

Physics-informed dynamic mode decomposition (piDMD)

Peter J. Baddoo^{*1}, Benjamin Herrmann², Beverley J. McKeon³, J. Nathan Kutz⁴, and Steven L. Brunton⁵

¹Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

²Department of Mechanical Engineering, University of Chile, Beauchef 851, Santiago, Chile

³Graduate Aerospace Laboratories, California Institute of Technology, Pasadena CA 91125, USA

⁴Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA

⁵Department of Mechanical Engineering, University of Washington, Seattle, WA 98195, USA

Abstract

In this work, we demonstrate how physical principles – such as symmetries, invariances, and conservation laws – can be integrated into the *dynamic mode decomposition* (DMD). DMD is a widely-used data analysis technique that extracts low-rank modal structures and dynamics from high-dimensional measurements. However, DMD frequently produces models that are sensitive to noise, fail to generalize outside the training data, and violate basic physical laws. Our physics-informed DMD (piDMD) optimization, which may be formulated as a Procrustes problem, restricts the family of admissible models to a matrix manifold that respects the physical structure of the system. We focus on five fundamental physical principles – conservation, self-adjointness, localization, causality, and shift-invariance – and derive several closed-form solutions and efficient algorithms for the corresponding piDMD optimizations. With fewer degrees of freedom, piDMD models are less prone to overfitting, require less training data, and are often less computationally expensive to build than standard DMD models. We demonstrate piDMD on a range of challenging problems in the physical sciences, including energy-preserving fluid flow, travelling-wave systems, the Schrödinger equation, solute advection-diffusion, a system with causal dynamics, and three-dimensional transitional channel flow. In each case, piDMD significantly outperforms standard DMD in metrics such as spectral identification, state prediction, and estimation of optimal forcings and responses.

1 Introduction

Integrating partial knowledge of physical principles into data-driven techniques is a primary goal of the scientific machine learning (ML) community [1]. Physical principles – such as conservation laws, symmetries, and invariances – can be incorporated into ML algorithms in the form of inductive biases, thereby ensuring that the learned models are constrained to the correct physics. Recent successful examples of ML algorithms that have been modified to respect physical principles include neural networks [2–10], kernel methods [11, 12], deep generative models [13], and sparse regression [14–18]. These examples demonstrate that incorporating partially-known physical principles into machine learning architectures can increase the accuracy, robustness, and generalizability of the resulting models, while simultaneously decreasing the required training data. In this work, we integrate knowledge of physical principles into one of the most widely-used methods in data-driven dynamical systems research: the dynamic mode decomposition [19–24].

^{*} Corresponding author (baddoo@mit.edu).

The dynamic mode decomposition (DMD) is a data diagnostic technique that extracts coherent spatial-temporal patterns from high-dimensional time series data [19, 24]. Although DMD originated in the fluid dynamics community [19], the algorithm has since been applied to a wealth of dynamical systems including in epidemiology [25], robotics [26, 27], neuroscience [28], quantum control [29], power grids [30], and plasma physics [31, 32]. Despite its widespread successes, DMD is highly sensitive to noise [33–35], fails to capture travelling wave physics, and can produce overfit models that do not generalize. Herein, we demonstrate that integrating physics into the learning framework can help address these challenges.

Suppose that we are studying a dynamical system defined by $\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x})$ (continuous time) or $\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k)$ (discrete time) where $\mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is unknown. DMD identifies the best low-rank linear approximation of \mathbf{F} given a collection of m pairs of measurements $\{\mathbf{x}_j, \mathbf{y}_j\}$; in other words, DMD seeks a rank r matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ such that

$$\mathbf{y}_j \approx \mathbf{A}\mathbf{x}_j \quad (1)$$

for $j = 1, \dots, m$. Arranging the data measurements into $n \times m$ matrices \mathbf{X} and \mathbf{Y} allows us to phrase the above formally as

$$\text{DMD regression: } \underset{\text{rank } (\mathbf{A})=r}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{AX}\|_F. \quad (2)$$

After approximately solving (2), the DMD process computes the dominant spectral properties of the learned linear operator [19–24]. The rank- r constraint in (2) is motivated by the assumed modal structure of the system but doesn't account for other important physical properties. For example, one limitation of DMD is that the solution of (2) lies within the span of \mathbf{Y} , so the learned model rarely generalizes outside the training regime. We address this limitation and others by embedding partial physics knowledge into the learning process.

In this work, we incorporate physical principles into the optimization (2) by constraining the solution matrix \mathbf{A} to lie on a matrix manifold $\mathcal{M} \subseteq \mathbb{R}^{n \times n}$:

$$\text{piDMD regression: } \underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{AX}\|_F. \quad (3)$$

The matrix manifold \mathcal{M} is dictated by the known physics of the system at hand. For example, we may select \mathcal{M} such that its members satisfy certain symmetries that we know are obeyed by the system at hand. We call (3) '*physics-informed* DMD' (piDMD) as the optimization integrates underlying knowledge of the system physics into the learning framework. Again, the spectral features of the solution of (3) can be computed to give insight into the modes that dominate the system and their dynamics. The low-rank DMD regression (2) is a special case of piDMD where \mathcal{M} is the manifold of rank- r matrices. However, piDMD models are not always low rank (see, for example, §4.4) but nevertheless have few degrees of freedom due to the physics-informed constraints. Constraining the solution matrix improves the generalizability of the learned model, reduces sensitivity to noise and reduces the demand for large training sets. More broadly, constraining a linear, dimensionality-reduced model to a manifold \mathcal{M} has been shown to result in a variety of optimization formulations and techniques which can improve low-rank approximations [36]. In figure 1 we compare the performance of standard (2) and physics-informed (3) DMD for a simple travelling wave. We constrain the piDMD model to the shift-invariant symmetry of the underlying system. As a result, the piDMD model produces an accurate prediction of the future dynamics,

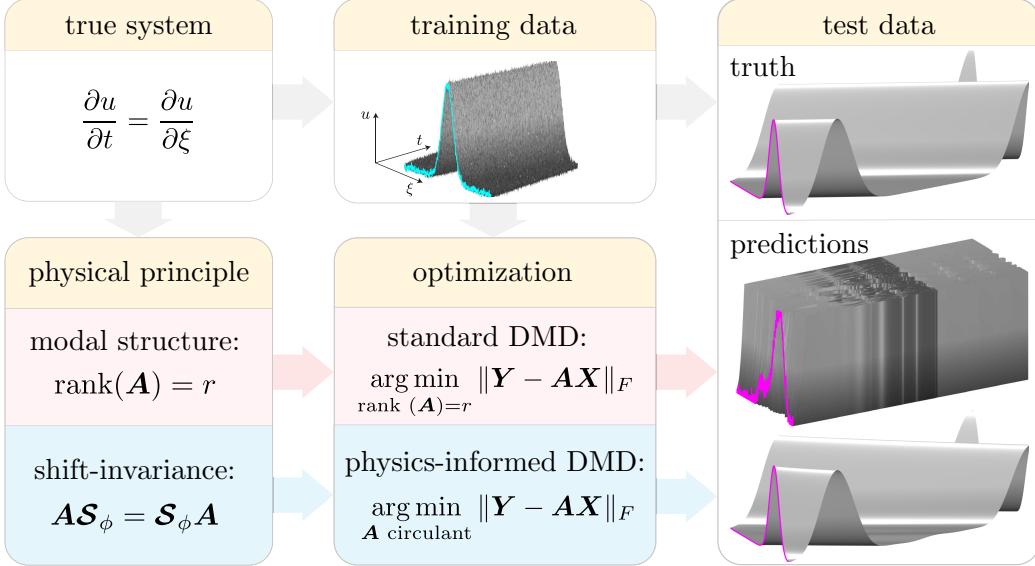


Figure 1: Comparing standard dynamic mode decomposition (specifically, optimized DMD [23]) and physics-informed dynamic mode decomposition applied to the advection equation. The data is contaminated with 2% additive Gaussian noise. Having trained DMD models, we then perform predictions for different initial conditions. Standard DMD fails whereas piDMD produces a faithful prediction. In the figure, \mathcal{S}_ϕ represents the shift operator (see §4.1).

whereas the standard DMD model blows up immediately. This simple example illustrates that embedding basic physics into DMD can substantially improve the algorithm’s performance.

The optimization problem (3) is well-known in the statistical literature as a *Procrustes problem* [37–49]. Our recasting of the DMD regression as a Procrustes problem is a new connection, and is the basis of piDMD. This perspective enables us to leverage the substantial extant work on Procrustes problems into new application areas. The literature contains many exact solutions for Procrustes problems, including notable cases of orthogonal matrices [45], and symmetric matrices [41].

A significant challenge in physics-informed ML architectures is to develop implementations that incorporate known physics but also scale well to higher dimensions [1]. A major contribution of this work is several exact solutions, in terms of standard linear algebra operations, of Procrustes problems of physical relevance. Where possible, we have made use of matrix factorizations and low-rank representations to alleviate the computational burden of implementation; it is rarely necessary to form \mathbf{A} explicitly when computing its spectral properties or other diagnostics. For complicated matrix manifolds, an exact solution of the Procrustes problem (3) may be intractable and we must resort to algorithmic approaches. Fortunately, optimization on matrix manifolds is a mature field [50], and many algorithms are implemented in the open source software package ‘Manopt’ [51, www.manopt.org].

The remainder of the paper is arranged as follows. In section 2 we provide background information on DMD and Procrustes problems. Then, in section 3, we describe the broad framework of piDMD. We consider a range of applications in section 4, with a focus on shift-invariant, conservative, self-adjoint, local and causal systems. Section 5 concludes with a discussion of the limitations of piDMD and suggests several future research directions. An open source implementation of piDMD with support for over 30 matrix manifolds is available in MATLAB at www.github.com/baddoo/piDMD.

2 Mathematical background

In this section we provide further details on DMD and Procrustes problems. Throughout the article, we assume that we have access to m snapshots pairs of n features each: $\{(\mathbf{x}_j, \mathbf{y}_j), j = 1, \dots, m\}$. For example, \mathbf{x}_j may be a discretized fluid flow field at time t_j and \mathbf{y}_j may be the flow field at the next time step t_{j+1} . It is convenient to arrange the data into $n \times m$ snapshot matrices of the form

$$\mathbf{X} = \begin{bmatrix} | & | & | \\ \mathbf{x}_1 & \cdots & \mathbf{x}_m \\ | & | & | \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{x}}_1 & \cdots \\ \vdots & \vdots \\ \tilde{\mathbf{x}}_n & \cdots \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} | & | & | \\ \mathbf{y}_1 & \cdots & \mathbf{y}_m \\ | & | & | \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{y}}_1 & \cdots \\ \vdots & \vdots \\ \tilde{\mathbf{y}}_n & \cdots \end{bmatrix} \quad (4)$$

so that rows i of \mathbf{X} and \mathbf{Y} are the measurements of the i -th features and columns j are the j -th temporal snapshots. Henceforth, we use $\tilde{\cdot}$ to represent a row vector.

2.1 Dynamic mode decomposition

DMD was initially proposed as a dimensionality reduction technique that extracts dominant spatio-temporal coherent structures from high-dimensional time-series data [19]. In particular, DMD identifies the leading-order spatial eigenmodes of the matrix \mathbf{A} in (1), along with a linear model for how the amplitudes of these coherent structures evolve in time. DMD has been applied to a range of systems, as summarized in the monograph [22] and the recent review by Schmid [24].

To address the challenges associated with DMD, researchers have derived many variations on the original algorithm, including sparsity promoting DMD [52], DMD with control [53], noise-robust variants [23, 33–35], recursive DMD [54], online DMD [55, 56], and versions for under-resolved data in space or time [57–59]. At present, the most widely used variant is ‘exact DMD’ [21], which phrases the DMD solution in terms of the Moore–Penrose pseudoinverse. This solution produces a more even distribution of errors between the terms. Due to its simplicity and widespread use, most of the comparisons in this paper are made between exact DMD and piDMD.

It is assumed that each \mathbf{x}_j and \mathbf{y}_j are connected by an unknown dynamical system of the form $\mathbf{y}_j = \mathbf{F}(\mathbf{x}_j)$; for discrete-time dynamics $\mathbf{y}_j = \mathbf{x}_{j+1}$ and for continuous-time dynamics $\mathbf{y}_j = \dot{\mathbf{x}}_j$. DMD aims to learn the dominant behaviour of \mathbf{F} by finding the best linear approximation for \mathbf{F} given the data and then performing diagnostics on that approximation. Thus, DMD seeks the linear operator \mathbf{A} that best maps the snapshots in the set $\{\mathbf{x}_j\}$ to those in the set $\{\mathbf{y}_j\}$:

$$\mathbf{y}_j \approx \mathbf{A}\mathbf{x}_j \quad \text{for } j = 1, \dots, m. \quad (5)$$

Expressed in terms of the snapshot matrices in (4), the linear system in (5) becomes

$$\mathbf{Y} \approx \mathbf{A}\mathbf{X}, \quad (6)$$

and the optimization problem for \mathbf{A} is given by (2). The minimum-norm solution for \mathbf{A} is given by

$$\mathbf{A} = \mathbf{Y}\mathbf{X}^\dagger = \mathbf{Y}\mathbf{V}\Sigma^\dagger\mathbf{U}^*, \quad (7)$$

where \dagger indicates the Moore–Penrose pseudoinverse [60] and $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^*$ is the singular value decomposition. In many applications, the state dimension n is very large, and forming or storing \mathbf{A} explicitly becomes impractical. Instead, we use a rank- r approximation for \mathbf{A} , denoted by $\hat{\mathbf{A}}$, where $r \ll n$. To form $\hat{\mathbf{A}}$, we construct the optimal rank- r approximation for \mathbf{X} using the

truncated singular value decomposition [61]: $\mathbf{X} \approx \mathbf{U}_r \boldsymbol{\Sigma}_r \mathbf{V}_r^*$. We then project \mathbf{A} onto the leading r principal components of \mathbf{X} as

$$\hat{\mathbf{A}} = \mathbf{U}_r^* \mathbf{A} \mathbf{U}_r = \mathbf{U}_r^* \mathbf{Y} \mathbf{V}_r \boldsymbol{\Sigma}_r^{-1}. \quad (8)$$

It is now computationally viable to compute the eigendecomposition of $\hat{\mathbf{A}}$ as

$$\hat{\mathbf{A}} \hat{\Psi} = \hat{\Psi} \Lambda. \quad (9)$$

The eigenvectors of \mathbf{A} can be approximated from the reduced eigenvectors Ψ by [21]

$$\Psi = \mathbf{Y} \mathbf{V} \boldsymbol{\Sigma}^{-1} \hat{\Psi}. \quad (10)$$

This eigendecomposition is connected to the Koopman operator of the system, and allows reconstructions and predictions [20, 22, 62, 63]. For example, for a discrete-time system (i.e., $\mathbf{y}_k = \mathbf{x}_{k+1}$) with evenly spaced samples in time, then the eigenvectors form a linearly independent set and

$$\mathbf{x}_j = \Psi \Lambda^{j-1} \mathbf{b} \quad (11)$$

where the vector \mathbf{b} contains the weights of the modes in the initial condition: $\mathbf{b} = \Psi^\dagger \mathbf{x}_1$. From the above, it is clear that the eigenvalues Λ govern the temporal behaviour of the system and the eigenvectors Ψ are the spatial modes.

2.2 Procrustes problems

Procrustes problems (3) comprise of finding the optimal transformation between two matrices subject to certain constraints on the class of admissible transformations [40, 41, 45]. According to Greek mythology, Procrustes was a bandit who would stretch or amputate the limbs of his victims to force them to fit onto his bed. Herein, \mathbf{X} plays the role of Procrustes' victim, \mathbf{Y} is the bed, and \mathbf{A} is the 'treatment' (stretching or amputation).¹ Procrustes problems (3) seek to learn the treatment \mathbf{A} that best represents the data measurements \mathbf{X} and \mathbf{Y} . The minimization is usually phrased in the Frobenius norm. Procrustes analysis finds relevance in a wealth of fields including sensory analysis [64], data visualization [65], neural networks [66], climate science [67], and solid mechanics [41]. A summary is available in the monograph [40].

The earliest and most common formulation is the 'orthogonal Procrustes problem' [40, 45]. Suppose that we have two sets of measurements \mathbf{X} and \mathbf{Y} that we know are related by an unitary (orthogonal) transformation (i.e., a rotation or reflection). The goal is to learn the best unitary transformation that relates the data measurements. Thus, \mathbf{A} is constrained to be a unitary matrix and the minimization problem is

$$\underset{\mathbf{A}^* \mathbf{A} = \mathbf{I}}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{A} \mathbf{X}\|_F. \quad (12)$$

The solution to (12) was derived by Schönemann [45] as

$$\mathbf{A} = \mathbf{U}_{YX} \mathbf{V}_{YX}^* \quad (13)$$

where $\mathbf{U}_{YX} \boldsymbol{\Sigma}_{YX} \mathbf{V}_{YX}^* = \mathbf{Y} \mathbf{X}^*$ is a full singular value decomposition. Alternatively, $\mathbf{A} = \mathbf{U}_P$, where $\mathbf{U}_P \mathbf{H}_P = \mathbf{Y} \mathbf{X}^*$ is a polar decomposition. The solution is unique if and only if $\mathbf{Y} \mathbf{X}^*$ is full-rank. As we show in section 4.2, a unitary matrix corresponds to an energy-preserving operator.

There are many solutions for Procrustes problems available in the literature for different matrix constraints [41]. When exact solutions are not possible, algorithmic solutions can be effective [51].

¹This terminology was first introduced by Hurley and Cattell [37].

3 Physics-informed dynamic mode decomposition (piDMD)

Incorporating physics into ML algorithms involves supplementing an existing technique with additional biases. Usually, ML practitioners take one of three approaches [1]. First, observational biases can be embedded through data augmentation techniques; however, augmented data is not always available, and incorporating additional training samples can become computationally expensive. Second, physical constraints can be included by suitably penalising the loss function [8]. Third, inductive biases can be incorporated directly into the machine learning architecture in the form of mathematical constraints [16]. As noted by [1], this third approach is arguably the most principled method since it produces model that strictly satisfy the physical constraints. piDMD falls into the final category.

piDMD incorporates physical principles by constraining the matrix manifold of the DMD regression problem (3). We abandon the low-rank representation of A typically sought in model order reduction in favour of alternative or additional matrix structures more relevant to the physical problem at hand. In figure 2 we illustrate the matrix structures used in this paper, along with their corresponding physical principle and references for the optimal solutions.

Analysing a system with the piDMD framework requires four steps: modeling, interpretation, optimization, and diagnostics:

1. Modeling: When studying a system from a data-driven perspective, it is common to have some partial knowledge of the underlying physics of the system. This first step in piDMD asks the user to summarise the known or suspected physical properties of the system at hand.
2. Interpretation: Having determined the physical principle we would like to enforce, we must translate these laws into the matrix manifold to which the linear model should be constrained. This step involves knowledge of the data collection procedure such as the spatial grid or the energy inner product.
3. Optimization: Equipped with a target matrix manifold, we may now solve the relevant Procrustes problem (3). By suitably modifying the optimization routine, we can guarantee that the resulting model satisfies the physical principle identified in step 1.
4. Diagnostics: The final step of piDMD involves extracting physical information from the learned model A . For example, one may wish to analyse the spectrum or modes, compute the resolvent modes [68], perform predictions, or investigate other diagnostics.

As an example, consider the travelling wave solution explored in the introduction (figure 1). We began with the physical principle that the system is shift invariant ('modeling'), and this lead us to seek a circulant matrix model ('interpretation'). We found a solution to the corresponding Procrustes problem ('optimization') and investigated the predictive behaviour of the model ('diagnostics').

While each of the four steps can be problematic, the optimization step is arguably the most conceptually and computationally difficult. In particular, finding a solution for the desired Procrustes problem can be challenging. Moreover, any solutions are quite peculiar to the matrix manifold under consideration: the symmetric Procrustes problem is largely unconnected to, say, the tridiagonal Procrustes problem. Thus, in this article, we present many new solutions for (3) with different physics-informed matrix manifold constraints. When exact solutions are not possible

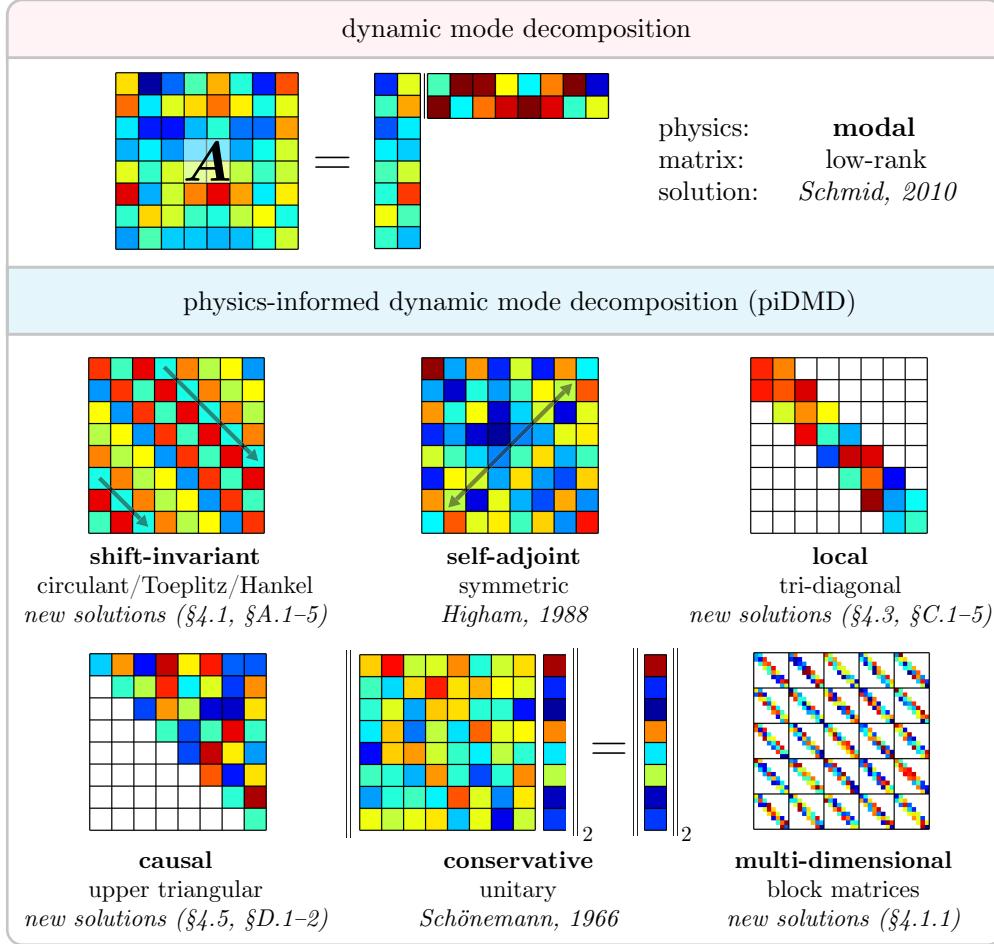


Figure 2: Visual illustrations of the types of matrices, their corresponding physical principles and references to the solutions of the corresponding optimization problem (3).

(e.g., if the manifold constraint is quite complicated) there are sophisticated algorithmic solutions available [51].

Standard DMD exploits the low-rank structure of A to efficiently perform diagnostics on the learned model. Some of the manifolds we consider (such as circulant, tridiagonal, upper triangular) do not have an obvious or useful low-rank approximation. Instead, these matrices often have an alternative structure that can be exploited to perform fast diagnostic tests. For example, tridiagonal matrices rarely have a meaningful low-rank approximation, but nevertheless admit a fast eigendecomposition and singular value decomposition [60].

Some of the matrix manifolds we consider (such as symmetric, triangular, tridiagonal or circulant) can be phrased as linear equality constraints and can, in principle, be solved with linear-equality constrained least squares [60]. However, the number of equality constraints needed is $\mathcal{O}(n^2)$ so the resulting least squares matrix will have $\mathcal{O}(n^2)$ rows, which is intractably large in most applications. Thus, we avoid phrasing the piDMD constraints in terms of linear equality constraints and instead exploit properties of the matrix manifold to find solutions that can be efficiently implemented (see sections 4.3 and 4.1, for example).

4 Applying piDMD to enforce canonical physical structures

We now present a range of examples of piDMD. Each of the following sections investigates a detailed application of piDMD for a specific physical principle. A summary is illustrated in figure 3. The physical principles, corresponding matrix structures, and optimal solutions and extensions are summarised below:

- §4.1: Shift-invariant: circulant (and symmetric, skew-symmetric, or unitary, §A.1), low rank (§A.2), non-equally spaced samples (§A.3), total least squares (§A.4), Toeplitz & Hankel (§A.5)
- §4.2: Conservative: unitary (§2.2)
- §4.3: Self-adjoint: symmetric (§4.3), skew-symmetric (§B.1), symmetric in a subspace (§4.3.1)
- §4.4: Local: tridiagonal (§4.4), variable diagonals (§C.1), periodic (§C.2), symmetric tridiagonal (§C.3), total least squares (§C.4), regularized locality (§C.5).
- §4.5: Causal: upper triangular (§D.1 D.2)

4.1 Shift-invariant systems

Shift invariance (also called ‘translation invariance’) is ubiquitous in the physical sciences. Spatially homogeneous systems – such as constant coefficient PDEs or convolution operators – are shift invariant, and thus appear identical with respect to any (stationary) observer. In view of Noether’s theorem [69], the conserved quantity related to shift invariance is linear momentum; as such, incorporating shift invariance into a DMD model means that the model preserves linear momentum. While it may be unusual for practical engineering problems to have a truly homogeneous direction, many problems of basic physical interest – such as wall-bounded turbulent flow – frequently possess a spatially homogeneous dimension. In the sequel, we derive a piDMD formulation for shift-invariant operators.

We define \mathcal{S}_ϕ as the ϕ -shift operator i.e. $\mathcal{S}_\phi v(\xi) = v(\xi + \phi)$ for all test functions v . We say that a space-continuous linear operator \mathcal{A} is shift invariant if \mathcal{A} commutes with the ϕ -shift operator for all shifts ϕ :

$$\mathcal{S}_\phi \mathcal{A} = \mathcal{A} \mathcal{S}_\phi. \quad (14)$$

An application of the shift operator shows that if \mathcal{A} is shift invariant then the quantity $e^{-\lambda\xi} \mathcal{A} e^{\lambda\xi}$ is constant for all ξ . Thus, neglecting boundary conditions for now, $\{e^{\lambda\xi}\}$ are eigenfunctions of \mathcal{A} . If the domain is normalised to $[-1, 1]$ and has periodic boundaries, then $\lambda = l\pi i$ for integer l . In this case the corresponding eigenfunctions form an orthogonal basis and the operator is diagonalised by the known eigenfunctions.

In the following, we assume that we are studying a shift-invariant system on a periodic domain $\xi \in [-1, 1]$. We move from the continuous formulation to a discretized space and assume that we are studying a function $u(\xi, t)$ and have access to evenly spaced samples of u at $\xi = [-1, -1 + \Delta\xi, \dots, 1 - \Delta\xi]$. We discuss the case of non-equally spaced samples in appendix A.3. If we define the state variable as $\mathbf{x}(t) = [u(\xi_1, t), u(\xi_2, t), \dots, u(\xi_n, t)]$ then the discrete-space linear operator \mathbf{A} that generates $\mathbf{x}(t)$ is diagonalized by its eigenvectors:

$$\mathbf{A} = \mathcal{F} \text{diag}(\hat{\mathbf{a}}) \mathcal{F}^{-1} \quad (15)$$

| system | physics (matrix structure) | data | standard DMD | physics-informed DMD |
|--|--------------------------------------|------|------------------|----------------------|
| convection-diffusion $\frac{\partial u}{\partial t} = v(\xi) \frac{\partial u}{\partial \xi} + \frac{\partial^2 u}{\partial \xi^2}$ | local (tri-diagonal) | | spectrum | spectrum |
| Schrödinger equation $i\hbar \frac{d}{dt} \Psi(t)\rangle = \hat{H} \Psi(t)\rangle$ | self-adjoint (Hermitian) | | spectrum | spectrum |
| cylinder flow $\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \frac{1}{Re} \nabla^2 \mathbf{u} - \nabla p$ | conservative (unitary) | | spectrum | spectrum |
| 2D advection $\frac{\partial \mathbf{u}}{\partial t} = \mathbf{c} \cdot \nabla \mathbf{u}$ | shift-invariant (block circulant) | | spectrum | spectrum |
| Volterra integro-differential equation $\frac{\partial u}{\partial t} = \int_{-1}^{\xi} K(\xi, \nu) u(\nu, t) d\nu$ | causal (upper triangular) | | spectrum | spectrum |
| channel flow $\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \frac{1}{Re} \nabla^2 \mathbf{u} - \nabla p$ | shift-invariant (block circulant) | | spectrum | spectrum |

Figure 3: A comparison of the models learned by exact DMD and piDMD for a range of applications. piDMD identifies the spectrum of the true operator with higher accuracy than exact DMD. The structure of the model matrices are also illustrated; piDMD models are generally more coherent than those learned by exact DMD. Details are given in the corresponding sections. In the spectrum subplots, we plot the true eigenvalues as \circ , the DMD eigenvalues as Δ , and the piDMD eigenvalues as \times .

where $\mathcal{F}_{j,k} = e^{2\pi i(j-1)(k-1)/n}/\sqrt{n}$ and $\mathcal{F}^{-1} = \mathcal{F}^*$ and $\{\hat{a}_j\}$ are the unknown eigenvalues. Equation (15) is equivalent to stating that \mathbf{A} is circulant:

$$\mathbf{A} = \begin{bmatrix} a_0 & a_{n-1} & \dots & a_1 \\ a_1 & a_0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & a_{n-1} \\ a_{n-1} & \dots & a_1 & a_0 \end{bmatrix}, \quad (16)$$

i.e. $\mathbf{A}_{j,k} = a_{(j-k) \bmod n}$. If the boundary conditions are instead Dirichlet or Neumann then \mathbf{A} is not circulant but Toeplitz; we solve the corresponding Procrustes problem in §A.5.

Substituting (15) into (3) and noting that the Frobenius norm is invariant to unitary transformations allows (3) to be transformed to

$$\operatorname{argmin}_{\hat{\mathbf{a}}} \|\operatorname{diag}(\hat{\mathbf{a}})\mathbf{X} - \mathbf{Y}\|_F, \quad (17)$$

where $\mathbf{X} = \mathcal{F}^* \mathbf{X}$ and $\mathbf{Y} = \mathcal{F}^* \mathbf{Y}$ are the spatial discrete Fourier transforms of \mathbf{X} and \mathbf{Y} . As such, \mathbf{X} and \mathbf{Y} can be efficiently formed in $O(mn \log(n))$ operations using Fast Fourier Transform (FFT) methods [70].

The rows of the cost function (17) now decouple to produce n minimization problems:

$$\operatorname{argmin}_{\hat{a}_j} \|\hat{a}_j \tilde{\mathbf{X}}_j - \tilde{\mathbf{Y}}_j\|_F \quad \text{for } 1 \leq j \leq n, \quad (18)$$

where $\tilde{\mathbf{X}}_j$ and $\tilde{\mathbf{Y}}_j$ are the j th rows of \mathbf{X} and \mathbf{Y} , respectively. An alternative view of this step is that spatial wavenumbers decouple in shift-invariant systems, so we may analyse each wavenumber individually. The optimal solution for each eigenvalue follows as

$$\hat{a}_j = \tilde{\mathbf{Y}}_j \tilde{\mathbf{X}}_j^\dagger = \tilde{\mathbf{Y}}_j \tilde{\mathbf{X}}_j^* / \|\tilde{\mathbf{X}}_j\|_2^2. \quad (19)$$

In principle, standard DMD methods [19, 21] could be encouraged to respect shift invariance by augmenting the data matrices with their shifted counterparts. However, one would strictly need to include n shifts, thus producing nm samples which is usually too large in applications. In contrast, the above solution is extremely efficient via the use of the FFT and the decoupling of wavenumbers.

The appendices include alternative solutions for cases where the system is shift-invariant and symmetric, skew-symmetric or unitary (§A.1), low rank (§A.2), the samples are not equally spaced (§A.3), the total least squares case (§A.4), and when the system is Toeplitz or Hankel (§A.5).

4.1.1 Example: plane channel flow

Our next example is the incompressible flow inside a plane channel of size $2\pi \times 2 \times 2\pi$ along the spatial x , y , and z coordinates that indicate the streamwise, wall-normal, and spanwise directions, respectively. The configuration considers a Reynolds number of $Re = 2000$ based on the channel half-height and the centerline velocity, and periodic boundary conditions in the open faces of the channel, hence the flow is homogeneous in the x and z directions. We use the spectral code *Channelflow* [71] to perform direct numerical simulations (DNS) with the same numerical configurations as that presented in [68].

The dataset investigated is generated from a DNS of the response of the laminar flow with parabolic velocity profile to a localized perturbation in the wall-normal velocity component of the form

$$v(x, y, z, 0) = \left(1 - \frac{r^2}{c_r^2}\right) (\cos(\pi y) + 1) e^{(-r^2/c_r^2 - y^2/c_y^2)}, \quad (20)$$

where $r^2 = (x - \pi)^2 + (z - \pi)^2$, the parameters are set to $c_r = 0.7$ and $c_y = 0.6$, and the amplitude of the perturbation was scaled to have an energy-norm of 10^{-5} to ensure that the effect of non-linearity is negligible. This initial condition, which was first studied in [72] and also in [68], is a model of a disturbance that could be generated in experiments using a spanwise and streamwise

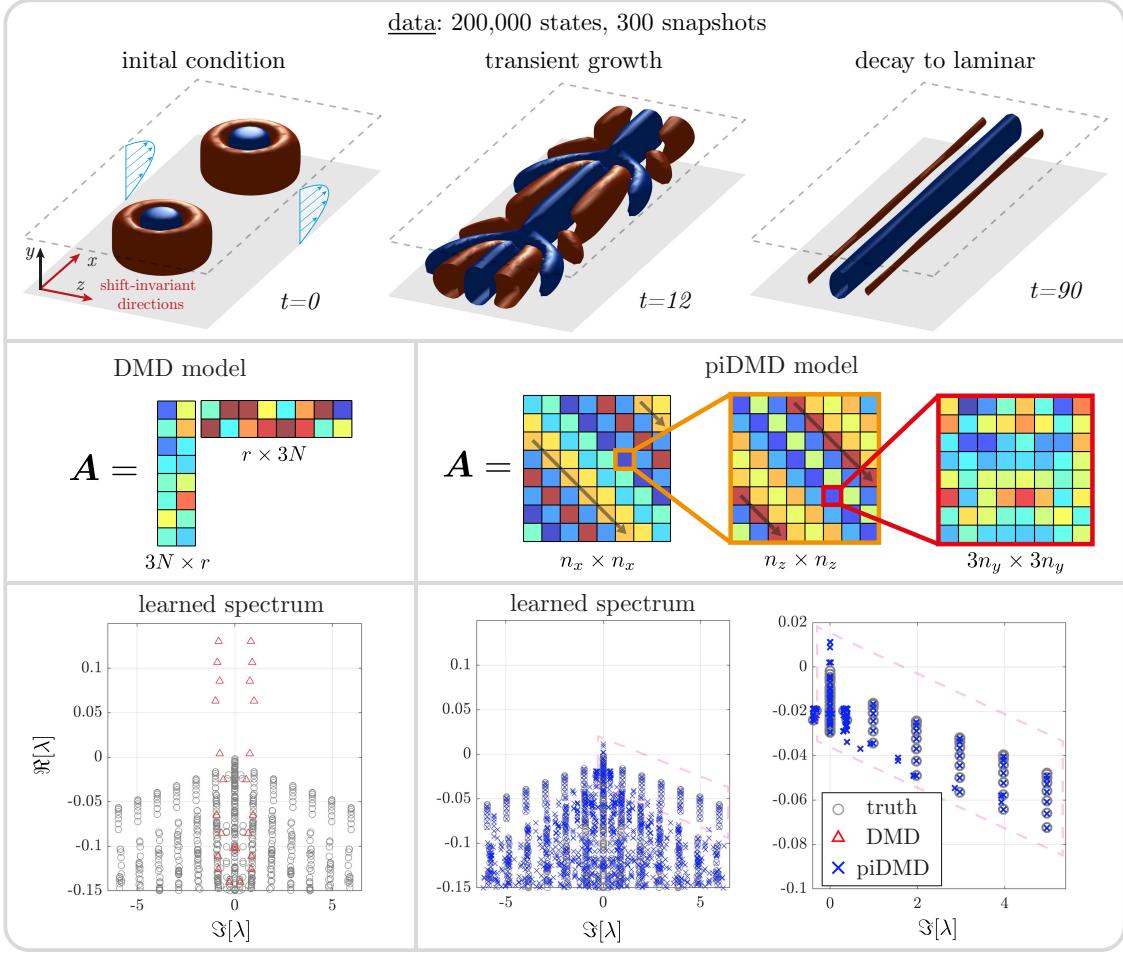


Figure 4: Learning the spectrum of the linearized Navier–Stokes equations from velocity measurements of the response to a localized disturbance in a channel flow. piDMD embeds the shift-invariant structure of the Navier–Stokes equations into the model learning framework, thus learning the spectrum of the linearized operator with improved accuracy over exact DMD.

periodic array of axisymmetric jets injecting fluid perpendicular to the wall. The dataset obtained from this simulation is then comprised of a sequence of 300 snapshots of the three-dimensional velocity-perturbation field recorded every 0.5 time units.

Our aim is to learn the spectrum of the underlying linear operator that *best* describes the evolution of the DNS snapshots. When directly applying exact DMD to this dataset, the algorithm attempts to find global modes for the three-dimensional velocity field and their spectrum. This is challenging because modes associated to different streamwise and spanwise wavenumbers may be mixed, leading to spurious eigenvalues, as shown in figure 4. However, we know that the flow is homogeneous in the x and z coordinates, and is therefore shift-invariant in these directions. Hence, we can leverage piDMD to incorporate this property, forcing the resulting matrix to respect, by construction, the structure with three nested levels shown in figure 4. In practice, this amounts to reshaping every snapshot, taking the FFT in the dimensions corresponding x and z , and performing DMD on the y -dependent Fourier amplitudes for every streamwise and spanwise

wavenumber tuple. Although simple, this has a tremendous impact on the quality of the learned spectrum. Shift-invariant piDMD produces modes that are associated with a single wavenumber tuple, which forces the modes to respect the spatial periodicity of the underlying system.

The spectra learned from both approaches are compared in figure 4, showing that piDMD results in a far more accurate eigenvalue spectrum. The DMD calculation considers a truncated SVD of rank 200. The piDMD results were obtained with a rank truncation of 50 at every wavenumber tuple.

4.2 Conservative systems

Conservation laws are foundational to all areas of science. Identifying quantities that remain constant in time – such as mass, momentum, energy, electric charge, and probability – allow us to derive universal governing equations, apply inductive reasoning, and understand basic physical mechanisms. In this section we demonstrate how conservation laws can be incorporated into the DMD framework.

Suppose that we are studying a system that we know conserves energy. In applications of DMD, it is implicitly assumed that measurements of the state have been suitably weighted so that the square of the 2-norm corresponds to the energy of the state [68]: $E(\mathbf{x}) = \|\mathbf{x}\|_2^2$. In these variables, the original optimization problems (2, 3) equate to finding the model \mathbf{A} that minimises the energy of the error between the true and predicted states (\mathbf{y}_k and \mathbf{Ax}_k respectively). Thus, if \mathbf{A} represents a discrete-time linear dynamical system ($\mathbf{y}_k = \mathbf{x}_{k+1} = \mathbf{Ax}_k$) then \mathbf{A} is an energy preserving operator if and only if

$$E(\mathbf{Ax}) = \|\mathbf{Ax}\|_2^2 = \|\mathbf{x}\|_2^2 = E(\mathbf{x}) \quad \text{for all } \mathbf{x} \in \mathbb{R}^n. \quad (21)$$

In words, (21) states that \mathbf{A} does not change the energy of the system but merely redistributes energy between the states. In matrix terminology, (21) holds if and only if \mathbf{A} is *unitary*. Therefore, for conservative systems the optimization problem (3) is the orthogonal Procrustes problem described in section 2.2, and the solution is given by (13). The eigenvalues of \mathbf{A} lie on the unit circle, and the eigenvectors are orthogonal. Thus, the eigenvectors oscillate in time, with no growth or decay. Since the solution of the orthogonal Procrustes problem (13) requires the full SVD of an $n \times n$ matrix, it can be more computationally efficient to first project onto the leading POD modes and build a model therein. In this case, the model is only energy preserving within the subspace spanned by the leading POD modes.

Noise is a substantial problem for most DMD methods, and a common remedy is to phrase the DMD optimization (2) as a total least squares problem [34]. Perhaps surprisingly, the solution to the orthogonal Procrustes problem (13) is also the solution to the total least squares problem when the solution is constrained to be orthogonal [43]. Thus, the solution (13) is optimal even when there is noise in both \mathbf{X} and \mathbf{Y} , as long as the noise has the same distribution in \mathbf{X} and \mathbf{Y} . There are various problems closely related to the orthogonal Procrustes problem, including the case where \mathbf{A} lies on the Steifel manifold (i.e., when \mathbf{A} is rectangular with orthonormal columns [48]), there is a weighting matrix [47], missing data [42], or high amounts of noise [46].

4.2.1 Example: $Re = 100$ flow past a cylinder

We now apply this energy-preserving piDMD to study the flow past a cylinder, which is a benchmark problem for modal decompositions [22]. The Reynolds number is $Re = DU/\nu = 100$ where

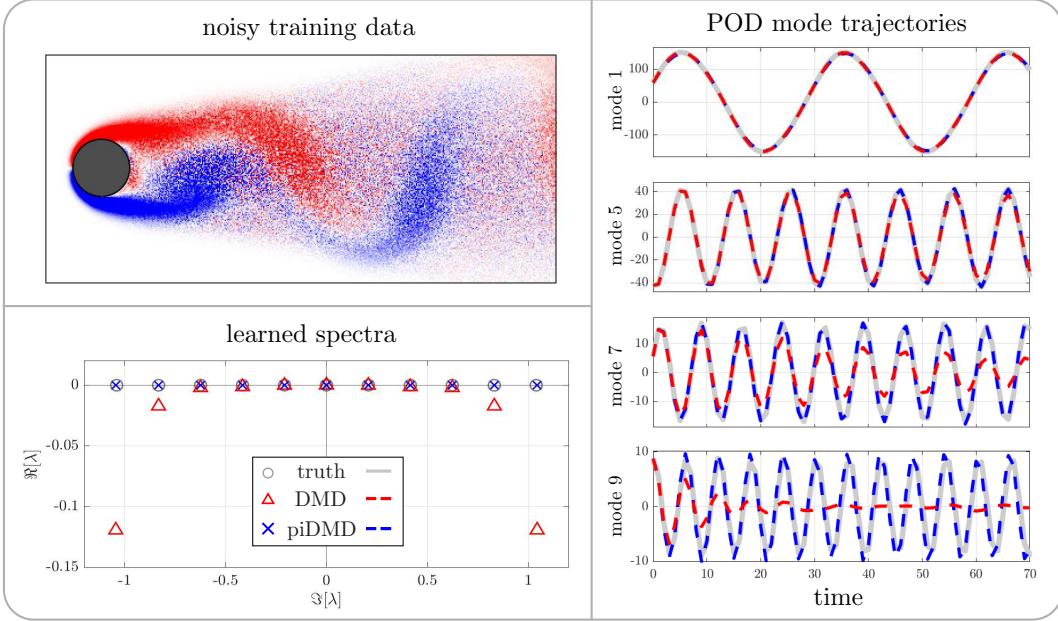


Figure 5: Learning the spectrum of the flow past a cylinder with noisy measurements with DMD and energy-preserving piDMD.

D is the diameter of the cylinder, U is the free-stream velocity and ν is the kinematic viscosity. The data consist of 151 samples of vorticity measurements, corresponding to five periods of vortex shedding, taken on a grid of 199×449 points. We note that the 2-norm of the measurements is approximately constant in time. The data are contaminated with 20% Gaussian noise, as illustrated in the top left panel of figure 5. We truncate the data to the first 15 POD modes and learn a standard DMD model, and an energy-preserving piDMD model.

piDMD learns a more accurate representation of the leading-order dynamics than standard DMD. For the high-order modes, the eigenvalues learned by DMD exhibit spurious damping whereas the eigenvalues learned by piDMD are purely oscillatory and remain on the imaginary axis. The trajectories of the POD coefficients are also more accurate for piDMD, and do not exhibit the spurious energy loss caused by noise.

4.3 Self-adjoint systems

Self-adjoint systems are another important class of linear systems that arise frequently in solid mechanics, heat and mass transfer, fluid mechanics, and quantum mechanics. When studying a system known to be self-adjoint we can use piDMD to constrain our model \mathbf{A} to lie in the manifold of symmetric (Hermitian) matrices, such that $\mathbf{A} = \mathbf{A}^*$. Symmetric matrices have real eigenvalues and, by the spectral theorem, are diagonalisable. This restriction places a significant constraint on the learned model and can substantially reduce the sensitivity of DMD to noise.

The minimum-norm solution of the symmetric Procrustes problem [73] is

$$\mathbf{A} = \mathbf{U}_X \mathbf{L} \mathbf{U}_X^*, \quad (22)$$

where the entries of \mathbf{L} are

$$\mathbf{L}_{i,j} = \overline{\mathbf{L}_{j,i}} = \begin{cases} \frac{\sigma_i \overline{\mathbf{C}_{j,i}} + \sigma_j \mathbf{C}_{i,j}}{\sigma_i^2 + \sigma_j^2} & \text{if } \sigma_i^2 + \sigma_j^2 \neq 0, \\ 0 & \text{otherwise,} \end{cases} \quad (23)$$

where $\mathbf{C} = \mathbf{U}_X^* \mathbf{Y} \mathbf{V}_X$ and $\mathbf{X} = \mathbf{U}_X \Sigma_X \mathbf{V}_X^*$ is the SVD of \mathbf{X} , and σ_i is the i th singular value. The solution of the skew-symmetric case is similar (§B.1), and algorithmic solutions are available when \mathbf{A} is instead constrained to be positive definite [44]. For large scale systems, a low-rank approximation to \mathbf{A} can be obtained by truncating the SVD of \mathbf{X} . If the SVD is truncated to rank r , then \mathbf{L} is an $r \times r$ matrix. This low-rank approximation preserves the self-adjointness of the model.

The symmetric Procrustes problem is connected to the orthogonal Procrustes problem (§4.2) via the fact that every unitary matrix \mathbf{U} can be expressed as the exponential of a Hermitian matrix \mathbf{H} as $\mathbf{U} = \exp(i\mathbf{H})$.

4.3.1 Learning energy states of the Schrödinger equation

Possibly the most famous self-adjoint operator is the quantum Hamiltonian, \hat{H} [74]. In quantum mechanics, physical quantities of interest, such as position, momentum, energy, and spin, are represented by self-adjoint linear operators called ‘observables’, of which one example is the Hamiltonian. The Hamiltonian describes the evolution of the probabilistic wave function via the time-dependent Schrödinger equation:

$$i\hbar \frac{d}{dt} |\Psi(t)\rangle = \hat{H} |\Psi(t)\rangle. \quad (24)$$

Solutions of (24) take the form of an eigenmode expansion

$$\Psi(\xi, t) = \sum_{j=1}^{\infty} \alpha_j \psi_j(\xi) e^{-iE_j t/\hbar}, \quad (25)$$

where E_j is an energy level (eigenvalue) of \hat{H} with corresponding eigenstate ψ_j , the coefficients α_j are determined by the initial distribution of the wave function, and ξ is the spatial coordinate.

We consider a wave function evolving according to the two-dimensional Schrödinger’s equation subject to an unknown Hamiltonian function. The unknown potential function is taken to be a finite well in a hexagonal shape. We collect measurements of the wave function and its velocity and train an exact DMD model and a piDMD model. In practice, empirical measurements of the wavefunction can be obtained through optical tomographic methods [75, 76] or sequential measurements of complementary variables [77]. The data matrices \mathbf{X} and \mathbf{Y} are formed from measurements of $|\Psi(t)\rangle$ and $i\hbar \frac{d}{dt} |\Psi(t)\rangle$ respectively, which are themselves formed by superposing eigenfunctions obtained by a finite difference method. Thus, the matrices \mathbf{X} and \mathbf{Y} are connected by an unknown self-adjoint Hamiltonian \hat{H} and we may apply our self-adjoint DMD optimization. The measurements of \mathbf{Y} are contaminated with Gaussian noise but the measurements of \mathbf{X} are clean.

Constraining the learned Hamiltonian to be self-adjoint reduces the sensitivity of the learning process to noise, as evidenced by figure 6. From the bottom left panel, we see that the energy levels

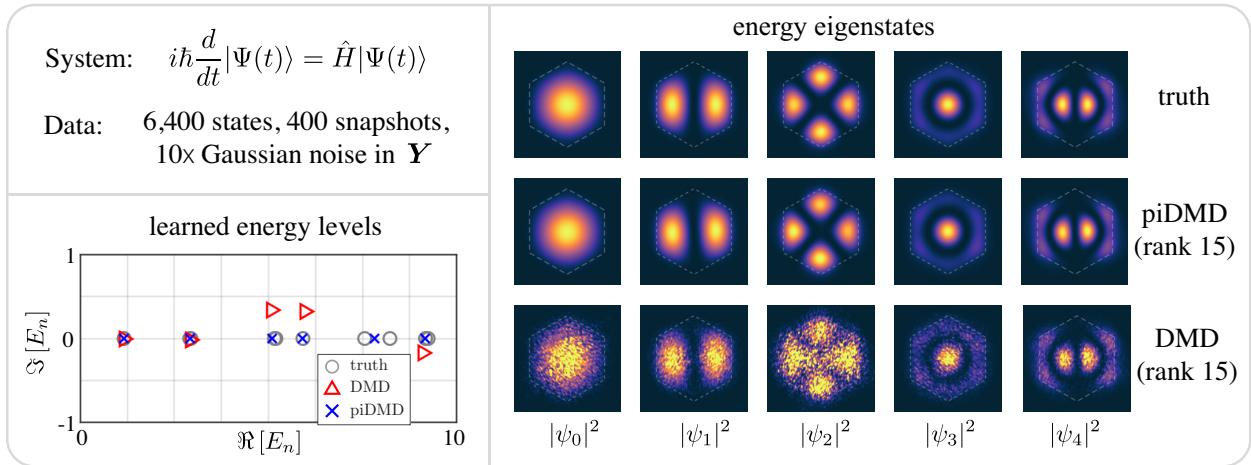


Figure 6: Learning the eigenstates of a quantum Hamiltonian of a hexagonal well with the self-adjoint piDMD.

of the piDMD model are physically consistent insofar as they are all real and positive. In contrast, the energy levels learned by standard DMD exhibit a spurious imaginary component, which will produce unrealistic growth/decay of the eigenstates. Both methods miss one eigenstate ($E_n \sim 7.5$) since this particular mode is not well-represented in the data. Additionally, the eigenstates learned by piDMD are much less noisy than those learned by standard DMD. This can be explained by noting that the eigenstates learned by piDMD lie in the span of \mathbf{X} , which is clean, whereas the eigenstates learned by standard DMD lie in the span of \mathbf{Y} , which is noisy. Lemma B.1 proves that the symmetric piDMD model (22) is less sensitive to noise than the exact DMD model (7). In summary, a physics-aware approach enables a more accurate identification of quantum energy levels and the associated eigenstates.

4.4 Spatially local systems

A hallmark of most physical systems is *spatial locality*. Points close to one another will usually have a stronger coupling than points that are far from one another. For example, basic physical processes such as convection and diffusion are spatially local. In most DMD applications, the data are collected from a spatial grid, but the grid rarely plays a role in the DMD analysis beyond forming the energy norm [68]. In fact, the output of the exact DMD algorithm remains invariant to unitary transformations on the rows and columns of the data matrices, so that randomly shuffling the rows (i.e., the spatial locations) will result in identical DMD models with the corresponding shuffling of the rows of the modes [57, 78]. Knowledge of the underlying grid from which the data were collected enables us to bias the learning process to prioritise relationships that we expect to have a strong local coupling.

Herein we consider the one-dimensional case; the analysis generalises straightforwardly to higher dimensions. We assume that the entries of \mathbf{x} are samples of a function u taken at n grid points $\{\xi_i\}$. The grid points can be arranged sequentially so that: $\xi_i < \xi_{i+1}$. By the spatial locality principle, we expect states that are spatially close to one another to have a stronger coupling than states that are spatially far from one another. For example, we may expect entries close to the

diagonal of \mathbf{A} to be larger than entries far from the diagonal. Then the entries of \mathbf{A} would satisfy

$$|A_{i,j}| \geq |A_{i,k}| \quad \text{for } |j - i| < |k - i|. \quad (26)$$

Equation (26) is merely a heuristic, and we would expect it to hold on average rather than for each (i, j, k) . Additionally, (26) is a difficult and expensive condition to implement in practice as it involves $\mathcal{O}(n^2)$ inequality constraints. Instead, we consider an alternative version of spatial locality where we only permit coupling between states that are sufficiently close:

$$|A_{i,k}| = 0 \quad \text{for } d < |k - i|. \quad (27)$$

Equation (27) describes a d -diagonal matrix. If $d = 0$ then states can only affect themselves and \mathbf{A} is diagonal (comparable to (17)). A more interesting case is if $d = 1$ and we only allow coupling between adjacent states. Then \mathbf{A} is tridiagonal: the entries of \mathbf{A} are zeros except the leading, upper and lower diagonals:

$$\mathbf{A} = \begin{bmatrix} \beta_1 & \gamma_1 & & & \\ \alpha_2 & \beta_2 & \gamma_2 & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & \gamma_{n-1} \\ & & & \alpha_n & \beta_n \end{bmatrix}. \quad (28)$$

We now solve the optimization problem (3) when \mathcal{M} is the manifold of tridiagonal matrices. Appendix C includes many solutions for more general diagonal-like structures including longer-range and variable coupling (§ C.1), spatially periodic local systems (§ C.2), self-adjoint (§ C.3), total-least squares (§ C.4), and weaker local structures (§ C.5). Since the rows of \mathbf{A} are decoupled, we can row-wise expand the Frobenius norm in (3) to obtain n smaller minimization problems:

$$\underset{\alpha_i, \beta_i, \gamma_i}{\operatorname{argmin}} \| \alpha_i \tilde{\mathbf{x}}_{i-1} + \beta_i \tilde{\mathbf{x}}_i + \gamma_i \tilde{\mathbf{x}}_{i+1} - \tilde{\mathbf{y}}_i \|_2 \quad \text{for } 1 \leq i \leq n, \quad (29)$$

with the convention $\alpha_1 = \gamma_n = 0$ and $\tilde{\mathbf{x}}_0 = \tilde{\mathbf{x}}_{n+1} = \mathbf{0}$. Recall that $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{y}}_i$ are the i -th rows of \mathbf{X} and \mathbf{Y} respectively (4). Each minimization problem (29) has the minimum-norm solution

$$[\alpha_i \quad \beta_i \quad \gamma_i] = \tilde{\mathbf{y}}_i \begin{bmatrix} \tilde{\mathbf{x}}_{i-1} \\ \tilde{\mathbf{x}}_i \\ \tilde{\mathbf{x}}_{i+1} \end{bmatrix}^\dagger \quad \text{for } 1 \leq i \leq n. \quad (30)$$

Each of the n minimizations costs $\mathcal{O}(m)$ so the total cost is $\mathcal{O}(mn)$.

4.4.1 Example: data-driven resolvent analysis

We consider a system of the form

$$u_t = \mathcal{A}u + f(\xi, t) \quad (31)$$

and aim to design a control strategy for the forcing f via input-output analysis [68, 79–81]. Fourier transforming in time and rearranging produces $\hat{u} = R_\omega \hat{f}$ where $R_\omega = (i\omega - \mathcal{A})^{-1}$ is the resolvent

of \mathcal{A} and the hats $\hat{\cdot}$ denote the Fourier transform in time. We wish to understand the form of the forcing \hat{f} that produces the largest response in \hat{u} . Formally, the optimal forcing (ϕ_1), response (ψ_1), and gain (σ_1) for each frequency component satisfy

$$\phi_1 = \underset{\|\phi_1\|_2=1}{\operatorname{argmax}} \|R_\omega \phi_1\|_2, \quad \sigma_1 = \|R_\omega \phi_1\|_2, \quad \psi_1 = \sigma_1^{-1} R_\omega \phi_1. \quad (32)$$

We may also investigate the higher-order singular triplets $(\sigma_j, \phi_j, \psi_j)$ that satisfy equivalent conditions with the additional requirement of orthogonality of $\{\phi_j\}$ and $\{\psi_j\}$. Equivalently, we wish to determine the Hilbert–Schmidt decomposition of R_ω [82]:

$$R_\omega = \sum_{j=1}^{\infty} \sigma_j \psi_j(\xi) \langle \phi_j, \cdot \rangle, \quad (33)$$

which is analogous to the SVD of the discretization of R_ω .

Recently, the authors have proposed ‘data-driven resolvent analysis’ as an efficient method for learning the singular triplets of a system purely from data measurements [68]. The method is analogous to DMD, except instead of computing the eigenmodes via eigendecomposition, we compute the resolvent modes via a singular value decomposition. In the following example, we demonstrate that piDMD can be integrated into data-driven resolvent analysis to provide physics-informed diagnostic information about the optimal forcings and responses of the system.

We consider the problem of designing a strategy to control the concentration of a solute that is governed by unknown dynamics. The solute concentration u is governed by the convection-diffusion operator $\mathcal{A}u = u_{\xi\xi} + a(\xi)u_\xi$ where a represents a non-constant convection term (here taken to be a random function) and boundary conditions $u_\xi(-1) = u(1) = 0$. A typical simulation is illustrated in the left panel of figure 7. The convection-diffusion operator is a canonical example of a non-normal system where the typical eigenvalue analysis fails to provide meaningful information [83, 84]. Our goal is to determine the forcings \hat{f} that produce the largest increase in solute in the frequency domain.

We solve the convection-diffusion equation in Chebfun [85] and collect 2,000 snapshots of u evaluated at 100 evenly spaced grid points. The snapshots matrices \mathbf{X} and \mathbf{Y} are formed from measurements of u and u_t respectively. We apply data-driven resolvent analysis with DMD and piDMD and compute the learned operators, resolvent norm, and forcing and response modes. For both algorithms, we approximate u_t from the measurements of u . The piDMD model is constrained to be tridiagonal, and the DMD model has rank 40. From the visualizations of the learned models in the center of figure 7, we see that the standard DMD identifies spurious relationships between the states, whereas piDMD, by construction, identifies a stronger coupling between adjacent states. The model learned by piDMD is reminiscent of the finite difference matrix; indeed, local piDMD may be viewed as a data-driven discovery of a finite difference stencil.

The benefit of piDMD is evidenced in its accurate identification of the resolvent modes (for $\omega = 1$), as seen in the right panel of figure 7. The true resolvent modes are computed in Chebfun (see section 12 of [86]). The modes learned by standard DMD are highly inaccurate for all j , and are completely incoherent for $j > 2$. In contrast, piDMD accurately determines the mode shapes, and associated gains for a range of frequencies. The physics-informed approach even determines the correct boundary conditions of the response and forcing modes, and resolves the complicated structures associated with the non-normality of the system. In particular, enforcing spatial locality enables piDMD to identify modes that are poorly represented by the data, which

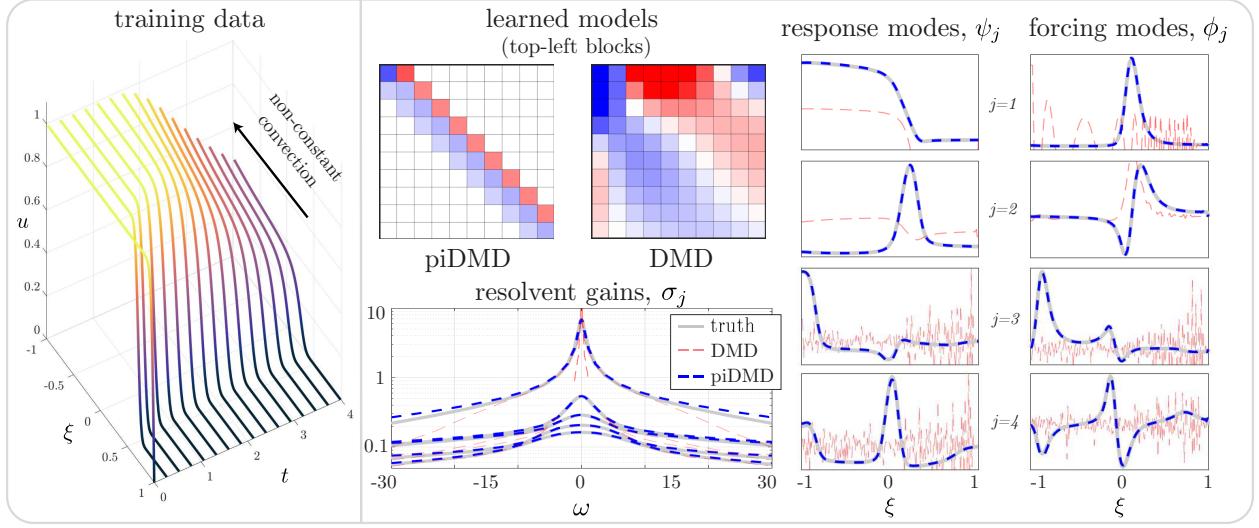


Figure 7: Data-driven resolvent analysis of the convection-diffusion equation for DMD and piDMD. The piDMD model is constrained to be tri-diagonal and therefore respects the spatial locality of the system. Incorporating this physical structure results a significant improvement in the identification of the resolvent norm, response modes, and forcing modes.

allows for diagnostic analysis that generalises outside of the training regime. Equipped with these data-driven modes, one could now design a low-rank control strategy for the solute concentration.

Analogously to finite difference methods, the accuracy of local piDMD involves a trade-off between the grid sizes in space and time. Future work should focus on clarifying these issues.

4.5 Causal systems

Causality is the process by which the behaviour of one state (the cause) influences another state (the effect). Many systems possess a spatially causal structure where each state depends only on ‘upstream’ states and are unaffected by ‘downstream’ states. Specifically, we may express the j th state as a function of the $n - j$ upstream states only: $y_j = f(x_j, x_{j+1}, \dots, x_n)$. For example, y_n depends only on x_n whereas y_1 depends only all the elements of \mathbf{x} . Accordingly, if we know that the system at hand possesses a spatially causal structure we may seek a piDMD model that is upper triangular. Such structures are typical of a system with causal features, for example, time-delay coordinates [87], feed-forward/strict-feedback systems [88, 89], and other control systems [90].

As in section 4.4, we may decouple the rows of \mathbf{A} in (3) and expand the cost function row-wise to obtain n smaller minimization problems:

$$\underset{\mathbf{A}_{j,j:n}}{\operatorname{argmin}} \|\mathbf{A}_{j,j:n} \mathbf{X}_{j:n,:} - \tilde{\mathbf{y}}_j\|_2 \quad \text{for } 1 \leq j \leq n. \quad (34)$$

where we have employed the MATLAB colon notation $j : n$ to indicate entries from j to n and $:$ to indicate entries from 1 to n . Here, $\tilde{\mathbf{y}}_j$ represents the j -th row of \mathbf{Y} . The (minimum norm) solution of (34) is obtained via the pseudoinverse of each sub-block of \mathbf{X} as

$$\mathbf{A}_{j,j:n} = \tilde{\mathbf{y}}_j \mathbf{X}_{j:n,:}^\dagger \quad \text{for } 1 \leq j \leq n. \quad (35)$$

Naively evaluating (35) for each j would require building n pseudoinverses for a total cost of $\mathcal{O}(n^2m \min(m, n))$. Fortunately, consecutive blocks of \mathbf{X} (e.g. $\mathbf{X}_{j-1:n,:}$ and $\mathbf{X}_{j:n,:}$) are related by a rank-1 update, so we can solve every problem in (35) for a total of $\mathcal{O}(nm \min(n, m))$ operations. (see appendix D.1 for further details). However, this approach is very numerically unstable when applied to realistic data with high condition number. Thus, in appendix D.2 we derive a more stable alternative.

We illustrate this causal structure on a simple example defined by the Volterra-type integro-differential equation

$$\frac{\partial u}{\partial t}(\xi, t) = \int_{-1}^{\xi} K(\xi, \nu)u(\nu, t), d\nu \quad -1 \leq \xi \leq 1. \quad (36)$$

Similar systems have been used to model transmission lines in neural networks during bursting activity [91] and the spread of disease in epidemics [92]. Taking $K(\xi, \nu) = \sqrt{1 - \xi^2}\sqrt{1 - \nu^2}$ and $u(\xi, 0) = \exp(-\xi^2)$, we numerically simulate (36) and construct \mathbf{X} and \mathbf{Y} from measurements of u only (velocity measurements are not used). The results of piDMD are illustrated in the penultimate row of figure 3, and indicate that piDMD is able to learn the leading eigenvalues while standard DMD fails. Note that the true spectrum here is continuous.

5 Conclusions and outlook

This work presents physics-informed dynamic mode decomposition (piDMD), a physics-aware modal decomposition technique that extracts coherent structures from high-dimensional time-series data. Rephrasing the DMD regression (2) as a Procrustes problem (3) shows that partially-known physics can be incorporated into the DMD framework by enforcing a matrix manifold constraint determined by known physics. We have applied piDMD to five of the most fundamental physical principles: conservation laws, self-adjointness, shift-invariance, locality, and causality. In several of these cases, we have presented new ‘exact’ solutions of the corresponding optimization problems. The examples presented demonstrate that piDMD can exhibit superior performance compared to classical DMD methods that do not account for the physics of the system.

The framework developed herein offers exciting new opportunities and challenges for data-driven dynamical systems, numerical linear algebra, convex optimization, and statistics. Fortunately, many research advances developed for classical DMD can also be leveraged to enhance piDMD. We conclude the paper with a brief discussion of the limitations of piDMD, and suggest promising future research directions.

5.1 Limitations and extensions

Each of the four steps of the piDMD framework outlined in section 3 pose distinct challenges. In some scenarios where the physics is poorly understood, determining suitable physical laws to impose on the model (step 1, ‘modeling’) can be challenging. For problems with intricate geometries and multiple dimensions, interpreting the physical principle as a matrix manifold (step 2, ‘interpretation’) can be the roadblock, as the manifold can become exceedingly complicated. We have already commented on the challenges and opportunities posed by step 3 (‘optimization’), and suggested algorithmic solutions [51]. Finally, the typically large state dimension and number

of samples can obfuscate step 4 when the matrix manifold does not easily admit the desired diagnostics. For example, if \mathbf{A} has a complicated banded structure, as in local piDMD, it may not be amenable to a fast eigendecomposition or SVD.

The solutions presented herein exhibit various levels of sensitivity to noise. As is typical of DMD methods, most solutions are unbiased with respect to \mathbf{Y} , but can be quite sensitive to noise in \mathbf{X} . Some of the solutions we have presented are relatively insensitive (§4.2) or can be reformulated in the total least squares sense (§A.4, §C.4). A full characterization of the sensitivity of piDMD should be performed on a case-by-case basis for different manifold constraints, and will be the subject of future investigations.

The Procrustes problem (3) is not the only means of incorporating physical principles into the DMD process. Closely related to Procrustes problems are ‘nearest matrix’ problems [93] where, given a matrix $\hat{\mathbf{A}}$, we seek the closest matrix on some matrix manifold \mathcal{M} :

$$\underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|\mathbf{A} - \hat{\mathbf{A}}\|. \quad (37)$$

Equation (37) can be viewed as a special case of the Procrustes problem (3) where $\mathbf{Y} = \hat{\mathbf{A}}$ and $\mathbf{X} = \mathbf{I}$. As such, nearest matrix problems are generally easier to solve than Procrustes problems. The nearest matrix problem (37) can be used to project any DMD model onto the nearest physically consistent model, or within optimization routines, such as proximal gradient methods or constrained gradient descent, to constrain the solution to the feasible set.

As presented, piDMD strictly requires that the model \mathbf{A} lies on the chosen matrix manifold \mathcal{M} . In many applications, such as when the data are very noisy or the physical laws and/or constraints are only approximately understood, it may be more appropriate to merely ‘encourage’ \mathbf{A} toward \mathcal{M} . In such cases, the piDMD regression (3) becomes

$$\underset{\mathbf{A}}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{AX}\|_F + \lambda R(\mathbf{A}), \quad (38)$$

where the first term is the reconstruction loss, R is a physics-informed regularizer, and λ is a user-defined constant that tunes the relative importance of the reconstruction and regularization. For example, $R(\mathbf{A})$ could represent the distance between \mathbf{A} and \mathcal{M} , which can be computed by solving the nearest matrix problem (37). An example of such a problem is solved in appendix C.5.

DMD models may be trained in discrete or continuous time. In discrete time with a constant sampling rate, the snapshots in \mathbf{Y} are $\mathbf{y}_k = \mathbf{x}_{k+1}$, whereas in the continuous time we have $\mathbf{y}_k = \dot{\mathbf{x}}_k$. If the continuous time operator is \mathbf{A}_C , then the discrete time operator is $\mathbf{A}_D = \exp(\Delta t \mathbf{A}_C)$. In special cases, the matrices \mathbf{A}_C and \mathbf{A}_D may lie on the same manifold, such as the upper-triangular and circulant manifolds; however, in general, they do not. For example, if \mathbf{A}_C is tridiagonal, then \mathbf{A}_D is not generally tridiagonal. Accordingly, if only discrete-time measurements are available but the matrix manifold is imposed on the continuous-time operator, then the optimization problem to solve is

$$\underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|\mathbf{Y} - \exp(\Delta t \mathbf{A}) \mathbf{X}\|_F, \quad (39)$$

which is a much more difficult optimization problem.

Acknowledgements

P.J.B. acknowledges insightful conversations with Suvrit Sra, Tasuku Soma and Andrew Horning. S.L.B. acknowledges valuable discussions with Jean-Christophe Loiseau. The authors acknowledge support from the Army Research Office (ARO W911NF-17-1-0306) and the National Science Foundation AI Institute in Dynamic Systems (Grant No. 2112085).

Appendices

The following appendices provide further details of the optimization problems considered in this paper.

A Further details of shift-invariant systems

Here we detail several extensions to the shift-invariant DMD problem introduced in section 4.1. Figure 8 compares these alternative solutions for training data from the 2-D advection equation.

A.1 Symmetric, skew-symmetric and unitary circulant matrices

Now consider the case where \mathbf{A} is symmetric and circulant. An application of (15) reveals that a circulant matrix is symmetric if and only if its eigenvalues $\{\hat{\mathbf{a}}_i\}$ are real. Thus, the minimization problem is

$$\underset{\Im[\hat{\mathbf{a}}]=\mathbf{0}}{\operatorname{argmin}} \|\operatorname{diag}(\hat{\mathbf{a}})\mathbf{X} - \mathbf{Y}\|_F^2. \quad (40)$$

Expanding the Frobenius norm row-wise and enforcing stationarity yields the optimal value as

$$\hat{a}_j = \frac{\Re \left[\tilde{\mathbf{y}}_j \tilde{\mathbf{x}}_j^* \right]}{\left\| \tilde{\mathbf{x}}_j \right\|_2^2}. \quad (41)$$

The case of a skew-symmetric circulant matrix is similar except the eigenvalues are now imaginary; the solution is

$$\hat{a}_j = i \frac{\Im \left[\tilde{\mathbf{y}}_j \tilde{\mathbf{x}}_j^* \right]}{\left\| \tilde{\mathbf{x}}_j \right\|_2^2}. \quad (42)$$

Finally, for a unitary circulant matrix (when the eigenvalues lie on the unit circle), we have

$$\hat{a}_j = \frac{\tilde{\mathbf{y}}_j \tilde{\mathbf{x}}_j^*}{\left\| \tilde{\mathbf{y}}_j \tilde{\mathbf{x}}_j^* \right\|_2}. \quad (43)$$

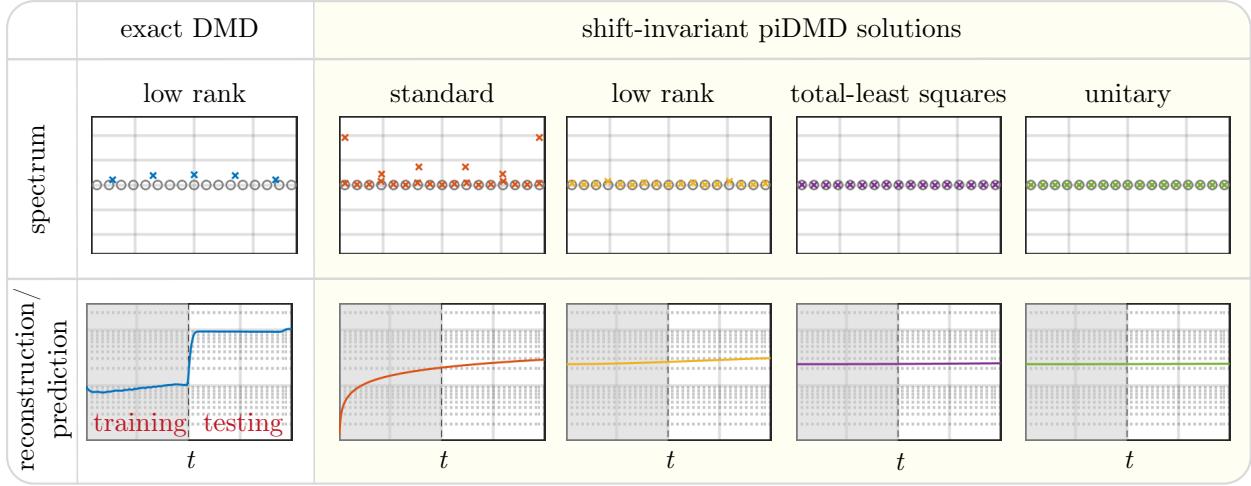


Figure 8: Comparison of exact DMD to shift-invariant piDMD and some variants. The model are trained on snapshots of a solution to the 2-D advection equation. The snapshots are contaminated by 2% Gaussian noise. The low-rank model has rank 45.

A.2 Best rank- r circulant approximant

Now consider the problem where \mathbf{A} is both circulant and rank $r \leq n$. Since the rank of a matrix is the number of non-zero eigenvalues, the rank of a circulant matrix is the number of non-zero entries in $\hat{\mathbf{a}}$. Thus, only r entries in $\hat{\mathbf{a}}$ are permitted to be non-zero and the minimization problem (17) becomes

$$\underset{\|\hat{\mathbf{a}}\|_0=r}{\operatorname{argmin}} \|\operatorname{diag}(\hat{\mathbf{a}})\mathbf{X} - \mathbf{Y}\|_F. \quad (44)$$

Equivalently, given the solution (15), we must select r eigenvalues to maintain and set all other eigenvalues to zero. To minimise the cost function, the r chosen eigenvalues should contribute to the largest reduction in the cost function. The extent to which each eigenvector minimises the cost function is characterised by their residual:

$$\rho_j = \sqrt{\|\tilde{\mathbf{y}}_j\|_2^2 - \|\tilde{\mathbf{x}}_j - \hat{a}_j \tilde{\mathbf{x}}_j\|_2^2} = \frac{\Re[\tilde{\mathbf{y}}_j \tilde{\mathbf{x}}_j^*]}{\|\tilde{\mathbf{x}}_j\|_2}. \quad (45)$$

A large residual indicates a large reduction in the cost function and vice versa. Therefore, the r non-zero eigenvalues to maintain are the r with the largest residuals.

A.3 Learning shift-invariant operators with non-equispaced samples

Here we extend the method of section 4.1 to allow spatial samples that are not equally spaced. We begin by formulating the Procrustes problem (3) in continuous space. Suppose that we have access to measurements of functions $u(\xi, t)$ and $v(\xi, t)$ for all $\xi \in [-1, 1]$ but at discrete points $t \in \{t_j | j = 1, \dots, m\}$. Then, the continuous-space piDMD regression (3) is

$$\underset{\mathcal{A} \in \mathcal{M}}{\operatorname{argmin}} \sum_{j=1}^m \|v(\xi, t_j) - \mathcal{A}u(\xi, t_j)\|_2^2, \quad (46)$$

where the continuous 2-norm is

$$\|f(\xi)\|_2^2 = \int_{-1}^1 |f(\xi)|^2 d\xi. \quad (47)$$

We now constrain \mathcal{A} to be a shift-invariant operator on a periodic domain. Such operators are diagonalized by exponential functions (§4.1) so we may write

$$\mathcal{A}f(\xi) = \sum_{k=-\infty}^{\infty} \hat{a}_k e^{i\pi k \xi} \hat{f}(k), \quad (48)$$

where $\hat{f}(k)$ is the finite Fourier transform of f at wavenumber k :

$$\hat{f}(k) = \int_{-1}^1 f(\eta) e^{-i\pi k \eta} d\eta. \quad (49)$$

Furthermore, if u and v are sufficiently smooth in ξ (e.g. they are Hölder continuous [94]) then we may write

$$u(\xi, t) = \sum_{k=-\infty}^{\infty} \hat{u}(k, t) e^{i\pi k \xi}, \quad v(\xi, t) = \sum_{k=-\infty}^{\infty} \hat{v}(k, t) e^{i\pi k \xi}. \quad (50)$$

Applying Parseval's theorem in concert with (48) and (50) shows that the minimization problem (46) is equivalent to

$$\operatorname{argmin}_{\hat{a}_k} \sum_{j=1}^m \left\| \sum_{k=-\infty}^{\infty} (\hat{v}(k, t_j) - \hat{a}_k \hat{u}(k, t_j)) e^{i\pi k \xi} \right\|_2^2 = \operatorname{argmin}_{\hat{a}_k} \sum_{k=-\infty}^{\infty} \sum_{j=1}^m |\hat{v}(k, t_j) - \hat{a}_k \hat{u}(k, t_j)|^2. \quad (51)$$

Furthermore, writing $\hat{\mathbf{u}} = [\hat{u}(k, t_1) \dots \hat{u}(k, t_m)]$ and $\hat{\mathbf{v}} = [\hat{v}(k, t_1) \dots \hat{v}(k, t_m)]$ transforms the right side of (51) to

$$\operatorname{argmin}_{\hat{a}_k} \sum_{k=-\infty}^{\infty} \|\hat{\mathbf{v}}(k) - \hat{a}_k \hat{\mathbf{u}}(k)\|_2^2,$$

whose solution is

$$\hat{a}_k = \hat{\mathbf{v}}(k) \hat{\mathbf{u}}^\dagger(k) = \hat{\mathbf{v}}(k) \hat{\mathbf{u}}^*(k) / \|\hat{\mathbf{u}}(k)\|_2^2. \quad (52)$$

Thus, (52) solves the continuous-space shift-invariant piDMD regression (46). Note that the solution (52) depends only on the Fourier coefficients of u and v (50).

Now, suppose that we only have measurements of u and v at arbitrary discrete grid points $\xi \in \{\xi_l\}$. We can then approximate the solution (52) by approximating the coefficients of the Fourier coefficients $\hat{\mathbf{v}}(k)$ and $\hat{\mathbf{u}}(k)$. The samples at $\{\xi_l\}$ correspond to the quadrature rule for the Fourier coefficients (49)

$$\hat{f}(k) \approx \sum_{l=1}^n w_l f(\xi_l) e^{-i\pi k \xi_l} \quad (53)$$

where $w_k = (\xi_{k+1} - \xi_{k-1})/2 \pmod{1}$ are the quadrature nodes for the grid. We can now use (53) to approximate $\hat{\mathbf{v}}$ and $\hat{\mathbf{u}}$ and thereby approximate the solution (52). The sums (53) can be computed efficiently with the non-uniform discrete Fourier transform (NUDFT-II, [70, 95]).

A.4 Total least-squares circulant problem

A equivalent statement of the Procrustes problem (3) is

$$\underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|\mathbf{R}\|_F \quad \text{subject to } \mathbf{AX} = \mathbf{Y} + \mathbf{R}. \quad (54)$$

As such, the formulation of (3) assumes assumes that any errors are in \mathbf{Y} only. Alternatively, we can consider errors in both \mathbf{X} and \mathbf{Y} and reformulate (3) as

$$\underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|[\mathbf{R} \ \mathbf{S}]\|_F \quad \text{subject to } \mathbf{A}(\mathbf{X} + \mathbf{S}) = \mathbf{Y} + \mathbf{R} \quad (55)$$

This approach is usually more principled as there will be inevitable errors in measurements of \mathbf{X} and \mathbf{Y} . Here we show that we can adapt the circulant solution of §4.1 to account for noise in both sets of measurements. Again, by the invariance of the Frobenius norm with respect to unitary transformations, (55) is equivalent to

$$\underset{\hat{\mathbf{a}}}{\operatorname{argmin}} \|[\mathcal{R} \ \mathcal{S}]\|_F \quad \text{subject to} \quad \operatorname{diag}(\hat{\mathbf{a}})(\mathcal{X} + \mathcal{S}) = \mathcal{Y} + \mathcal{R} \quad (56)$$

where $\mathcal{R} = \mathcal{F}^* \mathbf{R}$ and $\mathcal{S} = \mathcal{F}^* \mathbf{S}$. Again, the rows decouple to produce n smaller problems:

$$\underset{\hat{a}_i}{\operatorname{argmin}} \left\| \begin{bmatrix} \tilde{\mathcal{R}}_i & \tilde{\mathcal{S}}_i \end{bmatrix} \right\|_F \quad \text{subject to} \quad \hat{a}_i (\tilde{\mathcal{X}}_i + \tilde{\mathcal{S}}_i) = \tilde{\mathcal{Y}}_i + \tilde{\mathcal{R}}_i. \quad (57)$$

for $1 \leq i \leq n$. These problems may now be solved in the total least squares sense [35, 96].

We will briefly characterize the scenario where the total least-squares circulant model is statistically optimal. By optimal here, we mean that the model converges to the true model in probability as the number of samples increases [97]. In the unconstrained case, it is known that this is the case when the entries of \mathbf{R} and \mathbf{S} are random variables with zero mean and the covariance matrix of $\operatorname{vec}([\mathbf{R} \ \mathbf{S}])$ is a multiple of the identity matrix [97]. Therefore, the solution of the total least squares problem (57) will be statistically optimal if

$$\operatorname{Var}(\mathcal{R}_{i,j}) = \operatorname{Var}(\mathcal{S}_{i,j}) \quad \text{and} \quad \operatorname{Cov}(\mathcal{R}_{i,j}, \mathcal{S}_{i,j}) = 0. \quad (58)$$

In words the circulant-TLS problem is statistically optimal if errors in \mathbf{R} and \mathbf{S} have equal variance for a given wavenumber and the errors between $\operatorname{Cov}(\mathcal{R}_{i,j}, \mathcal{S}_{i,j})$ are independent. These conditions are met if \mathbf{R} and \mathbf{S} are independent random variables of equal variance σ , in which case we have $\operatorname{Var}(\mathcal{R}_{i,j}) = \operatorname{Var}(\mathcal{S}_{i,j}) = \sigma^2$. Therefore, if the unconstrained total least squares model is statistically optimal then so is the circulant piDMD model.

A.5 Toeplitz and Hankel matrices

Here we adapt the optimization of section §4.1 to Toeplitz matrices, which take the form

$$\mathbf{A} = \begin{bmatrix} a_0 & a_{-1} & \cdots & a_{-(n-1)} \\ a_1 & a_0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & a_{-1} \\ a_{n-1} & \cdots & a_1 & a_0 \end{bmatrix}. \quad (59)$$

Similarly to the circulant case, $A_{i,j} = a_{i-j}$ but now the subscripts not are considered modulo n . The Procrustes problem for a Toeplitz matrix was considered by [98]. We produce an alternative solution based on the fast Fourier transform, analogous to the solution for the circulant Procrustes problem (§4.1). A Toeplitz matrix can be embedded in a circulant matrix as

$$C = \left[\begin{array}{c|c} \mathbf{A} & \tilde{\mathbf{A}} \\ \hline \tilde{\mathbf{A}} & \mathbf{A} \end{array} \right] \quad \text{where} \quad \tilde{\mathbf{A}} = \begin{bmatrix} 0 & a_{n-1} & \cdots & a_1 \\ a_{-(n-1)} & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & a_{n-1} \\ a_{-1} & \cdots & a_{-(n-1)} & 0 \end{bmatrix} \quad (60)$$

Similarly to §4.1, the circulant matrix C is diagonalized by the DFT matrix \mathcal{F} (which is now $2n \times 2n$). Thus, we may write $C = \mathcal{F}^* \text{diag}(\hat{\mathbf{a}}) \mathcal{F}$ and the minimization (3) becomes

$$\underset{\hat{\mathbf{a}}}{\operatorname{argmin}} \left\| \hat{\mathbf{I}}^* \text{diag}(\hat{\mathbf{a}}) \mathcal{X} - \mathbf{Y} \right\|_F^2 \quad (61)$$

where $\hat{\mathbf{I}} = \mathcal{F} \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix}$ and $\mathcal{X} = \mathcal{F} \begin{bmatrix} \mathbf{X} \\ \mathbf{0} \end{bmatrix}$. It can be shown that the optimal coefficients $\hat{\mathbf{a}}$ solve the linear system

$$\mathbf{H}\hat{\mathbf{a}} = \mathbf{d} \quad \text{where} \quad \mathbf{H} = \left(\hat{\mathbf{I}} \hat{\mathbf{I}}^* \odot \overline{\mathcal{X} \mathcal{X}^*} \right), \quad \mathbf{d} = \text{diag} \left(\hat{\mathbf{I}} \mathbf{Y} \mathcal{X}^* \right) \quad (62)$$

and \odot represents element-wise multiplication (i.e. the Hadamard product). Note that a similar optimization arises in the computation of the amplitudes of DMD modes [99]. While mathematically correct, numerical implementations of (62) can be difficult. Specifically, issues arise because the Gramian matrix $\mathcal{X} \mathcal{X}^*$ will typically be large and ill-conditioned. Thus, we expect that one could find more efficient algorithms to calculate \mathbf{A} , possibly by adapting methods for fast solutions of Toeplitz systems [100–102].

This approach also applies to Hankel matrices, which take the form

$$\mathbf{A} = \begin{bmatrix} a_0 & a_1 & \cdots & a_{n-1} \\ a_1 & a_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & a_{2n-3} \\ a_{n-1} & \cdots & a_{2n-3} & a_{2n-2} \end{bmatrix} \quad (63)$$

i.e. each ascending skew-diagonal is constant: $A_{i,j} = a_{i+j-2}$. Hankel matrices are upside-down Toeplitz matrices: if \mathbf{A} is Hankel, then $\mathbf{J}\mathbf{A}$ is Toeplitz, where \mathbf{J} is the upside-down identity matrix. Thus, the Procrustes problem for a Hankel matrix is equivalent to the for a Toeplitz matrix with a suitable flipping of the data.

B Further details of self-adjoint systems

B.1 Skew-symmetric systems

If A is constrained to be skew-symmetric then (22) and still holds and $C = U_X^* Y V_X$ but now L is now defined by

$$L_{i,j} = -\bar{L}_{j,i} = \begin{cases} \frac{-\sigma_i \bar{C}_{j,i} + \sigma_j C_{i,j}}{\sigma_i^2 + \sigma_j^2} & \text{if } \sigma_i^2 + \sigma_j^2 \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (64)$$

B.2 Sensitivity of symmetric piDMD

In light of the results of section 4.3, we now compare the sensitivities of the symmetric piDMD model (22) and exact DMD model (7). In particular, we will show that the elements of the piDMD model have lower variance than those of the corresponding exact DMD model. Higham [41] also investigated the sensitivity of the symmetric Procrustes problem, but from the perspective of numerical stability; here we take a statistical perspective. In this section only, we denote the exact DMD solution (7) as A^e and the symmetric piDMD solution (22) as A^s .

We assume that measurements of Y are contaminated with I. I. D. Gaussian noise. For now, we neglect noise in X , which could be incorporated via a total least square approach [34, 35]. By linearity, it is sufficient to consider $Y_{i,j} \sim \mathcal{N}(0, 1)$. Again, since exact DMD and symmetric piDMD are linear models, they are unbiased estimators:

$$\mathbb{E}(A_{i,j}^s) = \mathbb{E}(A_{i,j}^e) = 0.$$

We now compare the variance of the models. Extensive algebra produces explicit expressions for the variances of $A_{i,j}^e$ and $A_{i,j}^s$ as

$$\text{Var}(A_{i,j}^e) = \sum_{l=1}^r \frac{U_{j,l}^2}{\sigma_l^2}, \quad (65)$$

$$\text{Var}(A_{i,j}^s) = \frac{1}{2} \sum_{k=1}^r \sum_{l=1}^r \frac{(U_{i,k} U_{j,l} + U_{i,l} U_{j,k})^2}{\sigma_k^2 + \sigma_l^2}, \quad (66)$$

where U is the matrix of left singular vectors of X , $\{\sigma_l\}$ are the associated singular values, and $\text{rank}(X) = r$. Figure 9 compares these theoretical values to empirical values for random data. These results are in agreement with figure 6 and indicate that the variance of the piDMD model is smaller than that of the exact DMD model. We formalize this result in the following lemma and proof.

Lemma B.1. *If the measurements of Y are contaminated with I. I. D. Gaussian noise, the variance of each element of the symmetric piDMD model (22) is bounded above by the variance of the corresponding element of the exact DMD model (7):*

$$\text{Var}(A_{i,j}^s) \leq \text{Var}(A_{i,j}^e). \quad (67)$$

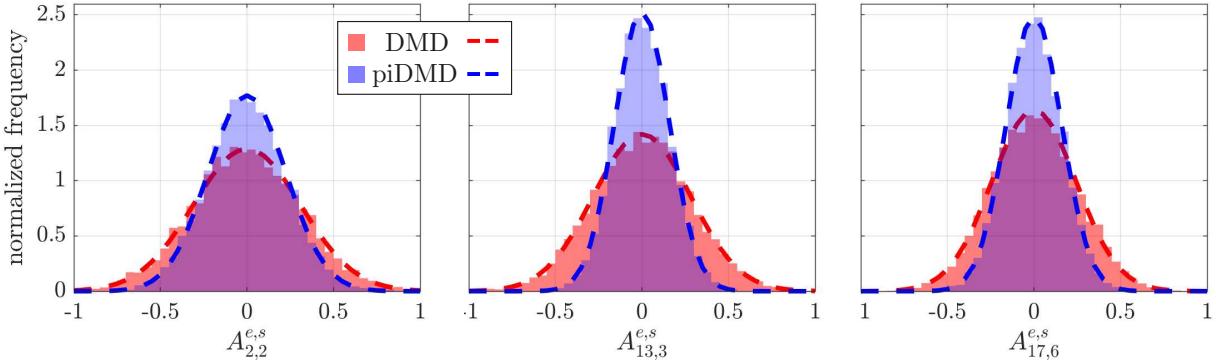


Figure 9: Distribution of errors in the entries of the learned models ($A_{2,2}^{e,s}$, $A_{13,3}^{e,s}$ and $A_{17,6}^{e,s}$) for exact DMD and self-adjoint piDMD. The block colors represent the frequency and the dashed lines are the theoretical curves (65, 66).

Proof. Sedrakyan's inequality [103] produces

$$\frac{(\mathbf{U}_{i,k}\mathbf{U}_{j,l} + \mathbf{U}_{i,l}\mathbf{U}_{j,k})^2}{\sigma_k^2 + \sigma_l^2} \leq \frac{\mathbf{U}_{i,k}^2\mathbf{U}_{j,l}^2}{\sigma_l^2} + \frac{\mathbf{U}_{i,l}^2\mathbf{U}_{j,k}^2}{\sigma_k^2}.$$

Therefore,

$$\text{Var}(\mathbf{A}_{i,j}^s) \leq \frac{1}{2} \sum_{k=1}^r \sum_{l=1}^r \left(\frac{\mathbf{U}_{i,k}^2\mathbf{U}_{j,l}^2}{\sigma_l^2} + \frac{\mathbf{U}_{i,l}^2\mathbf{U}_{j,k}^2}{\sigma_k^2} \right) = \sum_{k=1}^r \sum_{l=1}^r \frac{\mathbf{U}_{i,k}^2\mathbf{U}_{j,l}^2}{\sigma_l^2} = \left(\sum_{k=1}^r \mathbf{U}_{i,k}^2 \right) \text{Var}(\mathbf{A}_{i,j}^e). \quad (68)$$

The result now follows since the rows of \mathbf{U} have unit norm so $\sum_{k=1}^r \mathbf{U}_{i,k}^2 \leq 1$. \square

This result indicates that the symmetric piDMD model is less sensitive to noise than the exact DMD model when the noise is only in \mathbf{Y} . Future work should investigate the case when both \mathbf{X} and \mathbf{Y} are contaminated with noise, and clarify the effect of noise on the learned spectrum.

C Further details of local systems

C.1 More general local structures

The local solution of section 4.4 naturally extends to more general local structures. For example, it may be necessary to consider longer range interactions between states and allow, say, a state to depend on other states three grid points away. Additionally, there may be a variable local dependence where different states can depend on different numbers of elements. For example, a state near boundary (e.g. x_1 or x_n) may need to depend on more than one other state in order to properly resolve the boundary behaviour. To make this precise, suppose that state i depends on $l(i)$ states to the right and $u(i)$ states to the left above the leading diagonal (assuming that the data

is organised from left to right). Then, the solution for the optimal matrix entries is

$$\begin{bmatrix} A_{i-l(i),i} & \cdots & A_{i,i} & \cdots & A_{i+u(i),i} \end{bmatrix} = \tilde{\mathbf{y}}_i \begin{bmatrix} \tilde{\mathbf{x}}_{i-l(i)} \\ \vdots \\ \tilde{\mathbf{x}}_i \\ \vdots \\ \tilde{\mathbf{x}}_{i+u(i)} \end{bmatrix}^{\dagger} \quad \text{for } 1 \leq i \leq n \quad (69)$$

with the convention $A_{i,j} = 0$ if the entry does not exist and $\tilde{\mathbf{x}}_i = \mathbf{0}$ for $i < 1$ or $i > n$. We recover the tridiagonal solution (30) when $l(i) = u(i) = 1$.

C.2 Local and periodic

The tridiagonal solution in section 4.4 effectively assumed that the system had two endpoints that were disconnected from one another. However, the spatial domain could be periodic, in which case the first state can affect the n th state and vice versa. If we allow this structure then the matrix \mathbf{A} takes the form

$$\mathbf{A} = \begin{bmatrix} \beta_1 & \gamma_1 & & & \alpha_1 \\ \alpha_2 & \beta_2 & \gamma_2 & & \\ & \ddots & \ddots & \ddots & \\ & \alpha_3 & & & \\ & & \ddots & \ddots & \gamma_{n-1} \\ \gamma_n & & & \alpha_n & \beta_n \end{bmatrix}. \quad (70)$$

The solution for the optimal constants is identical to those stated in (30) except now the subscripts are evaluated modulo n .

C.3 Symmetric tridiagonal

Many spatially local systems are also self-adjoint. In such cases, we may enforce that \mathbf{A} is both tri-diagonal and symmetric:

$$\mathbf{A} = \begin{bmatrix} \alpha_1 & \beta_1 & & & \\ \beta_1 & \alpha_2 & \beta_2 & & \\ & \beta_2 & \ddots & \ddots & \\ & & \ddots & \ddots & b_{n-1} \\ & & & b_{n-1} & \alpha_n \end{bmatrix} \quad (71)$$

. The cost function (3) is then expressible as

$$\|\mathbf{Y} - \mathbf{AX}\|_F^2 = \|\tilde{\mathbf{y}}_1 - \alpha_1 \tilde{\mathbf{x}}_1 - \beta_1 \tilde{\mathbf{x}}_2\|_2^2 + \|\tilde{\mathbf{y}}_n - b_{n-1} \tilde{\mathbf{x}}_{n-1} - \alpha_n \tilde{\mathbf{x}}_n\|_2^2 \quad (72)$$

$$+ \sum_{i=2}^{n-1} \|\tilde{\mathbf{y}}_i - b_{i-1} \tilde{\mathbf{x}}_{i-1} - \alpha_i \tilde{\mathbf{x}}_i - b_i \tilde{\mathbf{x}}_{i+1}\|_2^2. \quad (73)$$

Enforcing stationarity in $\mathbf{c} = [\alpha_1 \ \cdots \ \alpha_n \ \beta_1 \ \cdots \ b_{n-1}]^T$ produces the linear system

$$\mathbf{T}\mathbf{c} = \mathbf{d} \quad (74)$$

where \mathbf{T} is the symmetric block-tridiagonal matrix

$$\mathbf{T} = \left[\begin{array}{c|c} \mathbf{T}_1 & \mathbf{T}_2 \\ \hline \mathbf{T}_2^* & \mathbf{T}_3 \end{array} \right]. \quad (75)$$

where the blocks are

$$\mathbf{T}_1 = \text{diag}(\|\tilde{\mathbf{x}}_1\|_2^2, \dots, \|\tilde{\mathbf{x}}_n\|_2^2), \quad (76)$$

$$\mathbf{T}_2 = \begin{bmatrix} \tilde{\mathbf{x}}_2^* \tilde{\mathbf{x}}_1 & & & & \\ \tilde{\mathbf{x}}_2^* \tilde{\mathbf{x}}_1 & \tilde{\mathbf{x}}_3^* \tilde{\mathbf{x}}_2 & & & \\ & \ddots & & & \\ & & \tilde{\mathbf{x}}_4^* \tilde{\mathbf{x}}_3 & & \\ & & & \ddots & & \\ & & & & \tilde{\mathbf{x}}_{n-1}^* \tilde{\mathbf{x}}_{n-2} & \tilde{\mathbf{x}}_n^* \tilde{\mathbf{x}}_{n-1} \\ & & & & & \tilde{\mathbf{x}}_n^* \tilde{\mathbf{x}}_{n-1} \end{bmatrix}, \quad (77)$$

$$\mathbf{T}_3 = \begin{bmatrix} \tilde{\mathbf{x}}_2^* \tilde{\mathbf{x}}_2 + \tilde{\mathbf{x}}_1^* \tilde{\mathbf{x}}_1 & \tilde{\mathbf{x}}_3^* \tilde{\mathbf{x}}_1 & & & \\ \tilde{\mathbf{x}}_3^* \tilde{\mathbf{x}}_1 & \tilde{\mathbf{x}}_3^* \tilde{\mathbf{x}}_3 + \tilde{\mathbf{x}}_2^* \tilde{\mathbf{x}}_2 & \tilde{\mathbf{x}}_4^* \tilde{\mathbf{x}}_2 & & \\ & \tilde{\mathbf{x}}_4^* \tilde{\mathbf{x}}_2 & \ddots & & \\ & & \ddots & \ddots & \\ & & & \ddots & \tilde{\mathbf{x}}_n^* \tilde{\mathbf{x}}_{n-1} \\ & & & & \tilde{\mathbf{x}}_n^* \tilde{\mathbf{x}}_{n-1} & \tilde{\mathbf{x}}_n^* \tilde{\mathbf{x}}_{n-1}, \end{bmatrix} \quad (78)$$

and

$$\mathbf{d} = [\tilde{\mathbf{x}}_1^* \tilde{\mathbf{y}}_1 \ \cdots \ \tilde{\mathbf{x}}_n^* \tilde{\mathbf{y}}_n \quad \tilde{\mathbf{x}}_1^* \tilde{\mathbf{y}}_2 + \tilde{\mathbf{x}}_2^* \tilde{\mathbf{y}}_1 \ \cdots \ \tilde{\mathbf{x}}_{n-1}^* \tilde{\mathbf{y}}_n + \tilde{\mathbf{x}}_n^* \tilde{\mathbf{y}}_{n-1}]^*.$$

Note that $\mathbf{T}_1 \in \mathbb{R}^{n \times n}$, $\mathbf{T}_2 \in \mathbb{R}^{n \times (n-1)}$, and $\mathbf{T}_3 \in \mathbb{R}^{(n-1) \times (n-1)}$. The system (74) can be solved very efficiently by exploiting the sparsity and symmetry of \mathbf{T} .

C.4 Local total least squares

This estimator is unbiased if there is only noise in \mathbf{Y} . If \mathbf{X} and \mathbf{Y} are both contaminated by noise then (29) can be suitably adapted and solved via total least squares [97], provided that the noise is normally distributed and has the same (diagonal) covariance matrix for \mathbf{X} and \mathbf{Y} .

C.5 Weaker spatial locality

A weaker version of spatial locality can also be enforced by penalising entries via

$$\underset{\mathbf{A}}{\text{argmin}} \|\mathbf{Y} - \mathbf{AX}\|_F + \lambda \|\mathbf{H} \odot \mathbf{A}\|_F \quad (79)$$

where \mathbf{H} is a matrix that penalises the solution for grid points that are spatially far from one another and λ is a regularization parameter. For example, a suitable choice for \mathbf{H} could be the Gaussian kernel matrix

$$\mathbf{H}_{i,j} = \exp\left(\|\xi_i - \xi_j\|_2/(2\sigma^2)\right) \quad (80)$$

for some parameter σ , where $\{\xi_i\}$ are the known grid points. For general \mathbf{H} , the solution to (79) is

$$\tilde{\mathbf{A}}_i = \tilde{\mathbf{y}}_i \mathbf{X}^* \left(\mathbf{X} \mathbf{X}^* + \text{diag}(\tilde{\mathbf{h}}_i) \right)^{-1}$$

for $1 \leq i \leq n$ where $\tilde{\mathbf{h}}_i$ is the i -th row of \mathbf{H} . Unfortunately, since the regularizer penalises every entry of \mathbf{A} , the implementation of (81) costs $\mathcal{O}(n^4)$ operations to compute the full model \mathbf{A} . Thus, we recommend pursuing inexact but efficient solutions, possibly using descent methods [104].

D Further details for causal systems

D.1 Updating equations for upper triangular system

Here we present an efficient implementation of the solution (35). Consecutive blocks of \mathbf{X} (e.g. $\mathbf{X}_{j-1:n,:}$ and $\mathbf{X}_{j:n,:}$) are related by the rank-1 update

$$\mathbf{X}_{j-1:n,:} = \begin{bmatrix} \mathbf{0} \\ \mathbf{X}_{j:n,:} \end{bmatrix} + \begin{bmatrix} 1 \\ \mathbf{0} \end{bmatrix} \tilde{\mathbf{x}}_{j-1}. \quad (81)$$

This observation enables us to efficiently calculate the pseudoinverse $\tilde{\mathbf{X}}_{j-1:n,:}^\dagger$ from $\tilde{\mathbf{X}}_{j:n,:}^\dagger$. The following formulas come from [105] (after some simplification). Let $\hat{\mathbf{x}}_{j-1} = \tilde{\mathbf{x}}_{j-1} (\mathbf{I} - \mathbf{X}_{j:n,:}^\dagger \mathbf{X}_{j:n,:})$ be the orthogonal projection of $\tilde{\mathbf{x}}_j$ onto the complement of the row space of $\mathbf{X}_{j:n,:}$. If $\tilde{\mathbf{x}}_{j-1}$ lies in the row space of $\mathbf{X}_{j:n,:}$ (i.e. $\|\hat{\mathbf{x}}_{j-1}\|_2 \neq 0$) then

$$\mathbf{X}_{j-1:n,:}^\dagger = \begin{bmatrix} \hat{\mathbf{x}}_{j-1}^\dagger & (\mathbf{I} - \hat{\mathbf{x}}_{j-1}^\dagger \tilde{\mathbf{x}}_{j-1}) \mathbf{X}_{j:n,:}^\dagger \end{bmatrix} \quad (82)$$

otherwise,

$$\mathbf{X}_{j-1:n,:}^\dagger = \begin{bmatrix} \mathbf{0} & \mathbf{X}_{j:n,:}^\dagger \end{bmatrix} (\mathbf{I} - \mathbf{v}_j^\dagger \mathbf{v}_j) \quad (83)$$

where $\mathbf{v}_j = [1 \ -\tilde{\mathbf{x}}_{j-1} \ \mathbf{X}_{j:n,:}^\dagger]$. Note that (82) and (83) can be formed using only matrix-vector products. It turns out that the above formulation is unstable for realistic data measurements (with high condition number). Below, we present an alternative algorithm that is as fast but more stable.

D.2 Alternative (more) stable solution

We now derive an alternative solution to the upper-triangular piDMD problem of section 4.5. In our experience, this version is more stable than that of the previous section. We use the economy RQ decomposition of \mathbf{X} and write $\mathbf{X} = \mathbf{R}\mathbf{Q}$ where $\mathbf{R} \in \mathbb{C}^{n \times n}$ is upper triangular, and $\mathbf{Q} \in \mathbb{C}^{n \times m}$ satisfies $\mathbf{Q}\mathbf{Q}^* = \mathbf{I}$. Note that

$$\|\mathbf{Y} - \mathbf{AX}\|_F^2 = \|\mathbf{Y}\|_F^2 - \|\mathbf{Y}\mathbf{Q}^*\|_F^2 + \|\mathbf{Y}\mathbf{Q}^* - \mathbf{AR}\|_F^2. \quad (84)$$

The first two terms in (84) are independent of \mathbf{A} and, by the submultiplicativity of the Frobenius norm², have a non-negative sum. Thus, the upper-triangular Procrustes problem may be phrased as

$$\underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{AX}\|_F = \underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|\mathbf{YQ}^* - \mathbf{AR}\|_F. \quad (85)$$

Now, since the product of upper triangular matrices is upper triangular, \mathbf{AR} is also upper triangular. Thus, the lower triangular component of \mathbf{YQ}^* is irrelevant to the minimization problem and we may write the problem as

$$\underset{\mathbf{A} \in \mathcal{M}}{\operatorname{argmin}} \|\operatorname{triu}(\mathbf{YQ}^*) - \mathbf{AR}\|_F, \quad (86)$$

where, in MATLAB notation, $\operatorname{triu}(\cdot)$ extracts the upper triangular portion of its argument. Since the rows of \mathbf{A} are independent, the Frobenius norm may be expanded row-wise and minimized separately. Thus, the rows of \mathbf{A} solve the n smaller minimization problems

$$\underset{\mathbf{A}_{i,:}}{\operatorname{argmin}} \|(YQ^*)_{i,i:n} - \mathbf{A}_{i,i:n} \mathbf{R}_i\|_F \quad \text{for } 1 \leq i \leq n, \quad (87)$$

where $\mathbf{R}_i = \mathbf{R}_{i:n,i:n} \in \mathbb{R}^{(n-i+1) \times n-i+1}$ is the i -th bottom right block of \mathbf{R} and $\mathbf{A}_{i,i:n} \in \mathbb{R}^{n-i+1}$ is the i -th row of \mathbf{A} with the $i-1$ leading zeros removed. Each minimization admits a unique solution if each \mathbf{R}_i is full rank, which occurs if the pivots of \mathbf{R} are non zero, i.e. \mathbf{X} is full rank. If the data is rank deficient, as is often the case in applications, there are infinitely many solutions. In any case, we seek the smallest solution (in the Frobenius norm sense), which is given by

$$\mathbf{A}_{i,i:n} = (YQ^*)_{i,i:n} \mathbf{R}_i^\dagger. \quad (88)$$

Calculating \mathbf{A} via (88) naively requires computing each pseudoinverse n times for a total of $\mathcal{O}(n^4)$ operations. As in §D.1, this complexity can be lowered with an updating procedure: following [105], we can efficiently compute \mathbf{R}_i^\dagger using \mathbf{R}_{i+1}^\dagger . Thus, we iterate through (88) backwards, starting at $i = n$. when the (scalar) pseudoinverse is given by $\mathbf{R}_{n,n}^\dagger = 1/\mathbf{R}_{n,n}$. Proceeding inductively, we note that the next matrix block can be expressed as a rank-one update to the previous block:

$$\mathbf{R}_i = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{i+1} \end{bmatrix} + \begin{bmatrix} 1 \\ \mathbf{0} \end{bmatrix} \mathbf{R}_{i,i:n}. \quad (89)$$

The updating formula takes one of two forms depending on whether the pivot $\mathbf{R}_{i,i}$ is smaller than some tolerance ϵ_{tol} close to machine precision:

Case 1: $|\mathbf{R}_{i,i}| > \epsilon_{\text{tol}}$

$$\mathbf{R}_i^\dagger = \begin{bmatrix} 1/\mathbf{R}_{i,i} & \mathbf{R}_{i,i:n}^\dagger \mathbf{R}_{i+1}^\dagger / \mathbf{R}_{i,i} \\ \mathbf{0} & \mathbf{R}_{i+1}^\dagger \end{bmatrix}. \quad (90)$$

²In particular, $\|\mathbf{CD}\|_F \leq \|\mathbf{C}\|_F \|\mathbf{D}\|_2$ for any matrices \mathbf{C} and \mathbf{D} . Note the 2-norm in the second term on the right.

Case 2: $|R_{i,i}| < \epsilon_{\text{tol}}$

$$\mathbf{R}_i^\dagger = \begin{bmatrix} 0 & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{i+1}^\dagger \end{bmatrix} \left(\mathbf{I} - \frac{\boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^*}{\|\boldsymbol{\alpha}_i\|_2^2} \right) \quad (91)$$

where $\boldsymbol{\alpha}_i = (\mathbf{R}_{i,i:n} \mathbf{R}_{i+1}^\dagger)^* - [1 \ 0]^T$. While mathematically correct, the accuracy of (90) and (91) can deteriorate over many iterations. Thus, in applications it may be sensible to reform the pseudoinverse \mathbf{R}_i^\dagger from scratch every so often.

The eigenvalues of \mathbf{A} can also be computed without forming \mathbf{A} explicitly. Note that if $R_{i,i} \neq 0$ then the upper-triangular structure of \mathbf{R}_i implies that $A_{i,i} = (\mathbf{Y} \mathbf{Q}^*)_{i,i} / R_{i,i}$. Further, the eigenvalues of a triangular matrix are its diagonal entries, so $\{A_{i,i}\}$ are the eigenvalues of \mathbf{A} (for $R_{i,i} \neq 0$) and can be computed without forming \mathbf{A} fully.

References

- [1] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nat. Rev. Phys.*, vol. 3, no. 6, pp. 422–440, 2021.
- [2] J. Behler and M. Parrinello, "Generalized neural-network representation of high-dimensional potential-energy surfaces," *Phys. Rev. Lett.*, vol. 98, no. 14, p. 146401, 2007.
- [3] S. Greydanus, M. Dzamba, and J. Yosinski, "Hamiltonian neural networks," in *Adv. Neural Inf. Process. Syst.*, vol. 32, 2019.
- [4] M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and Prabhat, "Deep learning and process understanding for data-driven Earth system science," *Nature*, vol. 566, no. 7743, pp. 195–204, 2019.
- [5] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, 2019.
- [6] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, "Lagrangian neural networks," *arXiv preprint arXiv:2003.04630*, 2020.
- [7] M. Raissi, A. Yazdani, and G. E. Karniadakis, "Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations," *Science*, vol. 367, no. 6481, pp. 1026–1030, 2020.
- [8] R. Wang, D. Maddix, C. Faloutsos, Y. Wang, and R. Yu, "Bridging physics-based and data-driven modeling for learning dynamical systems," *Proc. Mach. Learn. Res.*, vol. 144, pp. 385–398, 2020.
- [9] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis, "Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators," *Nat. Mach. Intell.*, vol. 3, no. 3, pp. 218–229, 2021.
- [10] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, "Fourier Neural Operator for Parametric Partial Differential Equations," in *9th Int. Conf. Learn. Represent.*, 2021.
- [11] P. J. Baddoo, B. Herrmann, B. J. McKeon, and S. L. Brunton, "Kernel Learning for Robust Dynamic Mode Decomposition: Linear and Nonlinear Disambiguation Optimization (LANDO)," *arXiv preprint arXiv:2106.01510*, 2021.
- [12] S. Klus, P. Gelß, F. Nüske, and F. Noé, "Symmetric and antisymmetric kernels for machine learning problems in quantum physics and chemistry," *Mach. Learn.: Sci. Technol.*, vol. 2, no. 4, 2021.
- [13] V. Shah, A. Joshi, S. Ghosal, B. Pokuri, S. Sarkar, B. Ganapathysubramanian, and C. Hegde, "Encoding Invariances in Deep Generative Models," *arXiv preprint arXiv:1906.01626*, 2019.
- [14] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proc. Natl. Acad. Sci.*, vol. 113, no. 15, pp. 3932–3937, 2016.

- [15] S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Data-driven discovery of partial differential equations," *Sci. Adv.*, vol. 3, no. 4, p. e1602614, 2017.
- [16] J. C. Loiseau and S. L. Brunton, "Constrained sparse Galerkin regression," *J. Fluid Mech.*, vol. 838, pp. 42–67, 2018.
- [17] Y. Guan, S. L. Brunton, and I. Novoselov, "Sparse nonlinear models of chaotic electroconvection," *Royal Society Open Science*, vol. 8, no. 8, p. 202367, 2021.
- [18] L. Zanna and T. Bolton, "Data-driven equation discovery of ocean mesoscale closures," *Geophysical Research Letters*, vol. 47, no. 17, p. e2020GL088376, 2020.
- [19] P. J. Schmid, "Dynamic mode decomposition of numerical and experimental data," *J. Fluid Mech.*, vol. 656, pp. 5–28, 2010.
- [20] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, "Spectral analysis of nonlinear flows," *J. Fluid Mech.*, vol. 641, pp. 115–127, 2009.
- [21] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz, "On dynamic mode decomposition: Theory and applications," *J. Comput. Dyn.*, vol. 1, no. 2, pp. 391–421, 2014.
- [22] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016.
- [23] T. Askham and J. N. Kutz, "Variable projection methods for an optimized dynamic mode decomposition," *SIAM J. Appl. Dyn. Syst.*, vol. 17, no. 1, pp. 380–416, 2018.
- [24] P. J. Schmid, "Dynamic Mode Decomposition and Its Variants," *Annu. Rev. of Fluid Mech.*, vol. 54, pp. 225–254, 2022.
- [25] J. L. Proctor and P. A. Eckhoff, "Discovering dynamic patterns from infectious disease data using dynamic mode decomposition," *International health*, vol. 7, no. 2, pp. 139–145, 2015.
- [26] E. Berger, M. Sastuba, D. Vogt, B. Jung, and H. B. Amor, "Estimation of perturbations in robotic behavior using dynamic mode decomposition," *Journal of Advanced Robotics*, vol. 29, no. 5, pp. 331–343, 2015.
- [27] I. Abraham and T. D. Murphey, "Active learning of dynamics for data-driven control using Koopman operators," *IEEE Transactions on Robotics*, vol. 35, no. 5, pp. 1071–1083, 2019.
- [28] B. W. Brunton, L. A. Johnson, J. G. Ojemann, and J. N. Kutz, "Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition," *Journal of Neuroscience Methods*, vol. 258, pp. 1–15, 2016.
- [29] A. Goldschmidt, E. Kaiser, J. L. DuBois, S. L. Brunton, and J. N. Kutz, "Bilinear dynamic mode decomposition for quantum control," *New J. Phys.*, vol. 23, no. 3, p. 033035, 2021.
- [30] Y. Susuki, I. Mezić, and T. Hikihara, "Coherent Swing Instability of Power Grids," *Journal of Nonlinear Science*, vol. 21, no. 3, pp. 403–439, 2011.
- [31] R. Taylor, J. N. Kutz, K. Morgan, and B. A. Nelson, "Dynamic mode decomposition for plasma diagnostics and validation," *Rev. Sci. Instrum.*, vol. 89, no. 5, p. 053501, 2018.
- [32] A. A. Kaptanoglu, K. D. Morgan, C. J. Hansen, and S. L. Brunton, "Characterizing magnetized plasmas with dynamic mode decomposition," *Physics of Plasmas*, vol. 27, p. 032108, 2020.
- [33] S. Bagheri, "Effects of weak noise on oscillating flows: Linking quality factor, Floquet modes, and Koopman spectrum," *Physics of Fluids*, vol. 26, no. 9, p. 094104, 2014.
- [34] S. T. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley, "Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition," *Exp. Fluids*, vol. 57, no. 3, p. 42, 2016.
- [35] M. S. Hemati, C. W. Rowley, E. A. Deem, and L. N. Cattafesta, "De-biasing the dynamic mode decomposition for applied Koopman spectral analysis," *TCFD*, vol. 31, no. 4, pp. 349–368, 2017.
- [36] J. P. Cunningham and Z. Ghahramani, "Linear dimensionality reduction: Survey, insights, and generalizations," *The Journal of Machine Learning Research*, vol. 16, no. 1, pp. 2859–2900, 2015.
- [37] J. R. Hurley and R. B. Cattell, "The Procrustes Program: Producing direct rotation to test a hypothesized factor structure," *Behav. Sci.*, vol. 7, no. 2, pp. 258–262, 1962.
- [38] I. Söderkvist, "Perturbation analysis of the orthogonal procrustes problem," *BIT*, vol. 33, no. 4, pp. 687–694, 1993.

- [39] N. T. Trendafilov and R. A. Lippert, "The multimode Procrustes problem," *Linear Algebra Appl.*, vol. 349, no. 1-3, pp. 245–264, 2002.
- [40] J. C. Gower and G. B. Dijksterhuis, *Procrustes Problems*. Oxford University Press, 2007.
- [41] N. J. Higham, "The symmetric Procrustes problem," *BIT*, vol. 28, no. 1, pp. 133–143, 1988.
- [42] J. M. ten Berge, H. A. Kiers, and J. J. Commandeur, "Orthogonal Procrustes rotation for matrices with missing values," *Br. J. Math. Stat. Psychol.*, vol. 46, no. 1, pp. 119–134, 1993.
- [43] K. S. Arun, "A Unitarily Constrained Total Least Squares Problem in Signal Processing," *SIAM J. Matrix Anal. Appl.*, vol. 13, no. 3, pp. 729–745, 1992.
- [44] N. Gillis and P. Sharma, "A semi-analytical approach for the positive semidefinite Procrustes problem," *Linear Algebra Appl.*, vol. 540, pp. 112–137, 2018.
- [45] P. H. Schönemann, "A generalized solution of the orthogonal procrustes problem," *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966.
- [46] T. Pumir, A. Singer, and N. Boumal, "The generalized orthogonal Procrustes problem in the high noise regime," *Inf. Inference A J. IMA*, vol. 10, no. 3, pp. 921–954, 2021.
- [47] T. Viklands, *Algorithms for the Weighted Orthogonal Procrustes Problem and other Least Squares Problems*. PhD thesis, Umeå University, 2006.
- [48] L. Eldén and H. Park, "A Procrustes problem on the Stiefel manifold," *Numer. Math.*, vol. 82, no. 4, pp. 599–619, 1999.
- [49] L. E. Andersson and T. Elfving, "A constrained Procrustes problem," *SIAM J. Matrix Anal. Appl.*, vol. 18, no. 1, pp. 124–139, 1997.
- [50] P. A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [51] N. Boumal, B. Mishra, P. A. Absil, and R. Sepulchre, "Manopt, a Matlab toolbox for optimization on manifolds," *J. Mach. Learn. Res.*, vol. 15, no. 42, pp. 1455–1459, 2014.
- [52] M. R. Jovanović, P. J. Schmid, and J. W. Nichols, "Sparsity-promoting dynamic mode decomposition," *Phys. Fluids*, vol. 26, no. 2, p. 024103, 2014.
- [53] J. L. Proctor, S. L. Brunton, and J. N. Kutz, "Dynamic mode decomposition with control," *SIAM J. Appl. Dyn. Syst.*, vol. 15, no. 1, pp. 142–161, 2016.
- [54] B. R. Noack, W. Stankiewicz, M. Morzynski, and P. J. Schmid, "Recursive dynamic mode decomposition of a transient cylinder wake," *J. Fluid Mech.*, vol. 809, pp. 843–872, 2016.
- [55] M. S. Hemati, M. O. Williams, and C. W. Rowley, "Dynamic mode decomposition for large and streaming datasets," *Physics of Fluids (1994-present)*, vol. 26, no. 11, p. 111701, 2014.
- [56] H. Zhang, C. W. Rowley, E. A. Deem, and L. N. Cattafesta, "Online dynamic mode decomposition for time-varying systems," *SIAM J. Appl. Dyn. Syst.*, vol. 18, no. 3, pp. 1586–1609, 2019.
- [57] S. L. Brunton, J. L. Proctor, J. H. Tu, and J. N. Kutz, "Compressed sensing and dynamic mode decomposition," *Journal of Computational Dynamics*, vol. 2, no. 2, pp. 165–191, 2015.
- [58] F. Gueniat, L. Mathelin, and L. Pastur, "A dynamic mode decomposition approach for large and arbitrarily sampled systems," *Physics of Fluids*, vol. 27, no. 2, p. 025113, 2015.
- [59] J. H. Tu, C. W. Rowley, J. N. Kutz, and J. K. Shang, "Spectral analysis of fluid flows using sub-Nyquist-rate PIV data," *Experiments in Fluids*, vol. 55, no. 9, pp. 1–13, 2014.
- [60] G. H. Golub and C. F. Van Loan, *Matrix Computations*, vol. 3. JHU Press, 2013.
- [61] C. Eckart and G. Young, "The approximation of one matrix by another of lower rank," *Psychometrika*, vol. 1, no. 3, pp. 211–218, 1936.
- [62] I. Mezić, "Analysis of fluid flows via spectral properties of the Koopman operator," *Ann. Rev. Fluid Mech.*, vol. 45, no. 1, pp. 357–378, 2013.
- [63] S. L. Brunton, M. Budišić, E. Kaiser, and J. N. Kutz, "Modern Koopman Theory for Dynamical Systems," *arXiv preprint arXiv:2102.12086*, 2021.
- [64] G. Dijksterhuis, "Procrustes analysis in studying sensory-instrumental relations," *Food Qual. Prefer.*, vol. 5, no. 1-2, pp. 115–120, 1994.

- [65] M. A. A. Cox and T. F. Cox, "Multidimensional Scaling," in *Handb. Data Vis.*, pp. 315–347, Springer, Berlin, Heidelberg, 2008.
- [66] X. Zhang, J. Zou, K. He, and J. Sun, "Accelerating Very Deep Convolutional Networks for Classification and Detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 10, pp. 1943–1955, 2016.
- [67] M. B. Richman, "Rotation of principal components," *J. Climatol.*, vol. 6, no. 3, pp. 293–335, 1986.
- [68] B. Herrmann, P. J. Baddoo, R. Semaan, S. L. Brunton, and B. J. McKeon, "Data-driven resolvent analysis," *J. Fluid Mech.*, vol. 918, p. A10, 2021.
- [69] E. Noether, "Invariante variationsprobleme," *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, vol. 1918, pp. 235–257, 1918.
- [70] A. Dutt and V. Rokhlin, "Fast Fourier Transforms for Nonequispaced Data," *SIAM J. Sci. Comput.*, vol. 14, no. 6, pp. 1368–1393, 1993.
- [71] J. F. Gibson, "Channelflow: A spectral Navier-Stokes simulator in C++," tech. rep., U. New Hampshire, 2014. Channelflow.org.
- [72] M. Ilak and C. W. Rowley, "Modeling of transitional channel flow using balanced proper orthogonal decomposition," *Phys. Fluids*, vol. 20, p. 034103, 2008.
- [73] N. J. Higham, *Functions of matrices: theory and computation*. Society for Industrial and Applied Mathematics, 2008.
- [74] D. J. Griffiths, *Introduction to quantum mechanics*. Pearson Prentice Hall, 1995.
- [75] D. T. Smithey, M. Beck, M. G. Raymer, and A. Faridani, "Measurement of the Wigner distribution and the density matrix of a light mode using optical homodyne tomography: Application to squeezed states and the vacuum," *Phys. Rev. Lett.*, vol. 70, no. 9, pp. 1244–1247, 1993.
- [76] K. Vogel and H. Risken, "Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase," *Phys. Rev. A*, vol. 40, no. 5, pp. 2847–2849, 1989.
- [77] J. S. Lundeen, B. Sutherland, A. Patel, C. Stewart, and C. Bamber, "Direct measurement of the quantum wavefunction," *Nature*, vol. 474, no. 7350, pp. 188–191, 2011.
- [78] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.
- [79] B. J. McKeon and A. S. Sharma, "A critical-layer framework for turbulent pipe flow," *J. Fluid Mech.*, vol. 658, p. 336382, 2010.
- [80] M. R. Jovanović, "From bypass transition to flow control and data-driven turbulence modeling: An input-output viewpoint," *Ann. Rev. Fluid Mech.*, vol. 53, pp. 311–345, 2021.
- [81] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, "Hydrodynamic stability without eigenvalues," *Science*, vol. 261, no. 5121, pp. 578–584, 1993.
- [82] T. Kato, *Perturbation Theory for Linear Operators*. Classics in Mathematics, Berlin, Heidelberg: Springer Berlin Heidelberg, 1966.
- [83] S. C. Reddy and L. N. Trefethen, "Pseudospectra of the convection-diffusion operator," *SIAM J. Appl. Math.*, vol. 54, no. 6, pp. 1634–1649, 1994.
- [84] L. N. Trefethen and M. Embree, *Spectra and pseudospectra: the behavior of nonnormal matrices and operators*. Princeton University Press, 2005.
- [85] T. A. Driscoll, N. Hale, and L. N. Trefethen, *Chebfun Guide*. Oxford: Pafnuty Publications, 2014.
- [86] J. L. Aurentz and L. N. Trefethen, "Block operators and spectral discretizations," *SIAM Rev.*, vol. 59, no. 2, pp. 423–446, 2017.
- [87] S. L. Brunton, B. W. Brunton, J. L. Proctor, E. Kaiser, and J. Nathan Kutz, "Chaos as an intermittently forced linear system," *Nat. Commun.*, vol. 8, no. 1, p. 19, 2017.
- [88] P. Krishnamurthy and F. Khorrami, "Generalized state scaling and applications to feedback, feedforward, and nontriangular nonlinear systems," *IEEE Trans. Automat. Contr.*, vol. 52, no. 1, pp. 102–108, 2007.
- [89] A. R. Teel, "A nonlinear small gain theorem for the analysis of control systems with saturation," *IEEE Trans. Automat. Contr.*, vol. 41, no. 9, pp. 1256–1270, 1996.

- [90] A. M. Annaswamy and J. Baillieul, "Adaptive Control of Nonlinear Systems with a Triangular Structure," *IEEE Trans. Automat. Contr.*, vol. 39, no. 7, pp. 1411–1428, 1994.
- [91] Z. Jackiewicz, M. Rahman, and B. D. Welfert, "Numerical solution of a Fredholm integro-differential equation modelling $\dot{\theta}$ -neural networks," *Appl. Math. Comput.*, vol. 195, no. 2, pp. 523–536, 2008.
- [92] J. Medlock and M. Kot, "Spreading disease: Integro-differential equations old and new," *Math. Biosci.*, vol. 184, no. 2, pp. 201–222, 2003.
- [93] N. J. Higham, "Matrix nearness problems and applications," *Appl. Matrix Theory*, pp. 1–27, 1989.
- [94] A. Zygmund, *Trigonometric series*, vol. I, II. Cambridge University Press, 1988.
- [95] D. Ruiz-Antolin and A. Townsend, "A nonuniform fast fourier transform based on low rank approximation," *SIAM J. Sci. Comput.*, vol. 40, no. 1, pp. A529–A547, 2018.
- [96] G. H. Golub and C. F. Van Loan, *Matrix computations*, vol. 3. JHU Press, 2012.
- [97] S. V. Huffel and J. Vandewalle, *The total least squares problem: computational aspects and analysis*. SIAM, 1991.
- [98] J. Yang and Y. B. Deng, "Procrustes problems for general, triangular, and symmetric toeplitz matrices," *J. Appl. Math.*, 2013.
- [99] M. R. Jovanović, P. J. Schmid, and J. W. Nichols, "Sparsity-promoting dynamic mode decomposition," *Physics of Fluids*, vol. 26, no. 2, p. 024103, 2014.
- [100] N. Levinson, "The Wiener (Root Mean Square) Error Criterion in Filter Design and Prediction," *J. Math. Phys.*, vol. 25, no. 1-4, pp. 261–278, 1946.
- [101] R. P. Brent, F. G. Gustavson, and D. Y. Yun, "Fast solution of toeplitz systems of equations and computation of Padé approximants," *J. Algorithms*, vol. 1, no. 3, pp. 259–295, 1980.
- [102] M. Stewart, "A superfast toeplitz solver with improved numerical stability," *SIAM J. Matrix Anal. Appl.*, vol. 25, no. 3, pp. 669–693, 2004.
- [103] N. Sedrakyan, "About the applications of one useful inequality," *Kvant J.*, vol. 97, no. 2, pp. 42–44, 1997.
- [104] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [105] C. D. Meyer, Jr., "Generalized Inversion of Modified Matrices," *SIAM J. Appl. Math.*, vol. 24, no. 3, pp. 315–323, 1973.