaker. This is a natural consequence of the base flow subtraction. Since the dynamics of the cylinder flow at $Re = 60$ do not depart drastically from the unstable equilibrium, we could expect stationary DMD modes to have some resemblance to the unstable equilibrium.

To the naked eye, these DMD modes from the three regimes look identical to the modes without base flow subtraction; see Figures 4, 5, and 7. We stress again, however, that the base flow subtraction is necessary to ensure that the modes have homogeneous boundary conditions. (The figures only show the innermost domain of the computational grid, and so the actual outer boundaries are not shown.)

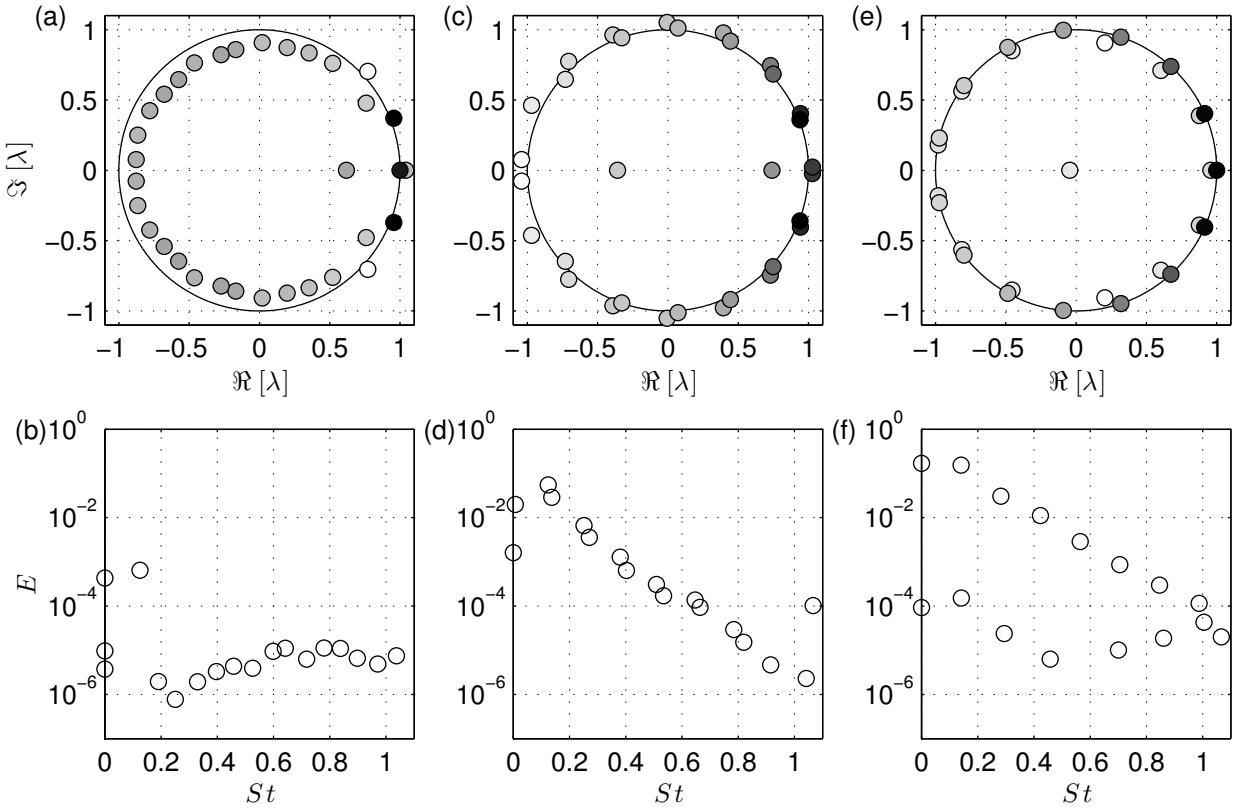

Figure 13: Ritz values λ computed from DMD on two periods of equilibrium-subtracted data. Top row: as the real and imaginary parts. Bottom row: as the normalized vector energy $E = \sqrt{m+1}\,\|\mathbf{v}\|_2 \,/\, \|[\mathbf{K}\,\mathbf{x}_m]\|_F$ versus the Strouhal number. Left column: near-equilibrium regime; middle column: transient regime; right column: limit cycle regime. Except for the Ritz value closest to $\lambda = 1$, the spectra are very close to ones in Figure 3, without the equilibrium subtraction.

6.7 Projection error

For the purpose of constructing reduced-order models (typically by Galerkin projection), we also assess the ability of the discussed modal decompositions to yield good projection subspaces. By construction, POD produces the modal subspace that retains the largest possible vector energy after projection (Holmes et al. 1996). Proper orthogonal decomposition, however, is fundamentally decoupled from any sense of time or dynamics, and therefore may not provide the best mode basis for constructing dynamical models.

Suppose that $\mathbf{V}$ is a linearly independent set of modes, and $\mathbf{P_V} = \mathbf{V}\,(\mathbf{V}^*\mathbf{V})^{-1}\,\mathbf{V}^*$ is the orthogonal projection onto the modes. We define the projection error of a data vector $\mathbf{x}$ as

$$\epsilon\,(\mathbf{x};\mathbf{V}) \triangleq \frac{\|\mathbf{x} - \mathbf{P_V}\mathbf{x}\|_2}{\|\mathbf{x}\|_2}. \tag{48}$$

A low projection error indicates that the set of modes in $\mathbf{V}$ captures a large amount of information in $\mathbf{x}$, from an energy standpoint. Again, we stress that modal decompositions with low projection errors do not necessarily create accurate reduced-order models. Nevertheless, this error is still a meaningful measure to consider. A large error suggests that reduced-order models created from the mode set cannot faithfully represent the true state of the dynamical system.

In Figure 14(a), we project the two-dimensional velocity data from the cylinder flow simulation onto different mode sets, and we show the resulting projection error $\epsilon\left(\mathbf{x}\left(\hat{t}\right);\mathbf{V}\right)$. Linear stability modes are computed from the near-equilibrium regime and orbital modes from the limit cycle, both using DMD. Three DMD modes—the $\lambda \approx 1$ mode and the most energetic complex conjugate pair—are used for projection, and the rest are truncated. (See Figures 4, 6, and 7.) As expected, linear stability modes yield low projection errors near the equilibrium, and orbital modes yield low projection errors in the limit cycle. The errors become large, however, away from the regimes from which we compute each set of modes. Noack et al. (2003) shows that the inclusion of a shift mode, defined in this case as the limit cycle mean minus the unstable equilibrium state, improves the projection error. This is also shown in Figure 14(a). We comment that there is little change in projection error if the mean flow and two temporal DFT modes are used instead of three DMD modes; therefore, temporal DFT results are not shown.

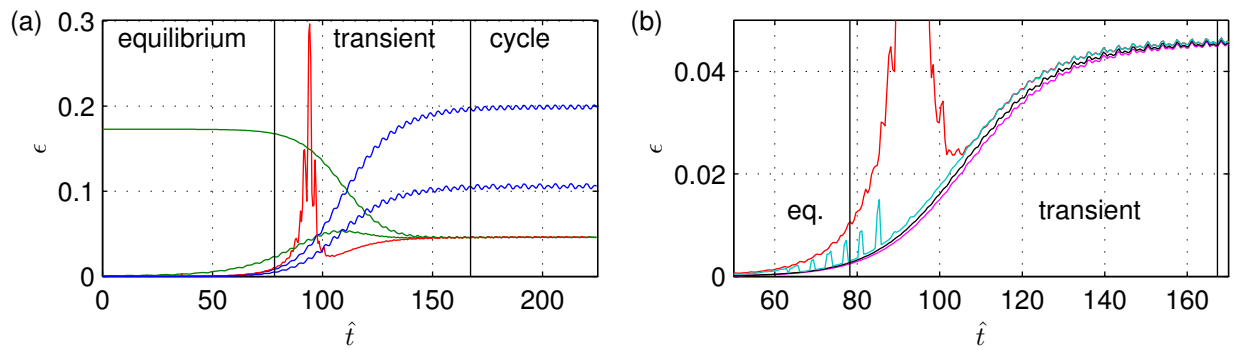

Figure 14: (a): The projection error using different DMD modes sets. (blue, top): linear stability modes; (blue, bottom): linear stability modes and the shift mode; (green, top): orbital modes; (green, bottom): orbital modes and the shift mode; (red): time-varying modes. (b): The projection error in the transient regime using different time-varying modes. (red): DMD, as in (a); (magenta): temporal DFT; (black): optimized DMD; (cyan): equilibirium-subtracted DMD. For reference, compare these two plots to Figure 2.

We might expect that the projection error could be reduced if data were projected onto modes computed only from their local temporal neighborhood. Thus, we also project each data vector onto three time-varying modes, which we compute by DMD on a data set beginning half a cycle before the data vector and terminating half a cycle after. As shown in Figure 14, this is effective at minimizing projection errors, except around $\hat{t} = 94$. There, the error becomes large because there exist two DMD modes with Ritz values $\lambda \approx 1$, and we only retain whichever is closer to $\lambda = 1$. Should the other be included, the spike in the error would not exist. Nonetheless, this shows that truncating DMD modes may cause unintended and harmful effects.

In Figure 14(b), we also compute the projection error onto three time-varying temporal DFT and optimized DMD modes. They are nearly equally good at minimizing the projection error, and they outperform any other set of modes tested. The figure also includes the projection error onto time-varying DMD modes with equilibrium subtraction. In this case, the projection is onto the span of the equilibrium, the $\lambda \approx 1$ mode, and the most energetic oscillatory pair of modes.

Finally, we remark that prior studies in the literature suggest the utility of DMD modes for Galerkin modeling. As mentioned in Sections 2 and 6.3, the DMD modes near the equilibrium are approximately the eigenmodes of the linearized dynamics, and DMD modes near the orbit are approximately the POD modes. A number of previous studies have already investigated the use of these modes in Galerkin modeling; see the first several references of Tadmor et al. (2010). For

a comprehensive study specifically relating to the Galerkin modeling of the cylinder wake, refer to Noack et al. (2003); this particular study shows that the wake is accuracy modeled using a combination of POD modes, linear stability modes, and the shift mode.

7 Conclusion

The Koopman operator is a linear operator that represents nonlinear dynamics without linearization. Dynamic mode decomposition is a data-based method for estimating Koopman eigenvalues and modes. It approximates a dynamical trajectory, often generated by a nonlinear system, as the result of a linear process. The Ritz values of the decomposition yield growth rates and frequencies, and the Ritz vectors yield corresponding directions. Unlike many previous decomposition techniques such as POD and balanced POD, DMD is not tightly constrained; the data need be neither periodic nor linearly generated for DMD to construct a meaningful modal decomposition.

We prove that the DMD algorithm of Rowley et al. (2009) and Schmid (2010) produces a unique decomposition if and only if all vectors except the last in a given data set are linearly independent, and the Ritz values are distinct. If either condition is violated, then non-unique Ritz values and vectors may still exist.

To use DMD modes in a reduced-order Galerkin model of a partial differential equation, a "base flow" that satisfies the appropriate boundary conditions must typically be included. The mean of the data is a common choice for a base flow, especially in POD-based analyses. Dynamic mode decomposition on a mean-subtracted set of data, however, is exactly equivalent to a temporal DFT and harmonic averaging. This implies that for $m+1$ data vectors, the Ritz values are precisely the $m+1$ roots of unity, excluding unity itself. The inability to compute growth rates, as well as the predetermination of frequencies regardless of data content, are fundamental properties of the temporal DFT. In addition, the temporal DFT cannot capture a particular frequency if the data set covers less than one corresponding period. Furthermore, the decay of temporal DFT vector energies is slow if the data are non-periodic. We can avoid these issues by choosing the equilibrium point instead as the base flow.

Next, we introduce an optimized DMD, in which the decomposition is tailored specifically to the desired number of output modes. In this decomposition, the representation of any data vector may contain a residual, but the total residual over all data vectors is minimized. In the absence of an analytic algorithm, we use simulated annealing and quasi-Newton iterators to compute the optimized DMD. Aside from these iterators, this decomposition generally does not suffer from numerical sensitivity as DMD does. Optimized DMD is also superior for computing a small number of modes, since the decomposition is computed for the given number of modes and does not require arbitrary truncation.

The two-dimensional incompressible flow over a cylinder is used as a test bed for the aforementioned decomposition methods. Dynamic mode decomposition recovers the leading unstable eigenmodes from data near the unstable equilibrium. It also recovers POD-like modes from data near the limit cycle. When a small window of snapshots traverses the cylinder data from the equilibrium to the limit cycle, the Ritz values and vectors morph gradually.

The optimized DMD and equilibrium-subtracted DMD of the cylinder flow reproduce the DMD results to a large extent. A significant improvement can be seen in the transient regime, however. Optimized DMD computes growth rates and frequencies more faithfully, and it avoids clusters of Ritz values, which are problematic for truncating mode sets. Equilibrium-subtracted DMD, on the other hand, largely preserves the DMD modes and spectra, while correctly accounting for boundary conditions. Temporal DFT modes of the cylinder flow look qualitatively like optimized

DMD modes, but the roll-off in the Ritz vector energies is slow when data are not taken from the periodic orbit. Unlike the temporal DFT, DMD is able to compute frequencies and growth rates from little over half a period of data; optimized DMD fares even better.

As the cylinder flow evolves from the unstable equilibrium to the stable limit cycle, time-varying temporal DFT, optimized DMD, and equilibrium-subtracted DMD modes are better able to represent data (in the sense of energy) than time-varying DMD modes. These mode sets maintain a low error through the entire trajectory of the cylinder flow, from the equilibrium to the limit cycle. Based on these findings, we hypothesize that the use of these time-varying modes may yield accurate Galerkin models.

Acknowledgements

This work was supported by the Department of Defense National Defense Science & Engineering Graduate (DOD NDSEG) Fellowship, the National Science Foundation Graduate Research Fellowship Program (NSF GRFP), and AFOSR grant FA9550-09-1-0257.

References

E. Åkervik, L. Brandt, D. S. Henningson, J. Hœpffner, O. Marxen, and P. Schlatter. Steady solutions of the Navier-Stokes equations by selective frequency damping. *Phys. Fluids*, 18(6): 068102, 2006.

T. Colonius and K. Taira. A fast immersed boundary method using a nullspace approach and multi-domain far-field boundary conditions. *Comput. Meth. Appl. Mech. Eng.*, 197(25–28):2131–2146, 2008.

F. Giannetti and P. Luchini. Structural sensitivity of the first instability of the cylinder wake. *J. Fluid Mech.*, 581:167–197, 2007.

E. A. Gillies. Low-dimensional control of the circular cylinder wake. *J. Fluid Mech.*, 371:157–178, 1998.

P. Holmes, J. L. Lumley, and G. Berkooz. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry.* Cambridge University Press, 1996.

M. Ilak and C. W. Rowley. Modeling of transitional channel flow using balanced proper orthogonal decomposition. *Phys. Fluids*, 20(3):034103, 2008.

O. Marquet, D. Sipp, and L. Jacquin. Sensitivity analysis and passive control of cylinder flow. *J. Fluid Mech.*, 615:221–252, 2008.

I. Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dyn.*, 41(1–3):309–325, 2005.

I. Mezić and A. Banaszuk. Comparison of systems with complex behavior. *Physica D*, 197(1–2): 101–133, 2004.

B. C. Moore. Principal component analysis in linear systems: controllability, observability, and model reduction. *IEEE Trans. Autom. Control*, 26(1):17–32, 1981.

B. R. Noack, K. Afanasiev, M. Morzyński, G. Tadmor, and F. Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.*, 497:335–363, 2003.

B. R. Noack, P. Papas, and P. A. Monkewitz. The need for a pressure-term representation in empirical Galerkin models of incompressible shear flow. *J. Fluid Mech.*, 523:339–365, 2005.

K. Petersen. *Ergodic Theory.* Cambridge University Press, 1983.

W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes.* Cambridge University Press, 3rd edition, 2007.

M. Provansal, C. Mathis, and L. Boyer. Bénard-von Kármán instability: transient and forced regimes. *J. Fluid Mech.*, 182:1–22, September 1987.

A. Roshko. On the development of turbulent wakes from vortex streets. Technical Report NACA 1191, National Advisory Committee for Aeronautics, 1954.

K. Roussopoulos. Feedback control of vortex shedding at low Reynolds numbers. *J. Fluid Mech.*, 248:267–296, 1993.

C. W. Rowley. Model reduction for fluids, using balanced proper orthogonal decomposition. *Int. J. Bifurc. Chaos*, 15(3):997–1013, 2005.

C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson. Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 641:115–127, 2009.

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

P. J. Schmid. Application of the dynamic mode decomposition to experimental data. *Exp. Fluids*, 50(4):1123–1130, 2011.

P. J. Schmid and D. S. Henningson. *Stability and Transition in Shear Flows.* Springer, 2000.

P. J. Schmid, L. Li, M. P. Juniper, and O. Pust. Applications of the dynamic mode decomposition. *Theor. Comput. Fluid Dyn.*, 25(1–4):249–259, 2011.

G. Tadmor, O. Lehmann, B. R. Noack, and M. Morzyński. Mean field representation of the natural and actuated cylinder wake. *Phys. Fluids*, 22(3):034102, 2010.

K. Taira and T. Colonius. The immersed boundary method: a projection approach. *J. Comput. Phys.*, 225(2):2118–2137, 2007.

C. H. K. Williamson. Vortex dynamics in the cylinder wake. *Annu. Rev. Fluid Mech.*, 28:477–539, 1996.