

OPEN ACCESS

Citation: Frame P, Towne A (2023) Space-time POD and the Hankel matrix. PLoS ONE 18(8): e0289637. https://doi.org/10.1371/journal. pone.0289637

Editor: Mohamed Kamel Riahi, Khalifa University of Science and Technology, UNITED ARAB EMIRATES

Received: December 20, 2022

Accepted: July 21, 2023

Published: August 10, 2023

Copyright: © 2023 Frame, Towne. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its Supporting information files.

Funding: The author received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

RESEARCH ARTICLE

Space-time POD and the Hankel matrix

Peter Frame *, Aaron Towne

Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI, United States of America

* pframe@umich.edu

Abstract

Time-delay embedding is an increasingly popular starting point for data-driven reduced-order modeling efforts. In particular, the singular value decomposition (SVD) of a block Hankel matrix formed from successive delay embeddings of the state of a dynamical system lies at the heart of several popular reduced-order modeling methods. In this paper, we show that the left singular vectors of this Hankel matrix are a discrete approximation of space-time proper orthogonal decomposition (POD) modes, and the singular values are square roots of the POD energies. Analogous to the connection between the SVD of a data matrix of snapshots and space-only POD, this connection establishes a clear interpretation of the Hankel modes grounded in classical theory, and we gain insights into the Hankel modes by instead analyzing the equivalent discrete space-time POD modes in terms of the correlation matrix formed by multiplying the Hankel matrix by its conjugate transpose. These insights include the distinct meaning of rows and columns, the implied norm in which the modes are optimal, the impact of the time step between snapshots on the modes, and an interpretation of the embedding dimension/height of the Hankel matrix in terms of the time window on which the modes are optimal. Moreover, the connections we establish offer opportunities to improve the convergence and computation time in certain practical cases, and to improve the accuracy of the modes with the same data. Finally, popular variants of POD, namely the standard space-only POD and spectral POD, are recovered in the limits that snapshots used to form each column of the Hankel matrix represent flow evolution over short and long times, respectively.

1 Introduction

Time series data, generated from simulations or experiments, are abundant in science and engineering, but analyzing or interpreting these data can be challenging. Often, such as in climate science or the analysis of financial data, researchers want an understanding of the governing laws underlying the time series. Other times, as in the case of fluid mechanics, a precise physical model exists, but it may be difficult to interpret the results, or simulating the full model may be computationally costly. In such cases, the goal is often to use data to uncover key physical mechanisms that contribute to the underlying dynamics [1] or to derive a less physical, but more computationally efficient, reduced-order model capable of approximating the dynamics at low cost.

Many techniques have emerged to address these challenges. At the heart of several of them, especially ones popular in the dynamical systems community, is analysis of a Hankel matrix. A

Hankel matrix has constant skew diagonals, i.e., the (i, j) entry of the matrix only depends on i + j. Common applications include system identification and minimal system realization [2], and the use of the Hankel matrix in this context goes back to the 1960s [3]. The Hankel matrix plays a key role in singular spectral analysis (SSA) [4] and in the eigensystem realization algorithm (ERA) [5]. Recently, the Hankel matrix has been used in the context of fluid dynamics. It is central to balanced truncation [6], which was made scalable to fluid dynamics problems with balanced POD [7, 8]. The Hankel matrix is also used in a variant of dynamic mode decomposition (DMD) [9] called Hankel DMD [10], where the goal is to extract the spectrum and modes of the Koopman operator for some dynamical system by performing DMD on a Hankel matrix of observables.

The entries of the Hankel matrix are taken from time series data, and moving down or right along the columns or rows of the matrix corresponds to moving forward in time. Therefore, the columns of the Hankel matrix are delay embeddings of the dynamical system that produced the time series. Delay embedding, which goes back to work from the 1980s [11, 12], is a method of encoding the state of a dynamical system by recording the time history of one (or a few) of its observables. Intuitively, the state of an *n*-dimensional system should be determined by *n* independent observables, which don't need to be the original degrees of freedom of the system [11]. This idea was later made rigorous by Takens [12], and, under weak conditions, 2n+ 1 entries of a time series are needed to determine the state of the system.

A recurring theme in applications of the Hankel matrix in dynamical systems is the singular value decomposition (SVD). Both SSA and ERA obtain their bases from the SVD of the Hankel matrix. In balanced truncation and balanced POD, the Hankel singular values are used. In the Hankel alternative view of Koopman (HAVOK) framework [13], the SVD of a Hankel matrix of data is used in order to form a low-rank linear model of the dynamics on some chaotic attractor. The modes in their linear model are the left singular vectors of the Hankel matrix formed from a time series of the dynamical system. Convolutional coordinates, used to represent the state of a dynamical system at a particular time in terms of its representation in some predefined temporal basis, can be defined in terms of the continuous SVD (Schmidt decomposition) of a continuous Hankel matrix [14]. The left singular vectors of the Hankel matrix have recently been called principal component trajectories and used for control [15].

A second, older, technique for analyzing time series data is proper orthogonal decomposition (POD). Originally introduced to the fluid dynamics community by Lumley [16] in 1967, it is known by a variety of names in other areas including principal component analysis, Karhunen-Loève decomposition, and empirical orthogonal functions. In POD, the flow data is analyzed statistically, and the objective is to search for the modes that most efficiently represent the data. Specifically, POD modes are defined to minimize the reconstruction error, as measured by the average square inner product, compared to any other basis of the same dimension. As introduced by Lumley, the most general version of POD seeks to describe the time evolution of the flow for a prespecified window of time, so the basis functions are functions both of space and of time. This general version is called space-time POD, and the reconstruction of a flow over the time window consists of these basis functions multiplied by constant coefficients. Space-time POD has been used infrequently in the literature. Notable exceptions include application of space-time POD to optimally describe transients [17], generalize dimension reduction methods [18], and study acoustic intermittency in the form of bursts in jets via conditioning [19, 20].

Following the work of Sirovich [21] and Aubry [22], today the most popular form of POD involves modes that are functions of space only. To represent a time-dependent flow, these modes are multiplied by time-varying coefficients. We refer to this form of POD as space-only POD [23], but note that it is often referred to simply as POD in the literature. Space-only POD

has been used extensively to form Galerkin-based reduced-order models [24–26], educe physically meaningful structures from flow data [27, 28], and reduce the data needed to store flow data [29]. We will show that space-only POD can be understood within the more general space-time POD framework as the limit as the time interval on which the optimization problem is defined goes to zero.

An increasingly popular variant of POD is spectral POD (SPOD). While also introduced in the original work of Lumley [16, 30], this form of POD was rarely employed until recently [23, 31-33]. Here, the objective is to optimally represent statistically stationary flows in the frequency domain. At each frequency, SPOD provides a set of spatial modes that capture the portion of the flow at that frequency more accurately, on average, than any other basis of the same order. SPOD can be formulated as the limit of space-time POD as the time interval on which the modes are defined goes to infinity.

Our aim in this paper is to show that the singular modes of the Hankel matrix are a discrete approximation of space-time POD modes and to demonstrate that understanding them as such is useful for analyzing them and improving their properties. This connection is established by observing that the Hankel matrix multiplied by its Hermetian transpose provides an approximation of the space-time correlation matrix whose eigendecomposition defines discrete space-time POD modes. Analyzing this space-time correlation matrix leads to insights into the modes and, in some cases, guidance about how to improve them. The Hankel modes are not the only approximation one can form of the space-time POD modes using the available data, and in many cases, they are not the most practical approximation.

With a surplus of available data, a sufficiently accurate space-time correlation matrix can be formed by throwing out many of the columns before taking the SVD, which can drastically reduce the cost and memory requirements of obtaining the modes. In fact, unless data is quite scarce, we show that it is inadvisable to form the entire Hankel matrix to obtain the modes for large systems. Discarding columns works because the Hankel matrix is a data matrix of trajectories, but each trajectory is very similar to the adjacent ones in the Hankel matrix, so a given trajectory provides little additional information. In the opposite limit, i.e., where data is severely limited, we derive a more accurate approximation of the space-time POD modes than that provided by the Hankel matrix SVD. The approximation works by exploiting the ergodicity of the system (ergodicity is assumed in forming the Hankel matrix) to form a more accurate space-time correlation matrix. We show that these modes are substantially more accurate for low dimensional systems than Hankel modes when data is severely limited as compared to converged Hankel modes.

The connection to space-time POD also makes clear the assumed inner product, which defines the sense in which the Hankel modes are optimal and clarifies the impact of the time step between successive snapshots on the approximation. Additionally, we show that the height and width of the Hankel matrix determine the extent to which temporal correlation is accounted for in defining the modes, and the convergence of the correlation, respectively. Finally, we show that in the limits of short and long delays, space-only and spectral POD are recovered, respectively.

The remainder of the paper is organized as follows. In section 2, we define the Hankel matrix and reiterate some of its applications. In section 3, we motivate and derive the continuous and discrete forms of space-only POD. We do this using a formalism that makes spectral and space-time POD follow easily, and that highlights the central role of the correlation matrix. In section 4, we derive the continuous and discrete forms of spectral POD. In section 5, we derive the continuous form of space-time POD and show that, in the limits of short and long time intervals, it reduces to space-only and spectral POD, respectively. In section 6, we show that the singular modes of the Hankel matrix provide a discrete approximation of space-time

POD modes by noting the Hankel matrix can be used to approximate the space-time correlation tensor. This connection motivates several improvements to the Hankel SVD procedure, and we also prove results about the convergence to space-only and spectral POD analogous to those for the continuous case. In section 7, we show that with the same time series data, we can construct a more accurate correlation matrix than the Hankel matrix multiplied by its conjugate transpose by exploiting the ergodicity of the system. The eigendecomosition of this correlation matrix gives more converged modes than the SVD of the Hankel matrix. In section 8, we demonstrate these results using a lid-driven cavity flow at Re = 22, 000. Finally, in section 9, we summarize the paper and present our conclusions.

2 The Hankel matrix

A Hankel matrix is a matrix whose skew-diagonals are constant, i.e., the (i, j) entry only depends on i + j,

$$H = \begin{bmatrix} g_1 & g_2 & \cdots & g_m \\ g_2 & g_3 & \cdots & g_{m+1} \\ \vdots & \vdots & \ddots & \vdots \\ g_d & g_{d+1} & \cdots & g_{m+d-1} \end{bmatrix}.$$
 (1)

In dynamical systems theory, these matrices are formed from a time series of observables g_1 , $g_2, g_3, \ldots, g_{m+d-1}$. This means that the columns of H are time-delay embeddings of the dynamical system. In this paper, we focus our attention on the more general block Hankel matrix

$$\mathbf{H} = \begin{bmatrix} \mathbf{q}_{1} & \mathbf{q}_{2} & \cdots & \mathbf{q}_{m} \\ \mathbf{q}_{2} & \mathbf{q}_{3} & \cdots & \mathbf{q}_{m+1} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{q}_{d} & \mathbf{q}_{d+1} & \cdots & \mathbf{q}_{m+d-1} \end{bmatrix},$$
(2)

formed from a time series of vector-valued observables $\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_d, \ldots, \mathbf{q}_{m+d-1}$.

Applications of the Hankel matrix include the singular spectrum analysis (SSA) [4], the eigensystem realization algorithm (ERA) [5] and the Hankel alternative view of Koopman (HAVOK) [13]. All of these methods depend on the SVD of the Hankel matrix,

$$\mathbf{H} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^*,\tag{3}$$

and in particular use the left singular vectors of the Hankel matrix as a basis, which have been called principal component trajectories [15]. The connections to time-delay embedding, Koopman theory, and dynamic mode decomposition have garnered increasing interest in analysis of the Hankel matrix.

3 Space-only POD

All forms of POD are statistical methods and view the dynamical system to which they are applied as random. This is a practical choice; though the system is not random, it might be chaotic, and without knowledge of the exact initial condition of the system, viewing it as random is the best we can do. If we think of a spatial realization of the flow, or a *snapshot* $q(\mathbf{x})$, as a vector in a vector space, space-only POD seeks to find the direction in the vector space along which there is the most variation between different snapshots, i.e., the most energy as defined

by some spatial inner product. Knowing the coordinate of a snapshot of the flow along this direction provides a better approximation, on average, than any other coordinate. To characterize a flow exactly, we need to specify the flow field everywhere in the domain, but if giving the flow's coordinates along just a few important directions approximates it to high precision, then it is useful to look for these important directions.

3.1 Formulation

The direction with the most variation is formalized as the mode which maximizes the expected value of the square of the projection of the flow snapshot, $q(\mathbf{x})$ onto the mode,

$$\lambda[\boldsymbol{\phi}(\mathbf{x})] = \frac{1}{\|\boldsymbol{\phi}(\mathbf{x})\|^2} \mathbb{E}\left[\|\langle \boldsymbol{q}(\mathbf{x}), \boldsymbol{\phi}(\mathbf{x})\rangle\|^2\right],\tag{4a}$$

$$\boldsymbol{\phi}_1(\mathbf{x}) = \arg \, \max \lambda[\boldsymbol{\phi}(\mathbf{x})]. \tag{4b}$$

Here, the expectation operator $\mathbb{E}[\cdot]$ acts over time, and λ is a functional which takes any function as input (in the appropriate function space) and returns the variation in the flow along it. The first POD mode ϕ_1 maximizes this functional. Typically, the modes are normalized to unity, but for clarity we will keep the magnitude in the formulae. The magnitude operator $\|\cdot\|$ is defined in the usual way by the inner product, and the inner product takes the form

$$\langle \boldsymbol{q}_1(\mathbf{x}), \boldsymbol{q}_2(\mathbf{x}) \rangle = \int_{\Omega} \boldsymbol{q}_2^*(\mathbf{x}) \boldsymbol{W}(\mathbf{x}) \boldsymbol{q}_1(\mathbf{x}) d\mathbf{x},$$
 (5)

where *W* is a weight matrix and Ω is the spatial domain of interest. The weight matrix is often used to make the norm correspond to some physical definition of energy, e.g., turbulent kinetic energy, or to give preference to certain flow variables or regions of the flow. The first POD mode ϕ_1 maximizes this functional. In solving (4b), it is helpful to rewrite (4a) as an inner product with the correlation tensor [16, 34],

$$\lambda[\boldsymbol{\phi}(\mathbf{x})] = \frac{1}{\|\boldsymbol{\phi}(\mathbf{x})\|^2} \langle \langle \boldsymbol{\phi}(\mathbf{x}'), \boldsymbol{C}(\mathbf{x}, \mathbf{x}') \rangle, \boldsymbol{\phi}(\mathbf{x}) \rangle,$$
(6)

where the correlation is defined as

$$\boldsymbol{C}(\mathbf{x}_1, \mathbf{x}_2) = \mathbb{E}[\boldsymbol{q}(\mathbf{x}_1)\boldsymbol{q}(\mathbf{x}_2)^*].$$
(7)

The correlation tensor is symmetric (under interchange of its two arguments), so it has orthogonal eigenfunctions { $\nu_1(\mathbf{x}), \nu_2(\mathbf{x}), \ldots$ }, which can be ordered so that the associated eigenvalues { $\lambda_1 \ge \lambda_2 \ge , \cdots \ge 0$ } are non-increasing. These are eigenfunctions in the sense that they satisfy

$$\langle \boldsymbol{\nu}_k(\mathbf{x}_2), \boldsymbol{C}(\mathbf{x}_1, \mathbf{x}_2) \rangle = \lambda_k \boldsymbol{\nu}_k(\mathbf{x}_1).$$
 (8)

Writing ϕ in the basis of these eigenfunctions provides insight into the maximization problem (4), and the energy of ϕ (6) can be written in terms of the expansion coefficients used to express it [16],

$$\lambda[\boldsymbol{\phi}(\mathbf{x})] = \sum_{k=1}^{\infty} |c_k|^2 \lambda_k \quad \text{for normalized } \boldsymbol{\phi}(\mathbf{x}), \tag{9}$$

where the expansion coefficient $c_k = \langle \phi(\mathbf{x}), \mathbf{v}_k(\mathbf{x}) \rangle$ is the projection of $\phi(\mathbf{x})$ on the k^{th} eigenfunction. The solution that maximizes (9) is $c_1 = 1$, $c_{\neq 1} = 0$, because $\lambda_1 > \lambda_{\neq 1}$. This tells us

that the first POD mode is the eigenvector of the correlation tensor with the greatest eigenvalue. To define the latter modes, we maximize $\lambda[\phi(\mathbf{x})]$ over all $\phi(\mathbf{x})$ orthogonal to previous modes. A simple inductive argument shows that if the first *k* POD modes are the first *k* eigenfunctions of the correlation tensor, then the $k + 1^{st}$ mode must be the $k + 1^{st}$ eigenfunction because this orthogonality condition tells us that $c_{\langle k+1} = 0$ for the $k + 1^{st}$ mode. The POD modes are therefore the eigenfunctions of the correlation tensor, $\phi_k = v_k$, and their energies are the eigenvalues, $\lambda[\phi_k(\mathbf{x})] = \lambda_k$. In terms of the inner product, the POD modes satisfy

$$\int_{\Omega} \boldsymbol{C}(\mathbf{x}_1, \mathbf{x}_2) \boldsymbol{W}(\mathbf{x}_2) \boldsymbol{\phi}_k(\mathbf{x}_2) d\mathbf{x}_2 = \lambda_k \boldsymbol{\phi}_k(\mathbf{x}_1).$$
(10)

3.2 With discrete data

In practice, space-only POD modes are approximated using data from a simulation or experiment defined on a discrete set of points and sampled in time. To approximate the true continuous space-only POD modes, the data is used to approximate a correlation matrix whose eigendecomposition gives the modes. Denoting one discrete snapshot as $\mathbf{q}_j \in \mathbb{R}^N$, we need a correlation matrix of the form $\mathbf{C}^{\mathbf{x}} \in \mathbb{R}^{N \times N}$, where each element represents the correlation between two components of the flow at different points. We label this spatial correlation matrix with an \mathbf{x} superscript to distinguish it from other correlations that arise later in the paper. It can be approximated as

$$\mathbf{C}^{\mathbf{x}} = \frac{1}{m} \mathbf{Q} \mathbf{Q}^{*}$$
 with $\mathbf{Q} = [\mathbf{q}_{1}, \mathbf{q}_{2}, \cdots, \mathbf{q}_{m}] \in \mathbb{R}^{N \times m}$. (11)

The data matrix **Q** contains an ensemble of snapshots, and $\frac{1}{m}$ **QQ**^{*} is an approximation of C because it gives each component of the correlation as a sample average over the realizations,

$$\left(\mathbf{Q}\mathbf{Q}^*\right)_{ij} = \sum_{k=1}^m (\mathbf{q}_k)_i (\mathbf{q}_k^*)_j.$$
(12)

Discrete space-only POD modes are given by the eigenvectors of the discrete correlation matrix multiplied by the weight,

$$\mathbf{C}^{\mathbf{x}}\mathbf{W}\mathbf{\Phi} = \mathbf{\Phi}\mathbf{\Lambda},\tag{13}$$

where **W** is the discrete weight matrix. This weight is used both as the discrete version of the continuous weight and to account for numerical quadrature of the integral in (10). The columns of the matrix $\Phi = [\phi_1, \phi_2, ...]$ are the discrete space-only POD modes, and the diagonal matrix Λ contains the corresponding eigenvalues.

The discrete space-only POD modes are also related to the left singular vectors of the data matrix \mathbf{Q} and the weight matrix, and the details of this relation will prove important to understanding the connection between the SVD of the Hankel matrix and POD. If we take the singular value decomposition to obtain

$$\frac{1}{\sqrt{m}}\mathbf{W}^{\frac{1}{2}}\mathbf{Q} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^{*},\tag{14}$$

then we can write the correlation tensor multiplied by the weight as

 $C^{x}W = W^{-\frac{1}{2}}U\Sigma V^{*}V\Sigma U^{*}W^{\frac{1}{2}}$. Because both U and V are orthonormal, multiplying by $W^{-\frac{1}{2}}U$ we

have

$$\mathbf{C}^{\mathbf{x}}\mathbf{W}\mathbf{W}^{-\frac{1}{2}}\mathbf{U} = \mathbf{W}^{-\frac{1}{2}}\mathbf{U}\boldsymbol{\Sigma}^{2}.$$
 (15)

This constitutes an eigendecomposition of $C^{x}W$, and therefore the POD modes are related to the left singular vectors of the data matrix and the energies to the singular values,

$$\Phi = \mathbf{W}^{-\frac{1}{2}}\mathbf{U} \quad \text{and} \quad \Lambda = \mathbf{\Sigma}^2. \tag{16}$$

It is important to remember that the correlation matrix is approximate because it is obtained from finite data. If the data are highly correlated, e.g., if the snapshots are taken from a time series whose length is on the same order as the characteristic timescale of the flow or shorter, then the correlation tensor, and hence the POD modes, will be inaccurate. Both will increase in accuracy with the number of realizations and the independence of the realizations.

4 Spectral POD

Spectral POD produces an optimal frequency domain representation of statistically stationary flows. At each frequency, SPOD modes reconstruct the Fourier transform of a flow more accurately, on average, than any other reconstruction of the same order.

4.1 Formulation

Spectral POD can also be cast as an optimization problem, analogous to that of space-only POD, as

$$\lambda_{\omega}[\boldsymbol{\phi}(\mathbf{x})] = \frac{1}{\|\boldsymbol{\phi}(\mathbf{x})\|^2} \mathbb{E}\left[\|\langle \hat{\boldsymbol{q}}_{\omega}(\mathbf{x}), \boldsymbol{\phi}(\mathbf{x}) \rangle\|^2\right], \tag{17a}$$

$$\boldsymbol{\psi}_{\omega,1}(\mathbf{x}) = \arg \max \lambda_{\omega}[\boldsymbol{\phi}(\mathbf{x})]. \tag{17b}$$

Here, the expectation operator acts over the Fourier transform of segments of the flow. These segments can be subsections of a longer time series or correspond to separate realizations of the flow. Similar to the space-only case, λ_{ω} is a functional that returns the energy in the flow at frequency ω captured by the argument. The Fourier transformed flow is defined as

$$\hat{\boldsymbol{q}}_{\omega}(\mathbf{x}) = \int_{-\infty}^{\infty} \boldsymbol{q}(\mathbf{x}, t) e^{-i\omega t} dt, \qquad (18)$$

and the inner product is the same as before, given in (5). Analogous to the space-only case, the optimization problem (17a) can be rewritten in terms of correlations,

$$\lambda_{\omega}[\boldsymbol{\phi}(\mathbf{x})] = \frac{1}{\|\boldsymbol{\phi}(\mathbf{x})\|^2} \langle \langle \boldsymbol{\phi}(\mathbf{x}'), \boldsymbol{S}_{\omega}(\mathbf{x}, \mathbf{x}') \rangle, \boldsymbol{\phi}(\mathbf{x}) \rangle,$$
(19)

where S_{ω} , called the cross-spectral density tensor, is the Fourier transform pair of the statistically stationary space-time correlation tensor,

$$\mathbf{S}_{\omega}(\mathbf{x}_1, \mathbf{x}_2) = \int_{-\infty}^{\infty} \mathbf{C}(\mathbf{x}_1, \mathbf{x}_2, \tau) e^{-i\omega\tau} d\tau$$
(20)

with

$$\boldsymbol{C}(\boldsymbol{x}_1, \boldsymbol{x}_2, \tau) = \mathbb{E}[\boldsymbol{q}(\boldsymbol{x}_1, t+\tau)\boldsymbol{q}^*(\boldsymbol{x}_2, t)].$$
(21)

Note that in statistically stationary flow, the correlation tensor C only depends on one time variable, τ , which represents the difference in the two times in a more general C that arises in space-time POD.

The mathematical structure of (19) is the same as in space-only POD, so to maximize (19) we solve the integral eigenvalue problem

$$\int_{\Omega} S_{\omega}(\mathbf{x}_1, \mathbf{x}_2) W(\mathbf{x}_2) \boldsymbol{\psi}_{\omega, k}(\mathbf{x}_2) d\mathbf{x}_2 = \lambda_k \boldsymbol{\psi}_{\omega, k}(\mathbf{x}_1),$$
(22)

where λ_k is the k^{th} largest eigenvalue of S_{ω} , and $\psi_{\omega,k}$ is the k^{th} SPOD mode at frequency ω .

4.2 With discrete data

As we saw in 3.2, data can be used to compute approximate space-only POD modes by forming a discrete spatial correlation matrix **C**. Analogously, approximate spectral POD modes can be computed by using data to form a discrete cross-spectral density matrix. To do this, instead of taking the Fourier transform of the space-time correlation matrix, the Wiener-Khinchin theorem is invoked, which states that the cross-spectral density tensor is equivalent to the correlation between points in the flow in Fourier space,

$$\mathbf{S}_{\omega}(\mathbf{x}_1, \mathbf{x}_2) = \mathbb{E}[\hat{\boldsymbol{q}}_{\omega}(\mathbf{x}_1)\hat{\boldsymbol{q}}_{\omega}^*(\mathbf{x}_2)].$$
(23)

To find \mathbf{S}_{ω} in this way, many realizations of the flow are needed in the frequency domain, at a particular frequency: $\hat{\mathbf{Q}}_{\omega} = [\hat{\mathbf{q}}_{\omega,1}, \hat{\mathbf{q}}_{\omega,2}, \dots \hat{\mathbf{q}}_{\omega,m}]$, where $\hat{\mathbf{q}}_{\omega,i} \in \mathbb{C}^N$ is the spatially discretized representation of the Fourier transform of the *i*th flow realization, at frequency ω . These frequency domain realizations come from applying the discrete Fourier transform to time series data, and to get more than one realization at each frequency from a single time series, the time series is broken up into *m* (possibly overlapping) blocks, and the discrete Fourier transform is taken of each block [39]. With this data matrix of Fourier realizations, the cross-spectral density matrix can be approximated as

$$\mathbf{S}_{\omega} = \frac{1}{m} \hat{\mathbf{Q}}_{\omega} \hat{\mathbf{Q}}_{\omega}^{*}.$$
 (24)

The product of the data matrix with its transpose approximates the discrete cross-spectral density tensor because it averages products of different components of $\hat{\mathbf{q}}_{\omega}$ over the realizations. The SPOD modes at frequency ω are the eigenvectors of this approximated cross-spectral density tensor multiplied by the weight matrix,

$$\mathbf{S}_{\boldsymbol{\omega}} \mathbf{W} \mathbf{\Psi}_{\boldsymbol{\omega}} = \mathbf{\Psi}_{\boldsymbol{\omega}} \mathbf{\Lambda}_{\boldsymbol{\omega}},\tag{25}$$

where $\Psi_{\omega} = [\psi_{\omega,1}, \psi_{\omega,2} \dots]$, and Λ_{ω} is the diagonal matrix of eigenvalues at frequency ω . As before, the modes can be obtained using the SVD of the (weighted) data matrix of realizations in Fourier space,

$$\frac{1}{\sqrt{m}} \mathbf{W}^{\frac{1}{2}} \hat{\mathbf{Q}}_{\omega} = \mathbf{U}_{\omega} \boldsymbol{\Sigma}_{\omega} \mathbf{V}_{\omega}^{*}.$$
(26)

The modes and energies are then given by

$$\Psi_{\omega} = \mathbf{W}^{-\frac{1}{2}} \mathbf{U}_{\omega} \quad \text{and} \quad \Lambda_{\omega} = \boldsymbol{\Sigma}_{\omega}^{2}. \tag{27}$$

5 Space-time POD

In this section, we describe a generalization of space-only POD and spectral POD called spacetime POD, in which the goal is to find modes that optimally represent the flow evolution on a finite-time window [0, T].

5.1 Formulation

The goal is formalized in the same way as the previous two cases, by maximizing the expected value of the square of an inner product,

$$\lambda_{\mathrm{xt}}[\boldsymbol{\phi}(\mathbf{x},t)] = \frac{1}{\|\boldsymbol{\phi}(\mathbf{x},t)\|^2} \mathbb{E}\left[\|\langle \boldsymbol{q}(\mathbf{x},t), \boldsymbol{\phi}(\mathbf{x},t) \rangle_{\mathrm{xt}}\|^2\right],$$
(28a)

$$\boldsymbol{\phi}_1(\mathbf{x}, t) = \arg \, \max \lambda_{\mathbf{x}\mathbf{t}}[\boldsymbol{\phi}(\mathbf{x}, t)]. \tag{28b}$$

This time, however, our inner product acts over both space and time,

$$\langle \boldsymbol{q}_1(\mathbf{x},t), \boldsymbol{q}_2(\mathbf{x},t) \rangle_{\mathrm{xt}} = \int_{0}^{T} \int_{\Omega} \boldsymbol{q}_2^*(\mathbf{x},t) \boldsymbol{W}(\mathbf{x},t) \boldsymbol{q}_1(\mathbf{x},t) d\mathbf{x} dt.$$
 (29)

The expectation operator acts over finite time segments of the flow. These segments can be separate realizations of the flow or can be extracted from a single long time series if the flow is ergodic. The space-time optimization problem (28) has a different physical meaning than the space-only optimization and the spectral optimization problems (4) and (17), respectively. The desired mode here will be a function of both space and time, and its accuracy is measured by an inner product over both space and time as well, whereas, in the previous two cases, the modes were functions of only space. Mathematically, however, the problems are quite similar; instead of the spatial domain, Ω , we now have the spatiotemporal domain $\mathcal{D} \equiv \Omega \times [0, T]$, and the modes are now a functions of both space and time, $\mathbf{z} \equiv [\mathbf{x}, t]^T$. With these definitions, the space-time optimization problem can be written

$$\lambda_{\mathbf{x}\mathbf{t}}[\boldsymbol{\phi}(\mathbf{z})] = \frac{1}{\|\boldsymbol{\phi}(\mathbf{z})\|^2} \mathbb{E}\left[\|\langle \boldsymbol{q}(\mathbf{z}), \boldsymbol{\phi}(\mathbf{z}) \rangle_{\mathbf{x}\mathbf{t}}\|^2\right], \tag{30a}$$

$$\boldsymbol{\phi}_1(\mathbf{z}) = \arg \, \max \lambda_{\mathrm{xt}}[\boldsymbol{\phi}(\mathbf{z})], \tag{30b}$$

where the inner product is

$$\langle \boldsymbol{q}_1(\boldsymbol{z}), \boldsymbol{q}_2(\boldsymbol{z}) \rangle_{\mathbf{xt}} = \int_{\mathcal{D}} \boldsymbol{q}_2^*(\boldsymbol{z}) \boldsymbol{W}(\boldsymbol{z}) \boldsymbol{q}_1(\boldsymbol{z}) d\boldsymbol{z}.$$
 (31)

The maximization is exactly the same mathematically as (4), so we know the solution will be that the optimal space-time modes are the eigenfunctions of the correlation between two points in the spatiotemporal domain,

$$\int_{\mathcal{D}} \boldsymbol{C}(\boldsymbol{z}_1, \boldsymbol{z}_2) \boldsymbol{W}(\boldsymbol{z}_2) \boldsymbol{\phi}_k(\boldsymbol{z}_2) d\boldsymbol{z}_2 = \lambda_k \boldsymbol{\phi}_k(\boldsymbol{z}_1).$$
(32)

Rewriting this again in terms of space and time, the space-time POD modes satisfy

$$\int_{0}^{1} \int_{\Omega} C(\mathbf{x}_1, t_1, \mathbf{x}_2, t_2) W(\mathbf{x}_2, t_2) \boldsymbol{\phi}_k(\mathbf{x}_2, t_2) d\mathbf{x}_2 dt_2 = \lambda_k \boldsymbol{\phi}_k(\mathbf{x}_1, t_1),$$
(33)

where the space-time correlation tensor is defined as

$$\boldsymbol{C}(\mathbf{x}_1, t_1, \mathbf{x}_2, t_2) = \mathbb{E}[\boldsymbol{q}(\mathbf{x}_1, t_1)\boldsymbol{q}(\mathbf{x}_2, t_2)^*].$$
(34)

In section 6, we will show that the Hankel SVD modes constitute a discrete approximation of space-time POD modes. Before doing so, we first show that spectral and space-only POD can be recovered from space-time POD in the limits of long and short time windows, respectively. The latter result is novel to the best of our knowledge.

5.2 Space-time POD on long times becomes spectral POD

Here, we show that in the limit that $T \to \infty$, the space-time modes become spectral POD modes. This limit has been shown in the literature, e.g., in [23, 30], but here we give a different proof. We restrict our attention to ergodic flows, though all that is needed for the results to be valid is that the flow is wide-sense statistically stationary, i.e., the first and second moments of the flow don't change with time. The integral eigenvalue problem that defines the modes when $T \to \infty$ is

$$\int_{-\infty}^{\infty} \int_{\Omega} C(\mathbf{x}_1, t_1, \mathbf{x}_2, t_2) W(\mathbf{x}_2) \boldsymbol{\phi}(\mathbf{x}_2, t_2) d\mathbf{x}_2 dt_2 = \lambda \boldsymbol{\phi}(\mathbf{x}_1, t_1).$$
(35)

For convenience, the bounds on time have been shifted to cover the entire real line. Because the flow is stationary, the correlation must only depend on the difference of the two times, not on the times themselves, so

$$\boldsymbol{C}(\mathbf{x}_1, t_1, \mathbf{x}_2, t_2) \to \boldsymbol{C}(\mathbf{x}_1, \mathbf{x}_2, t_1 - t_2).$$
(36)

Switching the order of integration, we have

$$\int_{\Omega} \int_{-\infty}^{\infty} \boldsymbol{C}(\mathbf{x}_1, \mathbf{x}_2, t_1 - t_2) \boldsymbol{W}(\mathbf{x}_2) \boldsymbol{\phi}(\mathbf{x}_2, t_2) dt_2 d\mathbf{x}_2 = \lambda \boldsymbol{\phi}(\mathbf{x}_1, t_1), \qquad (37)$$

which is a convolution of *C* with ϕ_k . Taking the Fourier transform of both sides and using the convolution theorem, we have

$$\int_{\Omega} \boldsymbol{S}_{\omega}(\mathbf{x}_{1},\mathbf{x}_{2}) \boldsymbol{W}(\mathbf{x}_{2}) \hat{\boldsymbol{\phi}}_{\omega}(\mathbf{x}_{2}) d\mathbf{x}_{2} = \lambda \hat{\boldsymbol{\phi}}_{\omega}(\mathbf{x}_{1}).$$
(38)

The modes must satisfy (38) for every ω . SPOD modes, $\psi_{\omega'}$ at some ω' , are solutions as they are defined by this equation for $\omega = \omega'$, and are zero for $\omega \neq \omega'$. The SPOD mode with the k^{th} greatest energy λ (over the modes at all frequencies) is the k^{th} space-time POD mode.

5.3 Space-time POD on short times becomes space-only POD

Here, we show an analogous result for space-only POD. Specifically, we investigate what happens to the space-time POD modes in the limit that the time *T* on which they evolve is short compared to the time scales of the flow. In this limit, the space-time correlation between two

points (\mathbf{x}_1, t_1) and (\mathbf{x}_2, t_2) is only a function of their locations in space because t_1 and t_2 are negligibly different. Intuitively, we can now think of this integral over space and time as an integral over space multiplied by time because there is no time dependence in the integrand, so long as the weight doesn't depend on time (which would be an unusual choice). Rewriting condition (33) with a constant weight and correlation, we have

$$\int_{0}^{1} \int_{\Omega} C(\mathbf{x}_{1}, \mathbf{x}_{2}) W(\mathbf{x}_{2}) \boldsymbol{\phi}_{k}(\mathbf{x}_{2}, t_{2}) d\mathbf{x}_{2} dt_{2} = \lambda_{k} \boldsymbol{\phi}_{k}(\mathbf{x}_{1}, t_{1}).$$
(39)

Because the left-hand side is not a function of t_1 , the modes must indeed be constant in time, and the condition on the modes becomes an integral over space multiplied by time,

$$T\int_{\Omega} \boldsymbol{C}(\mathbf{x}_1, \mathbf{x}_2) \boldsymbol{W}(\mathbf{x}_2) \boldsymbol{\phi}_k(\mathbf{x}_2) d\mathbf{x}_2 = \lambda_k \boldsymbol{\phi}_k(\mathbf{x}_1).$$
(40)

This condition is exactly that of space-only POD, as shown in (10). As a result, the space-time POD modes converge to the space-only POD modes and their energies are proportional to the space-only POD energies (with proportionality constant T).

5.4 Space-time POD with discrete data

Given time series data on a spatial grid from an experiment or simulation, how do we compute the space-time POD modes? Above, we have seen that the solution to the space-time POD problem is given by eigendecomposition of the space-time correlation, which becomes a matrix when space and time are discrete. Building this matrix requires any ensemble of finite-time realizations of the flow. If the time *T* that the space-time POD modes are to evolve on is represented by *d* time steps with our temporal discretization, then one realization of the flow can be written as a vector with *d* snapshots stacked on top of one another,

$$\mathbf{y} = \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{q}_2 \\ \vdots \\ \mathbf{q}_d \end{bmatrix} \in \mathbb{R}^{Nd}.$$
 (41)

With many of these realizations, we can approximate the correlation matrix $\mathbf{C} \in \mathbb{R}^{Nd \times Nd}$ in the same manner as before,

$$\mathbf{C} = \frac{1}{m} \mathbf{Y} \mathbf{Y}^* \quad \text{with} \quad \mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \cdots \mathbf{y}_m] \in \mathbb{R}^{(Nd) \times m}.$$
(42)

Analogous to the space-only POD and SPOD cases, this approximation works because the correlation between two points in the spatio-temporal domain is approximated as the product of the flow at those points averaged over the realizations. We will see later in section 7, however, that if the flow is assumed to be ergodic and the realizations are all to be formed from one long time series, this construction of the correlation matrix does not fully exploit the ergodicty of the system, and it is possible to create a more accurate approximation of the correlation. The discretized POD modes are obtained as

$$\mathbf{CW}\boldsymbol{\Phi} = \boldsymbol{\Phi}\boldsymbol{\Lambda},\tag{43}$$

where $\mathbf{W} \in \mathbb{R}^{Nd \times Nd}$ is the discretized weight matrix.

Similar to the space-only POD and SPOD cases, the modes on this domain can also be obtained from the SVD of the weighted data matrix,

$$\frac{1}{\sqrt{m}}\mathbf{W}^{\frac{1}{2}}\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{*}.$$
(44)

The space-time POD modes and energies are then

$$\mathbf{\Phi} = \mathbf{W}^{-\frac{1}{2}}\mathbf{U} \quad \text{and} \quad \mathbf{\Lambda} = \mathbf{\Sigma}^2. \tag{45}$$

6 Hankel singular vectors are space-time POD modes

Here, we demonstrate the connection between space-time POD modes and Hankel modes by showing that one way of computing the unweighted space-time POD modes is to take the SVD of a Hankel matrix. As we will show, this connection provides insight into the sense in which the Hankel SVD modes are optimal and the impacts of the time step, rows, and columns on the properties of the modes.

For ergodic systems, the flow realizations needed to construct the approximate correlation matrix can be extracted from a single long time series. Given a time series of snapshots, $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_d, \dots, \mathbf{q}_{m+d-1}$, we may extract *m* realizations of length *d* by creating the first realization $\mathbf{y}_1 = [\mathbf{q}_1^T \mathbf{q}_2^T, \dots, \mathbf{q}_d^T]^T$, then advancing one time step over to create the second realization, $\mathbf{y}_2 = [\mathbf{q}_2^T \mathbf{q}_3^T, \dots, \mathbf{q}_{d+1}^T]^T$, and so on. Stacking these as columns in the data matrix, the result is a block Hankel matrix,

$$\mathbf{H} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m] = \begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \cdots & \mathbf{q}_m \\ \mathbf{q}_2 & \mathbf{q}_3 & \cdots & \mathbf{q}_{m+1} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{q}_d & \mathbf{q}_{d+1} & \cdots & \mathbf{q}_{m+d-1} \end{bmatrix}.$$
(46)

Comparing (46) and (42) we see that the Hankel matrix **H** provides one way of generating the data matrix **Y**. Therefore, using (44) and (45) and taking the SVD of the (weighted) Hankel matrix,

$$\frac{1}{\sqrt{m}}\mathbf{W}^{\frac{1}{2}}\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{*},\tag{47}$$

we obtain the space-time POD modes as

$$\mathbf{\Phi} = \mathbf{W}^{-\frac{1}{2}}\mathbf{U} \quad \text{and} \quad \mathbf{\Lambda} = \mathbf{\Sigma}^2. \tag{48}$$

If the weight **W** is uniform, the modes come directly from the Hankel SVD,

$$\mathbf{\Phi} = \mathbf{U} = \mathrm{lsv}(\mathbf{H}). \tag{49}$$

That is, the left singular vectors of the Hankel matrix give a discrete approximation of the space-time POD modes *in the case of a uniform weight*.

This connection to space-time POD implies the sense in which the Hankel modes are optimal: they minimize the expected value of the square norm of the projection error, where the norm is the ℓ_2 norm, i.e., the weight is uniform. Uniformity is a reasonable choice for the weight matrix if the spatial grid is uniform. However, if the grid is non-uniform, a uniform weight matrix corresponds to a non-uniform continuous weight. For example, if the flow data is defined on a cylindrical grid, a uniform weight matrix will bias accuracy toward the center of the domain over the outside. If the desired weight is non-uniform, as, e.g., is likely to be the case if the grid is non-uniform in space, then the discrete space-time POD modes can be obtained using Eqs (48) and (47).

6.1 Height of the Hankel matrix corresponds to T

The connection between space-time POD and singular modes of the Hankel matrix reveals two interpretations of the number of rows *d* in the Hankel matrix. First, it determines, for a fixed time step Δt , the time window *T* on which the space-time POD modes optimally represent the flow. The window also depends on the time step, specifically, $T = (d - 1)\Delta t$. Second, the significance of the choice of *d* can also be understood in terms of the space-time correlation matrix implied by the Hankel matrix—the Hankel singular modes will account for a section of the space-time correlation of the system of width *T*, as shown in Fig 1. The properties of the correlation outside of this window are discarded; they play no role in the definition of the modes and cannot be represented by the modes.

Two interesting limits of this truncation occur when $(d-1)^* \Delta t < \tau_C$ and when $(d-1) \Delta t > \tau_C$, where τ_C is the maximum correlation time of the flow between any two locations in space. When the height of the Hankel matrix corresponds to a time much smaller than the correlation time, we should expect to get spatial modes identical to those from discrete space-only POD for the same reason discussed in section 5.3: the discrete space-time POD modes are eigenvectors of the correlation matrix, and if $(d-1)^* \Delta t = T < \tau_C$ then this correlation



Fig 1. The height of the Hankel matrix d, along with the time step Δt between snapshots, determines the time window T over which the space-time correlation C of the system is accounted for in the Hankel singular modes. Shorter Hankel matrices, therefore, severely truncate the correlations, whereas taller Hankel matrices retain more of these correlations.

https://doi.org/10.1371/journal.pone.0289637.g001

matrix is constant in time. In this limit, the correlation between two points in space is independent of the time difference, so

$$\mathbf{C}_{ij} = \mathbf{C}_{i+aN,j+bN},\tag{50}$$

where N is the number of spatial gridpoints and a and b are any integers, so long as the indices are valid. Incrementing one of the indices by N corresponds to moving one time step forward but not shifting in space, but we have assumed that the correlation matrix is constant in time. It can therefore be written in block form as

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}^{\mathbf{x}} & \mathbf{C}^{\mathbf{x}} & \cdots & \mathbf{C}^{\mathbf{x}} \\ \mathbf{C}^{\mathbf{x}} & \ddots & & \vdots \\ \vdots & & & \\ \mathbf{C}^{\mathbf{x}} & \cdots & \mathbf{C}^{\mathbf{x}} \end{bmatrix},$$
(51)

where $\mathbf{C}^{\mathbf{x}} \in \mathbb{R}^{N \times N}$ is the discrete correlation matrix at no time lag, as defined in 11. **C** has only *N* non-zero eigenvalues, and these eigenvalues are $d\Lambda^{\mathbf{x}}$, the space-only eigenvalues scaled by the number of time steps. The associated eigenvectors of **C** are

- /x -

$$\phi_{k} = \begin{pmatrix} \phi_{k}^{\mathbf{x}} \\ \phi_{k}^{\mathbf{x}} \\ \vdots \\ \phi_{k}^{\mathbf{x}} \end{pmatrix}, \qquad (52)$$

where ϕ_k^x is the k^{th} eigenvector of the space-only correlation \mathbf{C}^x . Since each $N \times 1$ block of ϕ_k is the same, the space-time POD modes are constant in time and have the spatial form of space-only POD. Thus, we've recovered space-only POD in the short-time limit, as we might expect from the discussion of the continuous case in section 5.3. The SVD of a very short Hankel matrix will therefore give modes with little time dependence, which could have been obtained with less computational effort using space-only POD instead.

Above, as well as in section 5.3, we assumed no temporal variation of the correlation function. Gibson et al. [35] showed that for the special case of a scalar Hankel matrix, if small variations in time to the correlation matrix are allowed, then the resulting POD basis is the Legendre polynomials. In intuitive terms, the Legendre polynomials appear because, over short times, the flow will be well approximated by its Taylor series truncated at a few terms, and each term is much more important than the next. Therefore, with *k* modes, the basis should span the set of degree k - 1 polynomials and also be orthogonal with respect to the inner product. If the inner product is uniform, then the basis satisfying these conditions is the Legendre polynomials. On a vector time series, we observe a full set of modes with no time dependence, consistent with the order zero Legendre polynomial, followed by full sets of modes with time dependence corresponding to each subsequent Legendre polynomial and negligible energy. The constant modes are the space-only POD modes, and their energies are the space-only POD energies multiplied by *d*. This confirms that our analysis assuming no temporal variation, and the resulting conclusions that the space-time modes converge to space-only modes, hold.

The other limit, $(d - 1)\Delta t = T > \tau_C$, occurs when the height of the Hankel matrix corresponds to a time much longer than any correlation time in the flow. From section 5.2, we would expect that the modes oscillate with a pure frequency as SPOD modes do. Indeed, taking the

SVD of a Hankel matrix becomes a discrete Fourier transform as the height of the Hankel matrix grows to infinity [36, 37]. This limit has been used in the context of time-delay embedding [38].

This connection to SPOD means that Hankel SVD is also related to two other popular methods, resolvent analysis and dynamic mode decomposition (DMD). Resolvent analysis is an input-output formulation of the Navier-Stokes equations in the frequency domain [39, 40] wherein the non-linear terms (about some base/mean flow) are interpreted as forcing inputs to a linearized system. Towne et al. [23] showed that if these forcings do not preference any spatial structure, i.e., they are white in space, then resolvent modes are equivalent to SPOD modes. Hence, due to the connection we have established, resolvent modes are SVD modes of an asymptotically tall Hankel matrix if the non-linear terms are white. DMD is a data-driven method that seeks the linear operator which most closely maps snapshots at one time to their value one timestep later, in the least-squares sense. SPOD modes are optimally averaged DMD modes [23], so too then are Hankel SVD modes in the case that the Hankel matrix is asymptotically tall.

Finally, we note that the time step between rows (along columns) signifies the temporal discretization of the continuous eigenvalue problem (33) which defines the modes. Next, we leverage the distinct meaning of the time step between rows and columns to approximate the modes at reduced cost.

6.2 Width of the Hankel matrix informs convergence

The connection with space-time POD also reveals the impact of the width of the Hankel matrix on its singular modes. The convergence of the discrete space-time POD modes is determined by the accuracy of the approximation of the correlation matrix formed from the data. Each element of the correlation is proportional to the sample average of products of one entry in the column with another, averaged over the columns of the Hankel matrix,

$$\mathbf{C}_{ij} = (\mathbf{H}\mathbf{H}^*)_{ij} = \sum_{k=1}^{m} \mathbf{H}_{ik} \mathbf{H}_{jk}.$$
(53)

Thus, the width of the Hankel matrix *m* determines the number of realizations that contribute to approximating the correlation.

However, because two adjacent columns in the Hankel matrix are only one time step apart, their contributions to each element of the correlation matrix are far from independent, so the accuracy of the correlation is not just a function of the number of columns. Indeed, one could imagine a Hankel matrix that has many columns but only represents data over a short time during which the underlying system does not fully explore its phase space. There are two criteria necessary for the modes to be accurate: there must be enough columns in the Hankel matrix so that their sample average converges, and the data must be representative of the underlying attractor. The latter may require the data to be taken over a long time, leading to a very wide Hankel matrix, thus a costly SVD. We point out here that the time step between columns signifies the time between two successive realizations, in contrast to the time step between rows.

Alternatively, we may throw out many of the columns of the Hankel matrix, so that the columns are less correlated, then take the SVD of this matrix,

$$\mathbf{Q}_{UC} = \begin{bmatrix} \mathbf{q}_{1} & \mathbf{q}_{s+1} & \mathbf{q}_{2s+1} & \cdots & \mathbf{q}_{m} \\ \mathbf{q}_{2} & \mathbf{q}_{s+2} & \mathbf{q}_{2s+2} & \cdots & \mathbf{q}_{m+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{q}_{d} & \mathbf{q}_{s+d} & \mathbf{q}_{2s+d} & \cdots & \mathbf{q}_{m+d-1} \end{bmatrix}.$$
(54)

Here, *s* represents the spacing in time between columns, and the Hankel matrix is recovered if s = 1. $\mathbf{Q}_{UC}\mathbf{Q}_{UC}^*$ still forms an approximation of the correlation matrix, so its singular vectors will approximate the space-time POD modes. If **H** is formed from data representative of the attractor, then \mathbf{Q}_{UC} will be as well, and if we retain enough columns, the correlation matrix will be accurate. An appropriate choice of *s*, one which negligibly impacts the accuracy of the modes, depends on the amount of data available, the system dimension, the embedding dimension, and the time step. The Hankel matrix is the data matrix with the most columns possible from the data, and in practice it is likely overkill, and overly costly, if the time series is long. If the data is lacking, however, the Hankel matrix will produce a more accurate correlation and hence more accurate modes than any other spacing.

This strategy of removing most of the columns provides a substantial cost reduction. The time complexity of the SVD scales quadratically with the smaller dimension of the matrix and linearly with the larger one. Therefore, this strategy reduces the computation time by a factor of s^2 if the Hankel matrix is taller than it is wide, which is usually the case for fluids and other high-dimensional problems, and by a factor of *s* otherwise. Critically, this reduction is achieved without changing the window *T* or the time step Δt along the columns, such that the meaning of the modes discussed in section 6.1 and the temporal discretization of the integral in (33) remain unchanged.

Obtaining the space-time POD modes from a long time series of data is, in some ways, analogous to finding the power spectrum of a signal from a single long time series. It is well known [41] in the latter case that one should break the long time series into many shorter, possibly overlapping blocks, and then average the specta obtained from each block. The ideal overlap is determined by the analyst's trade-off preferences between accuracy and cost and the particularities of the problem, but little is gained by going beyond, say, 90% overlap. Certainly, it would be ill-advised to use the maximum overlap, i.e., shifting each block by only one time step. But, in the case of obtaining space-time POD modes, using the full Hankel matrix is imposing such an excessive overlap. This maximum overlap should only be used if data is scarce and forming the Hankel matrix and taking its SVD do not pose time or memory problems.

A cartoon of the correlations produced from uncorrelated columns vs. Hankel columns with different amounts of data is shown in Fig 2. The top two correlation graphics show how the uncorrelated columns (left) and Hankel (right) approaches might approximate the correlation with limited data. In this case, the Hankel matrix is likely to produce more accurate correlations than the matrix with uncorrelated columns because the latter may have too few to converge averages. However, the limited data is not representative of the underlying attractor, so although sampling this data more and more (Hankel approach) and computing the correlations may converge, it will not converge to the true correlations. When the time series is long enough to be representative of the attractor (bottom), both approaches will produce accurate correlations so long as enough columns are included in the uncorrelated matrix to converge averages. We also note that in the case of a short time series, we derive a method in section 7 which produces a more accurate approximation of the correlations than the Hankel approach with the same data.

6.3 Summary of the connection

To summarize, we have shown that Hankel singular modes constitute an approximation of space-time POD modes. The block Hankel matrix is formed as a data matrix whose columns are discretized temporal flow realizations, and when multiplied by its conjugate transpose, gives the correlation matrix from space-time POD. The eigenvectors of this matrix, which are



Fig 2. Uncorrelated columns (left) and Hankel (right) approaches to approximating the correlation matrix from short (b,c) and long (d,e) time series. (a) A time series used to sample temporal realizations with different samplings. (b) Correlations from uncorrelated columns using a short time series. (c) Correlations from the Hankel matrix using the short times series. (d) Correlations from uncorrelated columns using a long time series. (e) Correlations from a Hankel matrix using the long time series. Though both (d) and (e) are accurate, the correlations in (d) come at significantly lower computational cost.

the space-time POD modes, are the left singular vectors of the Hankel matrix due to the wellknown equivalence of the SVD and eigendecomposition. This connection sheds light on the sense in which the Hankel modes are optimal: they are optimal in the case of a uniform weight, which is often an undesirable norm. Understanding Hankel singular modes as approximate space-time POD modes also reveals that the time step between rows and columns of the Hankel matrix need not be the same. Indeed, in many practical cases, discarding many of the columns of the Hankel matrix negligibly impacts the accuracy of the modes but significantly



Fig 3. The dimensions of the Hankel matrix can be interpreted in terms of the integral space-time eigenvalue problem (33). The height of the Hankel matrix corresponds to *T*, the length on which the space-time POD modes optimally represent the flow. The width, though there is subtlety about independence, informs how accurately the correlation matrix is approximated, and hence the accuracy of the modes.

reduces the time for the SVD. The time window corresponding to the height of the Hankel matrix $T = (d - 1)\Delta t$ is the time over which the Hankel modes optimally represent the flow as well as the window of space-time correlations accounted for in calculating the modes. Finally, in the limits of short and tall Hankel matrices, corresponding to evolution on short and long time windows, the Hankel modes become discrete space-only POD modes and discrete SPOD modes, respectively. Much of the discussion in this section is summarized in Fig 3.

7 Fully exploiting ergodicity: More accurate modes

As discussed in section 6, the Hankel matrix implies a particular approximation of the spacetime correlation, which in turn determines the approximation of space-time POD modes provided by singular modes of the Hankel matrix. In this section, we show how a more accurate approximation of the correlation matrix can be constructed by fully exploiting the ergodicity of the underlying system, which produces more accurate modes if there is a shortage of data, as we will demonstrate later in section 8.6. This approach has been used in the context of SSA [4], but to our knowledge, has never been employed for vector-valued data.

Ergodicity is assumed in constructing the Hankel matrix: instead of taking columns from different sample paths of the flow, columns are formed from different sections in one long sample path, so there is an assumption that these two are equivalent. However, in constructing the Hankel matrix (or a down-sampling thereof) and then taking its SVD, the assumed ergodicity is not fully exploited. To see this, look at the correlation $\mathbf{C} = \mathbf{H}\mathbf{H}^* \in \mathbb{R}^{Nd \times Nd}$. Broken up into its spatial block structure, this correlation is written

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} & \mathbf{C}_{13} & \cdots & \mathbf{C}_{1d} \\ \mathbf{C}_{12}^{T} & \mathbf{C}_{22} & \mathbf{C}_{23} & \cdots & \mathbf{C}_{2d} \\ \mathbf{C}_{13}^{T} & \mathbf{C}_{23}^{T} & \mathbf{C}_{33} & \cdots & \mathbf{C}_{3d} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_{1d}^{T} & \mathbf{C}_{2d}^{T} & \mathbf{C}_{3d}^{T} & \cdots & \mathbf{C}_{dd} \end{bmatrix},$$
(55)

where the matrix

$$\mathbf{C}_{ii} \in \mathbb{R}^{N \times N} \tag{56}$$

represents the correlations between all points in the spatial domain between times *i* and *j*. Because the flow is ergodic, these correlations should only depend on the difference in times, e.g., C_{12} should be equal to C_{56} , but by calculating this correlation with the Hankel matrix, this will not normally be the case. Indeed, writing each of these correlation blocks in terms of elements of the time series,

$$\mathbf{C}_{ij} = \sum_{k=1}^{m} \mathbf{q}_{i+k-1} \mathbf{q}_{j+k-1}^{*},$$
(57)

we see, for example, that the (1, 1) block of **C** does not depend on the last element of the time series while the (d, d) block does, so they will not be equivalent as we know they should be from ergodicity. Instead, the fully converged correlation should have the symmetric block Toeplitz structure,

$$\tilde{\mathbf{C}} = \begin{bmatrix} \mathbf{C}_{0} & \mathbf{C}_{1} & \mathbf{C}_{2} & \cdots & \mathbf{C}_{d-1} \\ \mathbf{C}_{1}^{T} & \mathbf{C}_{0} & \mathbf{C}_{1} & \cdots & \mathbf{C}_{d-2} \\ \mathbf{C}_{2}^{T} & \mathbf{C}_{1}^{T} & \mathbf{C}_{0} & \cdots & \mathbf{C}_{d-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_{d-1}^{T} & \mathbf{C}_{d-2}^{T} & \mathbf{C}_{d-3}^{T} & \cdots & \mathbf{C}_{0} \end{bmatrix}.$$
(58)

That the (block) diagonals of $C = HH^*$ are not constant indicates that ergodicity is not fully exploited in constructing the correlation.

To calculate each correlation C_i from a time series, $q_1, q_2, \ldots, q_{m+d-1}$, we simply take the sample average of the product of all points in the time series *i* time steps apart,

$$\mathbf{C}_{i} = \frac{1}{m+d-i-1} \sum_{k=1}^{m+d-i-1} \mathbf{q}_{k} \mathbf{q}_{k+i}^{*}.$$
(59)

After building \hat{C} as in (58), the modes are obtained as $\hat{C}W\Phi = \Phi\Lambda$. The difference in accuracy of the modes can be substantial when data is in short supply and when m/d is not large, as we demonstrate in section 8.6.

The entries of the Toeplitz matrix (59) are formed using more terms than are used in forming the Hankel-based correlations (57). The ratio depends on the entry calculated but is proportional to $\frac{m}{m+d}$; therefore as m/d increases, the advantage in the Toeplitz approach diminishes.

The relative cost of computing modes using the Hankel and Toeplitz approaches depends on the relative size of *Nd* and *m*. The time scaling for the Hankel approach comes from the SVD in both cases, which is $\mathcal{O}(N^2d^2m)$ for Nd < m and $\mathcal{O}(Ndm^2)$ otherwise. The time scaling for the Toeplitz approach comes from passing through the time series to calculate all blocks of the correlation and from taking the eigendecomposition of the correlation. The former scales as $\mathcal{O}((m + d)dN^2)$ and the latter as $\mathcal{O}(N^3d^3)$. The ratio of the scalings of the Toeplitz and Hankel algorithms is thus

$$\int \max\left(\frac{1}{d}, \frac{Nd}{m}\right) \quad \text{if } Nd < m,$$

$$\left(\frac{Nd}{m}\right)^{2} \quad \text{if } Nd > m.$$
(60)

In fluids applications, Nd is usually much greater than m, so this approach will scale much worse than the Hankel approach. However, in many other dynamics applications, including many classic applications of the Hankel matrix, N is small, often 1, and d substantially smaller than m [13, 15], so this approach both leads to more accurate modes and better scaling. In short, the Toeplitz method scales more favorably than building the Hankel matrix and taking the SVD when Nd < m, but when Nd > m the algorithm is slower.

8 Numerical experiments: Lid-driven cavity flow

In this section, we demonstrate our theoretical results with data from a 2D lid-driven cavity flow at Reynolds number Re = 22,000 [42]. First, we compare the convergence of modes with a true Hankel matrix against a data matrix whose columns are uncorrelated, showing that with the same number of columns (flow realizations) the uncorrelated matrix generates better modes. We then show a related, but more practical, example. Given an m_{H^-} column Hankel matrix, one can form an uncorrelated data matrix by sampling m of its columns. We show that for m significantly smaller than m_H the modes are nearly identical, which can be used to reduce computational cost. Next, we demonstrate that space-time POD converges to space-only POD on very short time intervals, and that the modes become Fourier in time on very long time intervals, indicating convergence to spectral POD. Finally, we show that with limited data, exploiting the ergodicity of the system produces more accurate modes than the Hankel approach.

8.1 Simulation description

We generate data for square lid-driven cavity flow at $Re = \frac{Uh}{v} = 22,000$, where U is the speed of the lid, h is the height of the square cavity, and v is the viscosity of the fluid. The incompressible Navier-Stokes equations are solved using a Crank-Nicolson method for the viscous term and an Adams-Bashforth method for the nonlinear term [43]. The domain is discretized with $N_x = N_y = 120$ grid-points. Data is generated by starting the simulation with zero velocity everywhere except the top, running until the initial transients have vanished and the statistics become stationary, and then collecting the time series data. The time is nondimensionalized so that in one unit of time the lid slides the cavity width, and the simulation time step is 5×10^{-4} . We sample this data with $\Delta t = 0.1$. Fig 4 shows a snapshot of the flow to illustrate the setup and spatial scale of variation of the flow.

8.2 Spaced columns yield better modes

In section 6.2, we argued that an *m*-column data matrix with uncorrelated columns will produce a more accurate correlation matrix, and hence more accurate modes, than an *m*-column Hankel matrix, when measured against converged modes. Here, this is demonstrated with the lid-driven cavity flow: we form two data matrices—one Hankel, and the other uncorrelated, with the structure in (54). For the sake of memory, which scales like *Nmd*, we choose a downsampled version of *u*, the *x* component of velocity, as our observable of the state **q**, with 24



Fig 4. Schematic of the lid-driven cavity flow. The lid moves to the right and drives the flow. Contours show a snapshot of the turbulent kinetic energy.

points in each direction. We choose T = 2, so that the modes evolve on the convective time scale of the flow, and $d = \frac{T}{\Delta t} + 1 = 21$. Each column of the matrices is an element of \mathbb{R}^{Nd} , where $Nd = 24 \times 24 \times 21 = 12$, 096. The space-time POD modes are obtained as the left singular vectors of the two matrices, and their convergence is evaluated by taking the square of their inner product with a fully converged mode, which is obtained from a data matrix with 40,000 columns, all far apart in time. The variation in different realizations of the chaotic system leads to variation in the accuracy of the modes calculated from them. To account for this, we repeat the process of finding the modes 400 times for each *m* and take the mean square inner product with the fully converged mode.

Fig 5 shows the results. As expected, the modes from the uncorrelated data matrix outperform those from the Hankel matrix. The *x*-axis is logarithmic, and the two curves are roughly a constant horizontal distance from one another. This means that to achieve the same accuracy as an uncorrelated data matrix with *m* columns, a Hankel matrix must have *cm* columns, where *c* is some constant independent of *m*; here it is roughly 8, though this number depends sensitively on the parameters used.

8.3 Faster computation by sampling the Hankel matrix

Choosing between taking the SVD of a Hankel matrix or an uncorrelated data matrix of the same size may be unrealistic—simulation or experimental data can be in short supply, and it is not trivial to simply generate the extra data needed to construct the uncorrelated data matrix. A more relevant question might be: given a time series of snapshots, should one form the entire Hankel matrix or can many of the columns be omitted? The latter has the advantage of reducing computational cost, in terms of both CPU time and memory requirements, while the former may be more accurate if data is scarce. The accuracy loss depends on how many columns are removed, but we demonstrate here that the loss of accuracy can be negligible for orders of magnitude computational speedup.

Again, we choose T = 2, so d = 21. Starting with a Hankel matrix with $m_H = 5000$ columns, we form a reduced data matrix by retaining only *m* columns with the maximum possible



Fig 5. Convergence of the modes with the width *m* of the matrix: (blue) Hankel matrix; (red) uncorrelated data matrix. The modes from the uncorrelated data matrix converge faster than those from the Hankel matrix. This improved convergence is observed because the columns of the uncorrelated data matrix represent data from more of the attractor than the columns of the Hankel matrix of the same size.

spacing. For example, if m = 51, we use every 100th column of the Hankel matrix to form the data matrix. Removing columns changes the time step along rows but not along columns, and our analysis in section 6.2 indicates that this only changes the rate of convergence of the modes, not what the modes converge to. We take the SVD of both the Hankel and data matrices and compare each mode from the data matrix to its respective Hankel matrix mode using the square of the inner product, i.e., $|\langle \phi_k^H, \phi_k^D \rangle|^2$ where ϕ_k^H is the *k*-th mode from the Hankel matrix and ϕ_k^D is the *k*-th mode from the data matrix; values of 1 and 0 indicate identical modes and orthogonal modes, respectively. We repeat this process 200 times to see how well, on average, the modes from the much smaller data matrix approximate the modes from the full Hankel matrix.

Fig 6 shows this accuracy metric as a function of m for the first three modes from each matrix, reporting both the mean (solid) and the median (dashed) of the 200 trials. Both the first and second modes from the m = 500 data matrix capture over 97% of the energy of the corresponding Hankel modes (mean), and over 99% (median). For m = 1000, the first and second modes both capture over 99.9% of the corresponding Hankel modes. Even these modest reductions in m offer significant computational savings, scaling here like the square of the reduction in m.

The ratio of the number of columns in the sampled Hankel matrix to that of the full Hankel matrix needed to achieve a given accuracy is not constant. For example, if the time step Δt is smaller, a sparser sampling of the Hankel matrix can be used to achieve the same accuracy because *k* columns of the Hankel matrix represent less of the underlying flow. The speedup gained by sampling the Hankel matrix can be significant; for problems with a complicated spatial domain, the height of the data matrices, *Nd*, will likely be much larger than the number of columns, *m*. Therefore the time complexity of the SVD is m^2 , and reductions in *m* save considerable time. We reiterate that this cost savings is enabled by the insight, provided by the



Fig 6. Similarity between modes from a downsampled Hankel matrix, retaining m/m_H columns, and the Hankel modes. Both the mean (solid) and the median (dashed) are reported for the 200 trials. The first mode (blue) and the second mode (red) are more similar than the third mode (green). For the parameters used, retaining 1/10 and 1/5 of the columns yields modes that capture over 97% and 99.9% of the energy of the Hankel modes, respectively, for both the first and second modes, on average (mean).

connection between Hankel singular modes and space-time POD, that the time step along rows and columns have different meanings and need not be equal.

8.4 Convergence to space-only POD for short times

Here, we demonstrate the convergence of the space-time POD modes to space-only POD modes as the time interval they are defined on approaches zero, $T \rightarrow 0$. Space-only POD modes are defined on Ω and space-time POD modes are defined on $\Omega \times [0, T]$, so in order to compare the two, we take the square of the space-only inner product, averaged over time. In Fig 7, we see that the spatial part of the space-time modes indeed converges to space-only modes as $T \rightarrow 0$. In fact, the square inner product of the two remains near unity even when the space-time modes are defined over a time on the order of the convective time-scale; when T = 1 it is 0.937, and it only drops below 0.9 when T > 1.4.

8.5 Convergence to spectral POD for long times

Here, we demonstrate that as the time interval becomes long, the time dependence of the modes becomes Fourier. Modes are computed via the SVD of a data matrix with large *T* and 800 well-spaced columns. We do this for different values of *T* and examine the frequency content of the modes. The expectation is that the modes will become delta functions in the frequency domain as *T* increases since we know that SPOD modes, which have purely Fourier time dependence, are the limits of space-time POD for long time windows. The spectral content of the space-time modes may vary over the spatial extent of the mode, so we use the norm of their Fourier component,

$$\|\hat{\phi}(\omega)\|^2 = \langle \hat{\phi}(\omega, \mathbf{x}), \hat{\phi}(\omega, \mathbf{x}) \rangle, \tag{61}$$





which is equivalent to the power spectral density of the mode integrated over the domain, as a measure of their overall spectral content. Fig 8 demonstrates the convergence with increasing T of the spectra of the leading mode to a delta function in frequency.

8.6 Fully exploiting ergodicity

Finally, we demonstrate that the modes from the Toeplitz correlation matrix offer a significant improvement over those from the Hankel data matrix. We extract the modes by taking the SVD and eigendecomposition of the Hankel and Toeplitz matrices, respectively, and calculate the energy of the flow captured by the first few modes of both methods. We repeat this many times to generate PDFs of the energy captured, observing that the Toeplitz modes capture more energy. We also compare the modes from both methods to fully converged Hankel modes and tabulate the results for different parameters, again finding that the Toeplitz modes are more accurate.

Energy of the space-time POD modes, (28), is an appropriate metric for evaluating the performance of the two methods for obtaining modes because it is the quantity optimized in the definition of the POD problem. Analogous to (9), we may rewrite the energy of a mode in terms of its projection coefficients onto the exact modes,

$$\lambda[\phi] = \sum_{k=1}^{\min(m,Nd)} c_k^2 \lambda_k \quad \text{with} \quad c_k = \langle \phi, \phi_k \rangle / \|\phi\|, \tag{62}$$

where ϕ_k is the k^{th} space-time POD mode. We calculate the energy of the modes produced by the Hankel and Toeplitz methods using the above formula, where the 'true' modes are calculated using a Hankel matrix built from a 40, 000-long time series with the same *N* and *d*.

Given *m* and *d*, and an m + d - 1-long time series, we compute the modes using the $Nd \times m$ Hankel matrix and the $Nd \times Nd$ Toeplitz matrix and calculate their energies. Forming the



Fig 8. Convergence of space-time POD to spectral POD, visualized in terms of the power spectral density of the space-time POD modes for various values of T: (blue) T = 50; (red) T = 100; (green) T = 200; (purple) T = 400. As T is increased, more and more of the energy of the modes is contained to an increasingly narrow band of frequencies, indicating convergence toward a discrete frequency, consistent with the time dependence of SPOD modes.

Toeplitz matrix and taking its eigendecomposition is infeasible for large systems, so here, we take the *N* points in the simulation with the highest variance for small values of *N*. Of course, the accuracy of the modes from either method depends on the time series used, so we repeat this process 1000 times, generating PDFs for different choices of *m*, *d*, and *N*. A PDF with significant support only near unity indicates a high probability of producing an accurate mode.

Fig 9 shows the PDF for the energy captured by the first mode from each method as a fraction of the energy of the true optimal mode (λ_1). Beginning with (a) in Fig 9, we see that the mean of the modes from the Toeplitz method is higher than that of the modes from the Hankel method. Both methods give a bimodal distribution, and these peaks represent instances where the method mistakes the true second mode for the first mode. There is a deep trough between the peaks of the Toeplitz method, indicating that it produces a mode close to the first or to the second true mode rather than some combination thereof. The Hankel method displays some of this behavior, but it is significantly less pronounced, indicating that it produces something between the two true modes more often than the Toeplitz method. This behavior is also present in (b), (c), and (d).

As expected, for a fixed value of *d* and *N*, increasing *m* yields a higher probability of more accurate modes, and the modes from both methods in (c) are more accurate than those in (b), which are more accurate than those in (a). We also see that as $\frac{m}{d}$ increases, the difference in the accuracy of the two methods decreases, and the PDFs of the two methods become closer going from (a) to (b) and from (b) to (c). We see this behavior because for $\frac{m}{d} >> 1$, all of the elements on the same diagonal of $\mathbf{C} = \mathbf{H}\mathbf{H}^*$ are sums of mostly the same products of terms from the time series, so **C** is nearly Toeplitz, whereas for $\frac{m}{d} \approx 1$ this is not true. In (e), $\frac{m}{d} = 50$, and the difference between the methods is hard to see by looking at the PDFs, though there is still some difference in the means and medians, especially for the second mode, as reported in Fig 10.



Fig 9. PDFs of energy captured by the first modes of the Hankel (blue) and Toeplitz (red) methods for various choices of m, d, and N. Both methods use the same time series of length m + d - 1 to approximate the modes. The PDFs are calculated with 1000 samples and the energies are divided by the maximum possible energy of one mode. These parameter values are, (a): m = 30, d = 30, N = 1. (b): m = 100, d = 30, N = 1. (c): m = 300, d = 30, N = 1. (d): m = 300, d = 30, N = 3. (e) m = 600, d = 30, N = 3. (f) m = 10, d = 10, N = 1.

For fixed values of d and m, increasing N worsens the accuracy of the modes, and the modes in (d) are less accurate on average than those in (b). Similarly, increasing d worsens the accuracy of the modes, as can be seen by comparing (a) and (f). In fact, the modes in (f) are more accurate from both methods than those in (a) despite being calculated with lower m. Increasing N or d decreases the quality of the modes because it causes them to be in a higher dimensional space, which makes them more difficult to approximate. If m is large enough, however, (signifying a long time series from which to calculate the modes), both methods will converge for any N and d.

To evaluate the accuracy of the latter modes, say the k^{th} , we add the energies of the modes up to the k^{th} . We do this instead of simply calculating $\lambda[\phi_k]$ because this would erroneously reward switching the first and the k^{th} mode. We denote the sum of the energies as

$$\lambda[\phi_{1,2,\cdots,k}] = \sum_{i=1}^{k} \lambda[\phi_i], \tag{63}$$

and this quantity is interpreted as the energy captured by the first k modes. Again, we generate



Fig 10. Mean (circles) and median (squares) accuracy of Hankel modes (blue) vs. Toeplitz modes (red) calculated from 1000 time series of the most energetic N points of the lid-driven cavity flow. Accuracy is measured as the square inner product against a 'true mode.' When the data is severely limited (*m* is relatively small), the difference is more pronounced.

PDFs of these energies using 1000 time series of length m + d - 1, computing modes using both methods, and calculating the energy they capture. Fig 11 shows the results for the first four modes using m = 30, d = 30, N = 1. The energy captured by the first two Toeplitz modes is significantly greater than that of the first two Hankel modes. The same can be said about the first three and first four Toeplitz modes. These observations hold for the majority of parameter combinations.

A second metric for evaluating the accuracy of the modes is to compare them directly to fully converged modes, via the square inner product. Again a value near unity indicates an accurate mode, though a value near zero for, e.g., the third mode does not indicate that that mode captures no energy, just that it captures energy orthogonal to that of the converged third mode. We calculate these converged modes from the 40, 000-long time series. For each choice of *m*, *d*, and *N*, we calculate the modes from the Hankel and Toeplitz matrices, take the inner product with the converged modes, and repeat 1000 times. We record the mean and median of the square inner products for the first three modes in Fig 10.

As *d* increases, the number of realizations *m* needed to get accurate modes also increases for both methods. Similarly, as *N* increases, *m* may need to be increased to retain accuracy. Intuitively, this tells us that we need more data to find a mode in a higher dimensional space. The improvement of the Toeplitz method decays with m/d, because if this quantity is large, the difference between **HH**^{*} and the Toeplitz \tilde{C} becomes small, as discussed in section 7. Again,



Fig 11. PDFs of the energy captured by the first four modes of the Toeplitz and Hankel methods, normalized by the maximum possible energy by four modes. These are calculated for m = 30, d = 30, N = 1. The difference is more severe for the latter modes, which we observed for the majority of the parameter choices that we tested.

we note that the Toeplitz modes are significantly more accurate when m is not significantly greater than d. Also, there is not a big difference when m is very large, because here the modes are converged. Because a large N necessitates a large m for accuracy, m/d will be large and thus the improvement gained by the Toeplitz method is marginal for large N. However, these are the cases for which the time complexity of the Toeplitz method is already prohibitive.

9 Conclusions

We have demonstrated that the singular modes of the Hankel matrix are a discrete approximation of the continuous space-time POD modes. That is, in the language of Dylewsky et al. [15], principal component trajectories are in fact an approximation of classical space-time POD modes defined on a finite temporal window. This connection is made by recognizing that the Hankel matrix multiplied by its conjugate transpose produces an approximation of a spacetime correlation matrix, and the eigendecomposition of this matrix gives discrete space-time POD modes. We are able to gain useful insights into the Hankel modes by analyzing this correlation matrix rather than the Hankel matrix itself. This analysis makes clear the sense in which the modes are optimal: they are optimal if the weight matrix is uniform, which is often not desirable due to a variety of reasons including grid non-uniformity, the need to weight certain flow variables differently, or preference for certain regions in space. We derive a formula for the case of a non-uniform weight. The analysis of the correlation matrix also allows us to distinguish between the meaning of the time step along rows vs. along columns: the time step along columns corresponds to the temporal discretization of the space-time integral eigenvalue problem (33), while the time step along rows corresponds to the time between successive flow realizations used to approximate the correlation matrix. Viewed in this way, it is clear that these time steps need not be the same. We leverage this insight to ease computation, omitting many of the columns of the Hankel matrix then taking the SVD of this new data matrix with uncorrelated columns. This produces modes with the same meaning at lower cost. We also show that two popular versions of POD, space-only POD and spectral POD, are recovered in the limits that the columns of the Hankel matrix represent short and long times, respectively. Finally, we improve the Hankel modes by fully exploiting the assumption of ergodicity to form a more accurate correlation matrix with the same data. This correlation matrix has Toeplitz structure owing to the fact that temporal correlations should only depend on the time separation, not on the times explicitly. We show that, especially if data is limited, the modes obtained as the eigenvectors of this Toeplitz matrix can be substantially more accurate than the Hankel modes.

We envision that this work will make Hankel-SVD-based methods applicable to new problems. In particular, the strategy of omitting many of the columns of the Hankel matrix may enable application of these methods to high-dimensional systems such as those found in fluid dynamics. Likewise, the improved modes obtained by exploiting ergodicity may aid applications with severely limited time series data, such as expensive simulations. Furthermore, the theoretical connection we have established between space-time POD and Hankel SVD could lead to additional algorithmic improvements that further broaden the problems to which these methods are applied. Finally, given its connection to widely used Hankel methods, we hope this work will increase interest in space-time POD.

A number of questions remain. First, how does one, a priori, choose the number of columns to retain in the data matrix (how many to throw out from the Hankel matrix)? This depends on the data available, *d*, *N*, and, of course, the computational resources available and/or trade-off preferences between accuracy and speed. It also depends on the time step relative to the time scales of the flow. Second, to what extent does the improved convergence offered by the Toeplitz modes over the Hankel modes impact performance in applications of interest? Finally, to our knowledge, this study is the first to examine the convergence of space-time POD to space-only and spectral POD. By further analyzing the correlation matrix, it may be possible to make statements about when this convergence is to take place, e.g., that space-time POD performed on a window *T* where $C(0, T) \approx 0$ will produce modes near to spectral POD modes.

Supporting information

S1 File. Scripts to generate figures in this paper. (ZIP)

Acknowledgments

We thank Ms. Peijing Liu for her contributions during the early stages of this project.

Author Contributions

Conceptualization: Peter Frame, Aaron Towne.

Formal analysis: Peter Frame, Aaron Towne.

Funding acquisition: Aaron Towne.

Investigation: Peter Frame.

Software: Peter Frame.

Supervision: Aaron Towne.

Writing - original draft: Peter Frame.

Writing - review & editing: Aaron Towne.

References

- 1. C. Rowley, T. Colonius, and R. Murray, POD based models of self-sustained oscillations in the flow past an open cavity, AIAA Paper #2000-1969.
- Fazel M., Pong T. K., Sun D., and Tseng P., Hankel matrix rank minimization with applications to system identification and realization, SIAM J. Matrix Anal. Appl, 34 (2013), pp. 946–977. https://doi.org/10. 1137/110853996
- Silverman L., Realization of linear dynamical systems, IEEE Trans. Automat. Contr., 16 (1971), pp. 554–567. https://doi.org/10.1109/TAC.1971.1099821
- Vautard R. and Ghil M., Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series, Phys. D, 35 (1989), pp. 395–424. https://doi.org/10.1016/0167-2789(89)90077-8
- Juang J.-N. and Pappa R. S., An eigensystem realization algorithm for modal parameter identification and model reduction, J. Guid. Control Dyn., 8 (1985), pp. 620–627. <u>https://doi.org/10.2514/3.20031</u>
- Moore B., Principal component analysis in linear systems: Controllability, observability, and model reduction, IEEE Trans. Automat. Contr., 26 (1981), pp. 17–32. https://doi.org/10.1109/TAC.1981. 1102568
- 7. Rowley C., Model reduction for fluids, using balanced proper orthogonal decomposition, Int. J. Bifurc. Chaos, 15 (2005), pp. 997–1013. https://doi.org/10.1142/S0218127405012429
- Willcox K. and Peraire J., Balanced model reduction via the proper orthogonal decomposition, AIAA J., 40 (2002), pp. 2323–2330. https://doi.org/10.2514/2.1570
- SCHMID P. J., Dynamic mode decomposition of numerical and experimental data, J. Fluid Mech., 656 (2010), pp. 5–28.
- Arbabi H. and Mezić I., Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the Koopman operator, SIAM J. Appl. Dyn. Syst., 16 (2017), pp. 2096–2126. <u>https://doi.org/10.1137/17M1125236</u>
- Packard N. H., Crutchfield J. P., Farmer J. D., and Shaw R. S., Geometry from a time series, Phys. Rev. Lett., 45 (1980), pp. 712–716. https://doi.org/10.1103/PhysRevLett.45.712
- Takens F., Detecting strange attractors in turbulence, in Dynamical Systems and Turbulence, Warwick 1980, Rand D. and Young L.-S., eds., Berlin, Heidelberg, 1981, Springer Berlin Heidelberg, pp. 366– 381.
- Brunton S., Brunton B., Proctor J., Kaiser E., and Kutz J., Chaos as an intermittently forced linear system, Nat. Commun., 8 (2016).
- Kamb M., Kaiser E., Brunton S. L., and Kutz J. N., Time-delay observables for Koopman: Theory and applications, SIAM J. Appl. Dyn. Syst., 19 (2020), pp. 886–917. https://doi.org/10.1137/18M1216572
- Dylewsky D., Kaiser E., Brunton S. L., and Kutz J. N., Principal component trajectories for modeling spectrally continuous dynamics as forced linear systems, Phys. Rev. E, 105 (2022), pp. 015312-1– 015312-14. https://doi.org/10.1103/PhysRevE.105.015312 PMID: 35193205
- LUMLEY J. L., The structure of inhomogeneous turbulent flows, Atmospheric Turbulence and Radio Wave Propagation, (1967).
- Gordeyev S. and Thomas F., A temporal proper decomposition (TPOD) for closed-loop flow control, Exp. Fluids, 54 (2013). https://doi.org/10.1007/s00348-013-1477-7
- 18. Rosario Z., Towne A., and laccarino G., Dimension reduction for shape design insight, 01 2018.

- Schmidt O. T. and Schmid P. J., A conditional space-time POD formalism for intermittent and rare events: example of acoustic bursts in turbulent jets, J. Fluid Mech., 867 (2019), pp. R2-1–R2-12. https://doi.org/10.1017/jfm.2019.200
- O. T. Schmidt, P. J. Schmid, A. Towne, and S. K. Lele, Statistical description of intermittency and rare events via conditional space-time POD: Example of acoustic bursts in turbulent jets, tech. rep., Proceedings of the Center for Turbulence Research Summer Program, 2018.
- Sirovich L., Turbulence and the dynamics of coherent structures. i—coherent structures. ii—symmetries and transformations. iii—dynamics and scaling, Quart. Appl. Math., 45 (1987). <u>https://doi.org/10.1090/ gam/910462</u>
- Aubry N., On the hidden beauty of the proper orthogonal decomposition, Theor. Comput. Fluid Dyn., 2 (1991), pp. 339–352. https://doi.org/10.1007/BF00271473
- Towne A., Schmidt O. T., and Colonius T., Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis, J. Fluid Mech., 847 (2018), pp. 821–867. https://doi.org/10.1017/jfm.2018.283
- Aubry N., Holmes P., Lumley J. L., and Stone E., The dynamics of coherent structures in the wall region of a turbulent boundary layer, J. Fluid Mech., 192 (1988), pp. 115–173. <u>https://doi.org/10.1017/</u> S0022112088001818
- Rowley C. W., Colonius T., and Murray R. M., Model reduction for compressible flows using POD and Galerkin projection, Phys. D, 189 (2004), pp. 115–129. https://doi.org/10.1016/j.physd.2003.03.001
- Rowley C. W. and Dawson S. T., Model reduction for flow analysis and control, Annu. Rev. Fluid Mech., 49 (2017), pp. 387–417. https://doi.org/10.1146/annurev-fluid-010816-060042
- Holmes P., Lumley J. L., Berkooz G., and Rowley C. W., Turbulence, coherent structures, dynamical systems and symmetry, Cambridge Monographs on Mechanics, Cambridge University Press, 2 ed., 2012.
- Moin P. and Moser R. D., Characteristic-eddy decomposition of turbulence in a channel, J. Fluid Mech., 200 (1989), pp. 471–509. https://doi.org/10.1017/S0022112089000741
- A. Pollard, L. Castillo, L. Danaila, and M. Glauser, Whither turbulence and big data in the 21st century?, 08 2016.
- 30. Lumley J. L., Stochastic tools in turbulence, (1970).
- Cavalieri A. V. G., Jordan P., and Lesshafft L., Wave-packet models for jet dynamics and sound radiation, Appl. Mech. Rev., 71 (2019). 020802. https://doi.org/10.1115/1.4042736
- 32. Schmidt O. T., Towne A., Rigas G., Colonius T., and Brès G. A., Spectral analysis of jet turbulence, J. Fluid Mech., 855 (2018), pp. 953–982. https://doi.org/10.1017/jfm.2018.675
- Symon S., Illingworth S. J., and Marusic I., Energy transfer in turbulent channel flows and implications for resolvent modelling, J. Fluid Mech., 911 (2021), pp. A3-1–A3-27. <u>https://doi.org/10.1017/jfm.2020.</u> 929
- C. W. Rowley, Modeling, simulation, and control of cavity flow oscillations, PhD thesis, California Institute of Technology, 2002.
- Gibson J. F., Doyne Farmer J., Casdagli M., and Eubank S., An analytic approach to practical state space reconstruction, Phys. D, 57 (1992), pp. 1–30. https://doi.org/10.1016/0167-2789(92)90085-2
- **36.** Bozzo E., Carniel R., and Fasino D., Relationship between singular spectrum analysis and Fourier analysis: Theory and application to the monitoring of volcanic activity, Comput. Math. Appl., 60 (2010), pp. 812–820. https://doi.org/10.1016/j.camwa.2010.05.028
- Broomhead D. and King G. P., Extracting qualitative dynamics from experimental data, Phys. D, 20 (1986), pp. 217–236. https://doi.org/10.1016/0167-2789(86)90031-X
- Kaiser E., Kutz J. N., and Brunton S. L., De-entanglement of complex flows with time delay coordinates, Bull. Am. Phys. Soc., (2020).
- McKEON B. J. and SHARMA A. S., A critical-layer framework for turbulent pipe flow, J. Fluid Mech., 658 (2010), pp. 336–382. https://doi.org/10.1017/S002211201000176X
- Trefethen L. N., Trefethen A. E., Reddy S. C., and Driscoll T. A., Hydrodynamic stability without eigenvalues, Science, 261 (1993), pp. 578–584. <u>https://doi.org/10.1126/science.261.5121.578</u> PMID: 17758167
- Welch P.; The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, IEEE trans. audio electroacoust., 15 (1967), pp. 70–73. https://doi.org/10.1109/TAU.1967.1161901
- Cazemier W., Verstappen R. W. C. P., and Veldman A. E. P., Proper orthogonal decomposition and low-dimensional models for driven cavity flows, Phys. Fluids, 10 (1998), pp. 1685–1699. https://doi.org/ 10.1063/1.869686
- **43.** michio, Cfd101: 2d lid driven cavity flow.