

Applied Koopmanism

Marko Budišić, Ryan Mohr, and Igor Mezić

Citation: *Chaos* **22**, 047510 (2012); doi: 10.1063/1.4772195

View online: <http://dx.doi.org/10.1063/1.4772195>

View Table of Contents: <http://aip.scitation.org/toc/cha/22/4>

Published by the [American Institute of Physics](#)

Articles you may be interested in

[Sparsity-promoting dynamic mode decomposition](#)

Chaos **26**, 024103024103 (2014); 10.1063/1.4863670

Welcome to a

Smarter Search 

PHYSICS
TODAY

with the redesigned
Physics Today Buyer's Guide

Find the tools you're looking for today!

Applied Koopmanism^{a)}

Marko Budišić, Ryan Mohr, and Igor Mezić

Department of Mechanical Engineering, University of California, Santa Barbara, California 93106-5070, USA

(Received 11 June 2012; accepted 30 November 2012; published online 21 December 2012)

A majority of methods from dynamical system analysis, especially those in applied settings, rely on Poincaré’s geometric picture that focuses on “dynamics of states.” While this picture has fueled our field for a century, it has shown difficulties in handling high-dimensional, ill-described, and uncertain systems, which are more and more common in engineered systems design and analysis of “big data” measurements. This overview article presents an alternative framework for dynamical systems, based on the “dynamics of observables” picture. The central object is the Koopman operator: an infinite-dimensional, linear operator that is nonetheless capable of capturing the full nonlinear dynamics. The first goal of this paper is to make it clear how methods that appeared in different papers and contexts all relate to each other through spectral properties of the Koopman operator. The second goal is to present these methods in a concise manner in an effort to make the framework accessible to researchers who would like to apply them, but also, expand and improve them. Finally, we aim to provide a road map through the literature where each of the topics was described in detail. We describe three main concepts: Koopman mode analysis, Koopman eigenquotients, and continuous indicators of ergodicity. For each concept, we provide a summary of theoretical concepts required to define and study them, numerical methods that have been developed for their analysis, and, when possible, applications that made use of them. The Koopman framework is showing potential for crossing over from academic and theoretical use to industrial practice. Therefore, the paper highlights its strengths in applied and numerical contexts. Additionally, we point out areas where an additional research push is needed before the approach is adopted as an off-the-shelf framework for analysis and design. © 2012 American Institute of Physics. [<http://dx.doi.org/10.1063/1.4772195>]

A majority of methods from dynamical systems analysis, especially those in applied settings, rely on Poincaré’s geometric picture that focuses on “dynamics of states.” While this picture has fueled our field for a century, it has shown difficulties in handling high-dimensional, ill-described, and uncertain systems, which are more and more common in engineered systems design and analysis of “big data” measurements. This overview article presents an alternative framework for dynamical systems, based on the “dynamics of observables” picture. We present an overview of several approaches to studying dynamical systems using the Koopman operator, which holds promise to resolve these issues. The dynamics are analyzed by looking at evolutions of functions on the state space, rather than directly at state space trajectories. The evolution can be understood by expanding the function into a basis of eigenfunctions of the Koopman operator. The first approach is based on the Koopman modes (KMs), which generalize linear mode analysis from linear systems to nonlinear systems, while preserving global nonlinear features of the system, unlike, e.g., linearizations based on Taylor- and Fourier-expansions. The second approach identifies coherent structures in flows. An equivalence relation between points in the state space can be defined using spectral properties of the Koopman operator, where equivalent points correspond to initial conditions that behave statistically the same with respect to any observable. The third

approach we present introduces continuous quantifications of ergodicity and mixing, concepts existing in ergodic theory are traditionally treated as binary notions. Throughout the paper, we highlight examples from the literature using each of these concepts. Examples are taken from diverse areas such as fluid mechanics, fluid mixing, energy efficiency of buildings, power systems, and unmanned aerial vehicle (UAV) path-planning for search-and-rescue. A common trait of all the methods is that they do not require access to an analytical model of the system; the spectral properties of the Koopman operator can be constructed from measured or simulated data.

I. INTRODUCTION

Currently, dynamical system analysis and design primarily use the geometric picture, as put forth by Poincaré in his work on the three body problem. Much of the framework is built around notions from differential geometry, trajectories, and invariant manifolds. Such an approach has met with success in a variety of settings and, at this point, one hardly needs to justify the use of geometric theory when working on a particular problem.

However, the geometric viewpoint is ill-suited to many of the situations that are of interest in real systems. For example, for systems possessing hyperbolic regimes, the unstable manifolds give rise to locally exponentially divergent trajectories. Any noise or uncertainty in the system will

^{a)}We dedicate this paper to the memory of Jerrold E. (Jerry) Marsden whose support of our research was invaluable.

lead to multiple possible trajectories for an initial condition, with the width of the set trajectories initially expanding exponentially. In such cases, questions about the behavior of a specific trajectory are difficult to answer.

Systems with a large number of dimensions can be problematic as well, since many of the geometric arguments are only valid in a low number of dimensions, e.g., Bendixson's Criterion for determining the non-existence of periodic orbits in the plane. In some cases, these arguments can be extended, with difficulty, to an arbitrary number of dimensions. Even in these cases, however, a practical implementation limits them to a moderate number of dimensions. To handle high-dimensional systems, special symmetries or other conditions are required in order to effectively reduce the dimension to a manageable size. Furthermore, without access to explicit ordinary differential equations (ODEs), even basic geometric analysis is difficult to apply. If dynamical systems theory is to become an important field in the context of pressing problems such as "big data," tools need to be developed that are capable of handling high-dimensional, uncertain, and ill-described systems, as well as systems for which past time-evolution data is available, but for which no simple mathematical description can be determined.²⁰

This article presents a viewpoint that is at the intersection of applied ergodic theory and operator theory. These two fields can be used in applied settings to analyze and design dynamical systems, with many of the aforementioned difficulties being handled with a certain amount of elegance. In fact, when we study dynamical systems through certain linear operators, the *full nonlinear dynamics can be captured within a linear setting*. This linear setting allows the power of spectral analysis to be brought to bear on a (nonlinear) problem without sacrificing any information as required by other linearization techniques. Contrast this with the traditional spectral approach that only determines geometry locally in the state space. Additionally, in theory, the operator-theoretic approach works equally well whether the original state-space is low- or high-dimensional; the same techniques apply to both cases. The framework is also well suited to studying noisy systems because the primary object of interest is no longer the trajectory. Finally, and perhaps most significantly, the operators involved can be constructed, approximated, or analyzed using only simulation or experimental measurement data. This allows a certain black-box approach to the analysis which is quite useful in real problems where the practitioner may not have full knowledge of the system's internals.

As with any technique, however, there is a tradeoff in order to gain the above advantages. The operator-theoretic picture has no immediate connection to our physical intuition, making its meaning more difficult to comprehend. One's viewpoint must change from considering the evolution of points in the state space to considering the evolution of functions. Additionally, the new approach is inherently infinite-dimensional, even when the state space is finite-dimensional. This sacrifice is what allows the full information of a nonlinear system to be contained within a linear setting. Because of this, the implementation of any approximation is a more delicate issue. Finally, the associated numerical techniques are underdeveloped. Most of our

approaches employ direct computations which are little more than numerical implementations of proofs.

The two main candidates for the study of systems via operators are the Koopman operator and the Perron-Frobenius operator. In appropriate function spaces, they are duals to each other, so theoretically there should not be any distinction in working with one as opposed to the other. However, as mentioned previously, we must always include applied considerations. Questions arise such as how do we construct or represent the chosen operator from the problem description and given data? How well does a finite approximation represent the ideal theoretical picture? What part of intuition gained is due to numerical artifacts and what is real?

The Perron-Frobenius operator represents a "dynamics of densities" picture; it looks at groups of trajectories. One can think of this as watching the evolution of a mass distribution under the action of a flow. From a numerical perspective, construction of the operator relies on selecting a set of initial conditions and simulating forward for only a short time period, thus avoiding the compounding of numerical time-integration errors. Due to these short bursts, transient dynamics can be captured very well. However, much attention has been focused on computing invariant densities,¹⁴ which are infinite-time objects, through approximating the Perron-Frobenius operator by a Markov chain. The number of simulated initial conditions is dictated by the need to sample the region of interest well. In high-dimensions, both short- and long-time dynamics simulations require a mesh on the entire space. This can be true even in the case of a low-dimensional attractor. If we have *a priori* knowledge of the low-dimensional subspace the attractor lives in, then the mesh size can be restricted. However, for an arbitrary system, this knowledge may not be initially available, thus requiring the full mesh.

On the other hand, the Koopman operator presents a picture for the "dynamics of observables." The difference in viewpoints between the Perron-Frobenius and Koopman operators is similar to the Eulerian versus the Lagrangian viewpoint in fluid mechanics, with the Koopman picture corresponding to the Lagrangian viewpoint. It meant that measurements are made along trajectories. For the Koopman operator, the numerical construction relies on potentially fewer initial conditions, but requires longer run-times, which is more suitable to physical experiments. For example, when testing a jet engine, it is started from a relatively small number of initial conditions and run it over a long time rather than preparing thousands of initial conditions and running the engine for a few seconds for each initial condition. As long run-times required, the asymptotics are well-understood. However, more research is needed to understand the transients.

To visualize high-dimensional dynamical systems, we often restrict our attention to one, or a few, two-dimensional cross-sections in the state space and look at the invariant structures intersecting that slice. With the Perron-Frobenius operator, it is difficult to directly compute invariant densities on the slice of interest, since, in principle, it requires a computation of the invariant density for the entire state space as an intermediate step. For the Koopman operator, invariant

objects are attached to initial conditions, making it well-suited to visualizing structures on an arbitrary 2D cross-section in the state space. Initial conditions can be easily prepared on the slice and the invariants directly computed. In such cases, the number of initial conditions required to understand the dynamics is significantly reduced.

While operator methods, and specifically the Koopman and Perron-Frobenius operators, have much potential to deal with applied problems, these methods are all but absent from the applied and industrial settings, with due exceptions.^{16,35,36} To speed up the adoption of operator techniques in these domains, any new methodology needs to be able to leverage already existing data, instead of proposing both a new methodology and a new way to collect data. The “dynamics of observables” perspective, and specifically the Koopman operator, is used as it deals with measurements, i.e., observables, which are well-understood both theoretically and computationally. On the other hand, the Perron-Frobenius techniques would require working with representations of densities, which are often singular, especially in well-behaved engineered systems.

In this paper, we intend to describe three concepts, all under the umbrella of “dynamics of observables,” that show how this theory can be made useful for analysis and design. Contributions can be split between theoretical and applied contributions.

A. Theoretical contributions

1. Dynamical evolution of a system can be studied by looking at what is termed *Koopman mode analysis*. The concept is similar to normal mode analysis familiar from linear vibration theory. Koopman mode analysis starts with a choice of a set of linearly independent observables, or equivalently a vector-valued observable. The Koopman operator U is then analyzed through its action on the subspace spanned by the chosen observables. The observables are decomposed into projections onto the eigenspaces of U , and the evolution is a sum of terms composed of a product of three terms: (i) a part that is time-dependent and is determined by the eigenvalue (or frequency) associated with the eigenspace; (ii) an eigenfunction of U , which is a function of the initial conditions; (iii) the vector of the coefficients of the projection of the observables onto the eigenspaces, with the coefficients only being functions of the chosen observables. In this way, spectral analysis can be performed on nonlinear systems. This analysis is also used for model reduction.^{34,36}
2. The notions of the *ergodic quotients and eigenquotients* allow the Koopman operator to be used for the extraction and analysis of invariant and periodic structures in the state space.⁶ The points in the state space are grouped into invariant sets using level sets of eigenfunctions of the Koopman operator. Instead of set-theoretic framework, this approach is lifted to the analytical setting using the ergodic quotient and eigenquotient formalism. The eigenquotients are studied as subsets of particular Sobolev spaces, where their geometry gives insight into the structure of the state space, in spirit similar to analysis of Hamiltonian systems via Morse theory of the associated energy functions.

3. In the standard interpretation of ergodic theory, mixing and ergodicity are treated as binary concepts: a system is either mixing/ergodic or it is not. Both mixing and ergodicity can be formulated using spectral invariants of the Koopman operator. From a finite-time evolution of an arbitrary system, we can quantify how close its spectral invariants are to the “ideal” case, e.g., mixing or ergodic, and in this way formulate *continuous indicators of ergodicity and mixing*. The ergodicity defect⁴⁸ and the mixing norm³⁰ are examples of continuous indicators that extend the corresponding binary notions. Such a relaxation brings the concepts of ergodicity and mixing into an engineering context, allowing, e.g., the use of the indicators as optimization criteria. We present a unified explanation of the concepts that have previously appeared in literature.^{28–31,45,48}

B. Numerical techniques and applications

Numerical computation of the objects in the theory uses elements from three different areas. Fourier analysis based methods are useful for computing Koopman modes for dynamics on the attractor in addition to being essential for the construction of the eigenquotient. A variant of the standard Arnoldi algorithm based on companion matrices is also useful for computing part of the spectrum of the Koopman operator, a basic element of Koopman mode analysis. This variant does not require an explicit representation of the operator and only requires data, sequences of vectors coming from either simulations or experiments.

Koopman mode analysis has seen applications in fluids mechanics to extract spatial structures for the flow.^{8,43,47,49} Koopman modes have also found applications in the analysis of coherency and instabilities for power systems^{53–56} and in the field of building energy efficiency where they have been used for model validation and data analysis.^{16,18}

To compute the eigenquotients numerically, a set of observables is averaged along trajectories started at different initial conditions, obtaining a finite-dimensional representation of any eigenquotient. Such representations are analyzed with the aid of a diffusion maps algorithm,^{9,10} which computes a change of coordinates, *the diffusion modes*, for the eigenquotient. In the limit where infinitely many initial conditions were simulated for infinite time, such coordinate change would render any consequent analysis independent of the choice of the observables averaged during the computation. The scale-ordering of diffusion coordinates makes it practical to obtain a low-dimensional approximation of the eigenquotients.

The averaging along trajectories is used again to formulate continuous indicators for ergodicity and mixing, where the rate of approach of the averages along finite-time trajectories to the infinite limit is indicative of the underlying dynamics. Based on such indicators, the dynamics of the flows can be designed to match a particular statistical behavior, with applications in path-planning for vehicles²⁸ and mixing of fluids on micro-scales.³¹

II. NOTATION AND TERMINOLOGY

We start off by fixing some notation and terminology. Let us denote the state space by M and define dynamics on it

by the iterated map $T : M \rightarrow M$. Note that the set M can be an arbitrary set (possessing no structure) and T can be an arbitrary map on this set. Then the abstract dynamical system is specified by the couple (M, T) . Note that standard texts on ergodic theory study a specific case when M is a measurable space, with a σ -algebra \mathfrak{B} , and is T \mathfrak{B} -measurable. Additionally, transformation T is typically assumed to be measure-preserving, i.e., there exists a measure μ , the *invariant measure*, such that for any $S \in \mathfrak{B}$

$$\mu(S) = \mu(T^{-1}S), \tag{1}$$

with $T^{-1}S$ understood as the pre-image of S . The measure μ does not necessarily have a density function associated with it. For the formulation of the theory in this paper, we do not require the measurable framework, although when answering more specific questions, we might restrict ourselves to it, as it is the one most commonly encountered in applied dynamical systems.

We will be concerned with the behavior of observables on the state space. To this end, we define an observable to be a function $f : M \rightarrow \mathbb{C}$, where f is an element of some function space \mathcal{F} . For now, it is not necessary to specify any structure for \mathcal{F} . A concrete interpretation of an observable is that of a sensor probe for the dynamical system in question; we can access information about the system via the evolution of the observable's values. Instead of tracking the trajectory $\{p, T(p), T^2(p), \dots\}$, we now track the trace $\{f(p), f(T(p)), f(T^2(p)), \dots\}$. The description of the dynamics can then be concisely written down in the form of state and output equations, familiar to control theorists,

$$\begin{aligned} p_{n+1} &= T(p_n) \\ v_n &= f(p_n). \end{aligned} \tag{2}$$

The dynamical systems community mainly focuses on state space trajectories $\{p_n\}$, while the control systems community usually studies systems that have an additional input or disturbance terms, with the functions T and f taking particular forms that are common in engineered systems.

We define the (discrete-time) Koopman operator, $U_T : \mathcal{F} \rightarrow \mathcal{F}$, as

$$[U_T f](p) = f(T(p)), \tag{3}$$

i.e., it is a composition, $U_T f = f \circ T$, of the observable f and the iterated map T . When it is obvious which transformation gives rise to the Koopman operator, we will drop the dependence on T from the notation and write U instead of U_T . When \mathcal{F} is a vector space, U is a linear operator.

When M is a finite set, U is a finite-dimensional operator and can be represented by a matrix. However, when M is finite- or infinite-dimensional, U is generally infinite dimensional. Much of the time, we only have access to a particular collection of observables $\{f_1, \dots, f_K\} \subset \mathcal{F}$; these could be physically relevant observables arising naturally from the problem, or a (subset of a) function basis for \mathcal{F} . We can extend the Koopman operator to this larger space in the natural way: If $F = (f_1, \dots, f_K)^T \in \mathcal{F}^K$, then $U_K : \mathcal{F}^K \rightarrow \mathcal{F}^K$ is defined as

$$[U_K F](p) := \begin{bmatrix} [Uf_1](p) \\ \vdots \\ [Uf_K](p) \end{bmatrix}. \tag{4}$$

Hence $U_K = \otimes_1^K U$. With an abuse of notation, we generally write U_K as U . The space \mathcal{F}^K is the space of \mathbb{C}^K -valued observables on M . In this context, \mathbb{C}^K is referred to as the output space. More generally, we can consider vector valued observables, $F : M \rightarrow V$, where V is some vector space. For example, when analyzing the heat equation on a periodic box \mathbb{B} , the state space can be regarded as the sequence space of Fourier coefficients and an observable $F : M \rightarrow L^2(\mathbb{B}, dx)$ can be regarded as mapping between a sequence of Fourier coefficients (the state space M) and a temperature distribution on \mathbb{B} (the real-valued space $L^2(\mathbb{B}, dx)$). We will revisit this setup in more detail in Example 4.

The above notion of the Koopman operator was defined in the context of discrete-time dynamical systems. Often though, working in the discrete-time setting poses an unneeded restriction; in many systems, the natural formulation of the dynamics is with respect to a continuous time variable. The Koopman operator can be extended to deal with continuous-time dynamical systems, or even more generally, event-based dynamical systems.

Assume we have the continuous-time dynamical system $\dot{p} = T(p)$. In this context, there is not just the Koopman operator, but a semigroup of operators $\{U^t\}_{t \in \mathbb{R}^+}$ given by a generator U . We call the semigroup $\{U^t\}$ the *Koopman semigroup*. We explicitly define the action of the semigroup on the observable $f \in \mathcal{F}$ as

$$[U^t f](p) = f(\Phi^t(p)). \tag{5}$$

Here $\Phi^t(p) \equiv \Phi(p, t)$ is the flow map that takes an initial condition $p \in M$ and maps it to the solution at time t of the initial value problem (IVP) having initial condition $p(0) = p$; i.e., for a fixed $p_0 \in M$, the trajectory $\{\Phi(p_0, \cdot)\}_{t \geq 0}$ is a solution of the IVP $\dot{p} = T(p)$, $p(0) = p_0$. The generator of the Koopman semigroup is defined by

$$[Uf] := \lim_{t \rightarrow 0} \frac{U^t f - f}{t}, \tag{6}$$

where the limit is taken in the strong sense.²⁴

The following examples describe the above concepts in certain simple, concrete cases.

Example 1 (Cyclic group). Let $M = \{e, a, a^2\}$ be a cyclic group of order 3 ($a^3 \equiv e$). Define $T : M \rightarrow M$ by $T(p) = a \cdot p$. Hence the entire state space is a periodic orbit of period 3. Let \mathcal{F} be the \mathbb{C} -valued functions on M . Clearly, the space of observables is \mathbb{C}^3 . Let f_1, f_2, f_3 be the indicator functions on e, a, a^2 , respectively,

$$\begin{aligned} f_1(p) &= \begin{cases} 1, & p = e, \\ 0, & p \neq e, \end{cases} \\ f_2(p) &= \begin{cases} 1, & p = a, \\ 0, & p \neq a, \end{cases} \\ f_3(p) &= \begin{cases} 1, & p = a^2, \\ 0, & p \neq a^2. \end{cases} \end{aligned} \tag{7}$$

These form a basis for \mathcal{F} . The action of the Koopman operator on this basis is given as

$$\begin{aligned} [Uf_1](p) &= f_1(a \cdot p) = f_3(p), \\ [Uf_2](p) &= f_2(a \cdot p) = f_1(p), \\ [Uf_3](p) &= f_3(a \cdot p) = f_2(p). \end{aligned} \tag{8}$$

For an arbitrary observable $f \in \mathcal{F}$ given by $f = c_1f_1 + c_2f_2 + c_3f_3$, with $c_i \in \mathbb{C}$, we have that

$$Uf = c_1f_3 + c_2f_1 + c_3f_2.$$

Then, the matrix representation of U in the $\{f_1, f_2, f_3\}$ basis is given by

$$U \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}. \tag{9}$$

In this case, Eq. (9) gives the full action of the Koopman operator on \mathcal{F} . \square

Example 2 (Linear, diagonalizable systems). Let $M = \mathbb{R}^d$ and define $T : M \rightarrow M$ by

$$(T(\mathbf{x}))_i = \mu_i x_i, \tag{10}$$

where $\mathbf{x} = (x_1, \dots, x_d)^T \in M$ and $\mu_i \in \mathbb{R}$. Let \mathcal{F} be the space of \mathbb{C} -valued functions on \mathbb{R}^d . Let $\{\mathbf{b}_1, \dots, \mathbf{b}_d\} \subset M$ be a basis for M and define $f_i(\mathbf{x}) = \langle \mathbf{b}_i, \mathbf{x} \rangle$, where $\langle \cdot, \cdot \rangle$ is the inner product on \mathbb{R}^d . The action of $U : \mathcal{F} \rightarrow \mathcal{F}$ on f_i is

$$\begin{aligned} [Uf_i](\mathbf{x}) &= \langle \mathbf{b}_i, T(\mathbf{x}) \rangle = [b_{i,1}, \dots, b_{i,d}] \begin{bmatrix} \mu_1 x_1 \\ \vdots \\ \mu_d x_d \end{bmatrix} \\ &= [b_{i,1}, \dots, b_{i,d}] \begin{bmatrix} \mu_1 & 0 & \cdots & 0 \\ 0 & \mu_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mu_d \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_d \end{bmatrix}. \end{aligned} \tag{11}$$

Let $\mathcal{F}^d = \otimes^d \mathcal{F}$ and define U_d on \mathcal{F}^d as in Eq. (4). Then for $F = (f_1, \dots, f_d)^T$,

$$[U_d F](\mathbf{x}) = \begin{bmatrix} b_{1,1} & \cdots & b_{1,d} \\ \vdots & \ddots & \vdots \\ b_{d,1} & \cdots & b_{d,d} \end{bmatrix} \begin{bmatrix} \mu_1 & 0 & \cdots & 0 \\ 0 & \mu_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mu_d \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_d \end{bmatrix}. \tag{12}$$

Note that Eq. (12) is just the action of the Koopman operator on the *particular* observable $F = (f_1, \dots, f_d)^T$, not the full action of the Koopman operator on the entire observable space, \mathcal{F} . The main point is that the Koopman operator is reducible, provided T leaves some subspace of \mathcal{F} invariant.

As a special case, we can take the \mathbf{b}_i 's to be the canonical basis vectors having zeros everywhere, except for a 1 in

the i^{th} entry. In this case, the functions $\{f_i\}$ represent the canonical projections onto the coordinates of \mathbf{x} ; i.e., $f_i(\mathbf{x}) = x_i$. Then the action of the Koopman operator with respect to these *particular* observables is given as

$$[U_d F](\mathbf{x}) = \begin{bmatrix} \mu_1 & 0 & \cdots & 0 \\ 0 & \mu_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mu_d \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_d \end{bmatrix}. \tag{13}$$

\square

The next example shows that the Koopman operator formalism can easily handle state spaces that are mixtures of discrete and continuous domains.

Example 3 (Mixed state space). Let $\mathbb{T}^2 = [0, 2\pi) \times [0, 2\pi)$ and $G = \{0, 1, 2\}$. Define $M = \mathbb{T} \times G$ and let the dynamics be given by

$$\begin{aligned} I_{k+1} &= I_k + \left(\frac{S_k}{2}\right)K \sin \theta_k, \quad \text{mod } 2\pi \\ \theta_{k+1} &= \theta_k + I_{k+1}, \quad \text{mod } 2\pi \\ s_{k+1} &= s_k + 1, \quad \text{mod } 3 \end{aligned} \tag{14}$$

where $\mathbf{p}_k = (I_k, \theta_k) \in \mathbb{T}^2$, $s_k \in G$, and $K > 0$. The dynamics are given by the standard map cycling between the unperturbed, shear-flow case and two perturbed cases. The ‘‘perturbation dynamics’’ are driven by a group action.

Let $\mathcal{F}_{\mathbb{T}^2} = \{f : \mathbb{T}^2 \rightarrow \mathbb{C}\}$ be the set of all functions mapping the torus into the complex numbers. We do not assume that the functions in $\mathcal{F}_{\mathbb{T}^2}$ have any type of regularity or algebraic properties; for the moment they are completely arbitrary. Similarly, let $\mathcal{F}_G = \{g : G \rightarrow \mathbb{C}\}$ be the set of all functions from the group G into the complex numbers. One possible choice for the space of observables \mathcal{F} on M is the set $\mathcal{F} = \{h = g \cdot f | f \in \mathcal{F}_{\mathbb{T}^2}, g \in \mathcal{F}_G\}$. Hence observables on M are pointwise products of functions on \mathbb{T}^2 and G and map the mixed state space into \mathbb{C} . The Koopman operator can easily be defined as

$$[Uh](\mathbf{p}_k, s_k) = g(s_{k+1}) \cdot f(\mathbf{p}_{k+1}).$$

Another possible choice for \mathcal{F} could be the set of all the observables that are functions of only I and θ . This is a subset of the previous choice by taking g to be a constant function. This is a natural choice in the case that the ‘‘perturbation dynamics’’ (the dynamics on G), cannot be measured. \square

Example 4 (Partial differential equations). Consider the 2D heat equation on $\mathbb{B} = [-\frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$ with periodic boundary conditions,

$$\frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial t} = c^2 \nabla^2 \mathbf{u}(\mathbf{x}, t). \tag{15}$$

Assuming $\mathbf{u}, \nabla^2 \mathbf{u} \in L^2(\mathbb{B}, d\mathbf{x})$, $\mathbf{u}(\mathbf{x}, t)$ can be expanded in a trigonometric basis,

$$\mathbf{u}(\mathbf{x}, t) = \sum_{\mathbf{j} \in \mathbb{Z}^2} a_{\mathbf{j}}(t) e^{i2\pi \mathbf{j} \cdot \mathbf{x}}, \tag{16}$$

where $\mathbf{j} \cdot \mathbf{x}$ is the dot product of \mathbf{j} and \mathbf{x} . A Galerkin projection onto this basis yields

$$\dot{a}_j(t) = -4\pi^2 c^2 \|j\|_2^2 a_j(t). \tag{17}$$

Thus, we have the continuous-time, infinite-dimensional analogue of Example 2. We could proceed with exhibiting the Koopman semigroup or induce a discrete-time evolution from the continuous-time flow. In the latter case, fix a time step $h > 0$ and get

$$a_j(t_{n+1}) = \exp(-4\pi^2 c^2 \|j\|_2^2 h) a_j(t_n), \tag{18}$$

where $t_n := nh$.

The state space, $M = \ell^2(\mathbb{Z}^2)$, is the space of Fourier coefficients; if $\mathbf{a} \in M$, then $\mathbf{a} = (a_{j_1}, a_{j_2}, \dots)$, where (j_1, j_2, \dots) is some ordering of \mathbb{Z}^2 . Let Eq. (18) define the induced discrete-time evolution map, $T_h : M \rightarrow M$. At this point, we could exactly reduce this problem to Example 2 by restricting our attention to a finite-dimensional subspace of M .

Let the observable space, \mathcal{F} , be the \mathbb{C} -valued functions on M . A family of observables on M , parameterized by $\mathbf{x} \in \mathbb{B}$, is given by (16); namely, fixing $\mathbf{x} \in \mathbb{B}$, we have for any $\mathbf{a} \in M$

$$f_{\mathbf{x}}(\mathbf{a}) = \sum_{j \in \mathbb{Z}^2} a_j e^{i2\pi j \cdot \mathbf{x}}. \tag{19}$$

Hence, the temperature at a point $\mathbf{x} \in \mathbb{B}$ is a linear observable on the space of Fourier coefficients.

Assume the temperature can be measured at a finite number of points $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K)$ in \mathbb{B} and take the finite collection of observables $(f_{\mathbf{x}_1}, \dots, f_{\mathbf{x}_K})$ defined by Eq. (19). Then the action of the Koopman operator on this set of observables is

$$U \begin{bmatrix} f_{\mathbf{x}_1}(\mathbf{a}) \\ \vdots \\ f_{\mathbf{x}_K}(\mathbf{a}) \end{bmatrix} = \sum_{j \in \mathbb{Z}^2} \mu_j a_j \begin{bmatrix} e^{i2\pi j \cdot \mathbf{x}_1} \\ \vdots \\ e^{i2\pi j \cdot \mathbf{x}_K} \end{bmatrix}, \tag{20}$$

where $\mu_j = \exp(-4\pi^2 c^2 \|j\|_2^2 h)$. □

III. KOOPMAN MODE ANALYSIS

A. Eigenfunctions and Koopman modes of U

Thus far, we have avoided putting structure on the function space \mathcal{F} . When \mathcal{F} is a vector space, the Koopman operator is linear. It, therefore, makes sense to study its spectral properties as this will give us insight into the dynamics of the system, similar to the case of linear finite-dimensional systems. We make the further assumptions that \mathcal{F} is a Banach space under some norm, $\|\cdot\|$, and that U is a bounded, and hence continuous, operator on this space.

Let $\{\phi_1, \dots, \phi_n\}$ be a set of eigenfunctions of U , where $n = 1, 2, \dots$, or ∞ , not necessarily forming a complete basis set for \mathcal{F} . In the discrete-time case, we have that

$$[U\phi_i](p) = \lambda_i \phi_i(p). \tag{21}$$

In the continuous-time case, the λ 's are eigenvalues of the generator U of the Koopman semigroup, $\{U^t\}$. The eigencondition is then

$$[U^t \phi_i](p) = e^{\lambda_i t} \phi_i(p), \tag{22}$$

so that $\{e^{\lambda_i t}\}$ are the eigenvalues for the Koopman semigroup.

We first note two simple properties of eigenfunctions: their algebraic structure (Prop. 5) and their role in spectral equivalence of the systems (Prop. 7).

Proposition 5 (Algebraic structure of eigenfunctions under products). Assume \mathcal{F} is a subset of all \mathbb{C} -valued functions on M that forms a vector space which is closed under pointwise products of functions. Then, the set of eigenfunctions forms an Abelian semigroup under pointwise products of functions. In particular, if $\phi_1, \phi_2 \in \mathcal{F}$ are eigenfunctions of U with eigenvalues λ_1 and λ_2 , then $\phi_1 \phi_2$ is an eigenfunction of U with eigenvalue $\lambda_1 \lambda_2$.

Furthermore, if $p \in \mathbb{R}^+$ and ϕ is an eigenfunction with eigenvalue λ , then ϕ^p is a eigenfunction with eigenvalue λ^p , where $\phi^p(x) := (\phi(x))^p$. If ϕ is an eigenfunction that vanishes nowhere and $r \in \mathbb{R}$, then ϕ^r is an eigenfunction with eigenvalue λ^r . The eigenfunctions that vanish nowhere form an Abelian group.

Proof. Assume $U\phi_1 = \lambda_1 \phi_1$ and $U\phi_2 = \lambda_2 \phi_2$ and put $\psi(x) = \phi_1(x)\phi_2(x)$. In discrete time,

$$\begin{aligned} [U\psi](x) &= \psi(T(x)) = \phi_1(T(x))\phi_2(T(x)) \\ &= [U\phi_1](x)[U\phi_2](x) = \lambda_1 \lambda_2 \phi_1(x)\phi_2(x) \\ &= \lambda_1 \lambda_2 \psi(x). \end{aligned}$$

Hence, the set of eigenfunctions is closed under pointwise products. An analogous computation holds for continuous time.

Note that constant functions are eigenfunctions at eigenvalue 1. Hence the constant function that is equal to 1 everywhere is an eigenfunction of U and acts as the identity element. Combining this with the above closure property and standard properties of pointwise products of functions shows that the set of eigenfunctions is an Abelian semigroup.

Let $U\phi = \lambda\phi$ and fix $p \in \mathbb{R}^+$. Then

$$\begin{aligned} [U\phi^p](x) &= \phi^p(Tx) = (\phi(Tx))^p = (\lambda\phi(x))^p \\ &= \lambda^p \phi^p(x). \end{aligned}$$

If ϕ vanishes nowhere, then $\phi^{-1}(x) := 1/\phi(x)$ is well-defined. Then, the above chain of identities remains valid for $r \in \mathbb{R}$ replacing $p \in \mathbb{R}^+$. Hence, the Abelian semigroup of eigenfunctions also contains all of its inverses. Therefore, the set of eigenfunctions vanishing nowhere is an Abelian group. □

Example 6 (Analytic observables of stable/unstable systems). Let $\dot{x} = \lambda x$, with $x, \lambda \in \mathbb{C}$ and $|\lambda| \neq 1$. Then $\Phi^t(x) = e^{\lambda t} x$. Let $\phi(x) = x$. Then

$$\begin{aligned} [U^t \phi](x) &= \phi(\Phi^t(x)) = \phi(e^{\lambda t} x) \\ &= e^{\lambda t} x = e^{\lambda t} \phi(x), \end{aligned}$$

which implies that ϕ is an eigenfunction of U . By Proposition 5, any $\phi_n(x) := (\phi(x))^n = x^n$ is an eigenfunction of U with eigenvalue λ^n .

Let $f(x)$ be an analytic function. Then $f(x) = \sum c_n x^n = \sum c_n \phi_n(x)$, where $c_n = \frac{1}{n!} \frac{d^n f(0)}{dx^n}$. Therefore,

$$[Uf](x) = \sum c_n [U\phi_n](x) = \sum \lambda^n c_n \phi_n(x). \tag{23}$$

□

The second property shows the spectral equivalence of topologically conjugate transformations.

Proposition 7 (Spectral equivalence of topologically conjugate systems). Let $S : M \rightarrow M$ and $T : N \rightarrow N$ be topologically conjugate; i.e., there exists a homeomorphism $h : N \rightarrow M$ such that $S \circ h = h \circ T$. If ϕ is an eigenfunction of U_S with eigenvalue λ , then $\phi \circ h$ is an eigenfunction U_T at eigenvalue λ .

Proof. Fix $x \in M$ and let $y \in N$ be such that $x = h(y)$. The result follows from the chain of equalities:

$$\begin{aligned} \lambda(\phi \circ h)(y) &= \lambda\phi(x) = [U_S\phi](x) = \phi[S(x)] \\ &= \phi\{S[h(y)]\} = \phi\{h[T(y)]\} \\ &= [U_T(\phi \circ h)](y). \end{aligned}$$

□

Example 8 (Topological conjugacy of diagonalizable systems). Let $\mathbf{y}^{(k)} = (y_1^{(k)}, y_2^{(k)})^T$, where the superscript (k) indexes time, and let $\mathbf{y}^{(k+1)} = T\mathbf{y}^{(k)}$, where T is a matrix. Assume that T has eigenvectors $\mathbf{v}_1, \mathbf{v}_2$ at eigenvalues λ_1, λ_2 such that $\mathbf{v}_i \neq \mathbf{e}_j$, where \mathbf{e}_j is the canonical basis vector. If $V = [\mathbf{v}_1, \mathbf{v}_2]$, then after defining new coordinates $\mathbf{x}^{(k)} = (x_1^{(k)}, x_2^{(k)})^T = V^{-1}\mathbf{y}^{(k)}$, we get

$$\begin{bmatrix} x_1^{(k+1)} \\ x_2^{(k+1)} \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} x_1^{(k)} \\ x_2^{(k)} \end{bmatrix} =: \Lambda \begin{bmatrix} x_1^{(k)} \\ x_2^{(k)} \end{bmatrix}.$$

The maps Λ and T are topologically conjugate by $\Lambda V^{-1} = V^{-1}T$.

Note that $\phi_1(\mathbf{x}^{(k)}) = x_1^{(k)}$ and $\phi_2(\mathbf{x}^{(k)}) = x_2^{(k)}$ are eigenfunctions of U_Λ at eigenvalues λ_1 and λ_2 , respectively. By Proposition 5 and Example 6, we have that $\phi_{m,n}(\mathbf{x}^{(k)}) := [\phi_1(\mathbf{x}^{(k)})]^m [\phi_2(\mathbf{x}^{(k)})]^n \equiv [x_1^{(k)}]^m [x_2^{(k)}]^n$ is an eigenfunction of U_Λ at eigenvalue $\lambda_1^m \lambda_2^n$. By Proposition 7, $\phi_{m,n} \circ V^{-1}$ is an eigenfunction of U_T at eigenvalue $\lambda_1^m \lambda_2^n$, where V^{-1} has taken the place of h in Proposition 7. □

Now, assume $f \in \mathcal{F}$ is an observable in the closed, linear span of a set of linearly independent eigenfunctions $\{\phi_i\}_1^n$ (recall n could be finite or infinite). Then

$$f(p) = \sum_{i=1}^n c_i(f)\phi_i(p), \tag{23}$$

for some constants $c_i(f) \in \mathbb{C}$. The dynamics of f are particularly simple,

$$\begin{aligned} [Uf](p) &= f(T(p)) = \sum_{i=1}^n c_i(f)\phi_i(T(p)) \\ &= \sum_{i=1}^n c_i(f)[U\phi_i](p) \\ &= \sum_{i=1}^n \lambda_i c_i(f)\phi_i(p), \end{aligned} \tag{24}$$

and similarly

$$[U^m f](p) = \sum_{i=1}^n \lambda_i^m c_i(f)\phi_i(p). \tag{25}$$

The extension to vector-valued observables $F = (f_1, \dots, f_K)^T$, where each f_i is in the closed linear span of the eigenfunctions, is trivial,

$$\begin{aligned} [U^k F](p) &= \sum_{i=1}^n \lambda_i^m \phi_i(p) \begin{bmatrix} c_i(f_1) \\ \vdots \\ c_i(f_K) \end{bmatrix} \\ &= \sum_{i=1}^n \lambda_i^m \phi_i(p) C_i(F), \end{aligned} \tag{26}$$

where $C_i(F) := [c_i(f_1), \dots, c_i(f_K)]^T$. Motivated by Eq. (26), we have the following definition.

Definition 9. Let ϕ_i be an eigenfunction for the Koopman operator corresponding to the eigenvalue λ_i . Given a vector-valued observable $F : M \rightarrow V$, the Koopman mode, $C_i(F)$, corresponding to ϕ_i is the vector of the coefficients of the projection of F onto $\text{span}\{\phi_i\}$.

Remark 10. The importance of defining Koopman modes with respect to eigenfunctions, rather than eigenvalues, becomes apparent when we consider vector-valued observables and non-simple eigenvalues. For example, let λ have a two-dimensional eigenspace E_λ and let ϕ_1 and ϕ_2 be a basis for it. Let $f_1 = c_1\phi_1 + c_2\phi_2$ and $f_2 = c_3\phi_2$ be scalar-valued. Define $F = (f_1, f_2)^T$. The Koopman modes corresponding to ϕ_1 and ϕ_2 are

$$C_1(F) = \begin{bmatrix} c_1 \\ 0 \end{bmatrix} \quad \text{and} \quad C_2(F) = \begin{bmatrix} c_2 \\ c_3 \end{bmatrix},$$

respectively. Note that both of these Koopman modes have λ as the associated eigenvalue. Therefore, if the Koopman mode was defined with respect to the eigenvalue λ , then it would not be a well-defined object. However, when the eigenspace is one-dimensional, there is no confusion in saying “the Koopman mode corresponding to λ .” □

Remark 11. The definition of Koopman modes can be carried over with a slight modification to generalized eigenfunctions. When an observable can be expanded in terms of only eigenfunctions, the Koopman modes are time-invariant objects. However, when a generalized eigenfunction is present in the expansion, the Koopman modes become time-dependent objects. For example, let ϕ be an eigenfunction and ψ a generalized eigenfunction of U corresponding to $\lambda \neq 0$,

$$U\phi = \lambda\phi \quad \text{and} \quad U\psi = \phi + \lambda\psi.$$

Let $F = C_1(F)\phi + C_2(F)\psi$ be a vector-valued observable. Then

$$\begin{aligned} U^k F &= (C_1(F)\lambda^k + C_2(F)k\lambda^{k-1})\phi + \lambda^k C_2(F)\psi \\ &= \lambda^k \left(C_1(F) + \frac{k}{\lambda} C_2(F) \right) \phi + \lambda^k C_2(F)\psi \end{aligned}$$

for $k \geq 0$. The Koopman mode for ϕ at time k is the time-dependent quantity $C_1(F) + k\lambda^{-1}C_2(F)$.

However, since Koopman modes of generalized eigenfunctions have not been treated in the literature, *whenever we*

refer to Koopman modes in this paper, we implicitly mean a Koopman mode corresponding to an eigenfunction. \square

To be completely explicit, the eigenfunctions are \mathbb{C} -valued observables on the state space M and the eigenpairs (λ_i, ϕ_i) depend only upon the dynamics (M, T) and the function space \mathcal{F} , not on a particular observable. The $C_i(\cdot)$'s can be thought of as a mapping from the observable space into a vector space V ; for example, in Eq. (26) above, C_i maps \mathcal{F} into \mathbb{C}^K . The map $F \mapsto \phi_j C_j(F)$ is then a vector-valued projection operator onto the subspace $\text{span}\{\phi_j\}$.

Remark 12. Given \mathcal{F} as a Banach space of scalar functions, one can ask for conditions on the geometric multiplicity of λ , i.e., dimension of the eigenspace E_λ . The general answer to this question depends on the dynamics T and on the particular space of observables \mathcal{F} chosen. We can give an answer for the case most studied in literature, when T preserves a measure μ with $\mathcal{F} = L^2(M, \mu)$. In this case, all the eigenvalues of the associated Koopman operator U are on the unit circle.

When T is an ergodic transformation (i.e., when any measurable set S invariant under T is either of zero or full measure), all eigenvalues of U are simple (see Petersen, Sec. 2.4 (Ref. 40)). When T is not ergodic, the state space can be partitioned into ergodic sets: minimal invariant sets S such that the restriction $T|_S : S \rightarrow S$ to any S is ergodic. Since all ergodic sets S are disjoint, they support mutually singular functions from \mathcal{F} . As a result, the number of linearly independent eigenfunctions of U at any particular eigenvalue λ is bounded from above by the number of ergodic sets in the state space. The number of such ergodic sets is highly dependent on the character of dynamics. The partition into ergodic sets, the ergodic partition, will be discussed in more detail in Sec. IV A.

Note that the computational method discussed later in Sec. III B 2 assumes fixing an initial condition $p_0 \in M$, which effectively selects the ergodic set $S \subset M$ which contains the point p_0 . Then the Koopman operator $U|_S$ acting on $L^2(S, \mu|_S)$ has simple eigenvalues, where $\mu|_S$ is the ergodic measure on the ergodic component S .

There is an interesting relation between ergodic dynamics and Proposition 5 when the space of observables $\mathcal{F} = L^2(M, \mu)$ is defined with respect to an ergodic measure μ . Given an eigenfunction $\phi \in \mathcal{F}$ with eigenvalue λ , Proposition 5 guarantees that ϕ^n is an eigenfunction with eigenvalue λ^n , as long as $\phi^n \in \mathcal{F}$. If the eigenvalue is periodic ($\lambda^k = \lambda$ for some $k \geq 2$), then ϕ^k lies in the eigenspace E_λ . Ergodicity guarantees the simplicity of E_λ , and hence there exists a non-zero $c \in \mathbb{C}$, such that $\|\phi^k - c\phi\|_2 = 0$, in $L^2(M, \mu)$ norm. \square

Example 13 (Linear systems⁴³). We first look at the case when the dynamics are given by a linear map, $A : M \rightarrow M$, on some finite-dimensional, inner-product space M ; i.e., $\mathbf{x}_{m+1} = A\mathbf{x}_m$. Suppose A has a complete set of eigenvectors, denoted by $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, with corresponding eigenvalues $\{\lambda_1, \dots, \lambda_n\}$. Let $\{\mathbf{w}_j\}_1^n$ be the eigenvectors of the adjoint A^* with eigenvalues $\{\bar{\lambda}_j\}_1^n$, normalized so that $\langle \mathbf{v}_j, \mathbf{w}_k \rangle = \delta_{jk}$. Consider the observable defined as $\phi_j(\mathbf{x}) = \langle \mathbf{x}, \mathbf{w}_j \rangle$. Then

$$\begin{aligned} [U\phi_j](\mathbf{x}) &= \phi_j(A\mathbf{x}) = \langle A\mathbf{x}, \mathbf{w}_j \rangle = \langle \mathbf{x}, A^*\mathbf{w}_j \rangle \\ &= \langle \mathbf{x}, \bar{\lambda}_j \mathbf{w}_j \rangle = \bar{\lambda}_j \langle \mathbf{x}, \mathbf{w}_j \rangle = \bar{\lambda}_j \phi_j(\mathbf{x}). \end{aligned} \tag{27}$$

We see that ϕ_j is an eigenfunction of the Koopman operator. However, the functions $\{\phi_j\}_1^n$ do not exhaust all of the eigenfunctions of U . For example, by Proposition 5, $\phi_j(\mathbf{x})\phi_k(\mathbf{x}) = \langle \mathbf{x}, \mathbf{w}_j \rangle \langle \mathbf{x}, \mathbf{w}_k \rangle$ is an eigenfunction of U . In particular, $g(\mathbf{x}) := \langle \mathbf{x}, \mathbf{w}_j \rangle^k$ is an eigenfunction of U with eigenvalue λ^k for any $k \in \mathbb{N}$.

Let F be the vector-valued observable defined as $F(\mathbf{x}) = \mathbf{x}$ when $\mathbf{x} \in \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_\ell\}$, where $\ell < n$, and zero otherwise; i.e., F acts as the identity on a subspace of M spanned by the first ℓ eigenvectors and has the complement of that subspace as its kernel. Then

$$F(\mathbf{x}) = \sum_{j=1}^{\ell} \langle \mathbf{x}, \mathbf{w}_j \rangle \mathbf{v}_j = \sum_{j=1}^{\ell} \phi_j(\mathbf{x}) \mathbf{v}_j \tag{28}$$

and

$$[U^m F](\mathbf{x}) = \sum_{j=1}^{\ell} \lambda_j^m \phi_j(\mathbf{x}) \mathbf{v}_j. \tag{29}$$

From these expressions, we see that the eigenvector, \mathbf{v}_j , of the linear map A is the Koopman mode, $C_j(F)$, corresponding to ϕ_j . \square

Example 14 (Rotations of the circle). Let the state space be the interval $M = [0, 1)$. Let $\omega \in (0, 1)$ and define $T : M \rightarrow M$ by

$$T(p) = p + \omega, \quad \text{mod } 1. \tag{30}$$

Note that there is a natural identification of M with the circle $\mathbb{T} = \mathbb{R}/2\pi\mathbb{Z}$ and the functions on M with 2π -periodic functions on \mathbb{R} . It is well-known that if $\omega \in \mathbb{Q}$, then every initial condition is periodic and if ω is irrational, then the trajectory starting from any initial condition densely fills M . Note that the dynamics preserve the Lebesgue measure. Let $\mathcal{F} = L^1_{\mathbb{C}}(M)$, be the space of Lebesgue integrable \mathbb{C} -valued functions on M and consider the observable $\phi_n(p) = e^{i2\pi np}$, $n \in \mathbb{Z}$. Then

$$\begin{aligned} [U\phi_n](p) &= \phi_n(T(p)) = e^{i2\pi n(p+\omega)} \\ &= e^{i2\pi n\omega} \phi_n(p). \end{aligned} \tag{31}$$

Therefore, for any $n \in \mathbb{Z}$, ϕ_n is an eigenfunction of U with eigenvalue $\lambda_n = e^{i2\pi n\omega}$. Since the trigonometric polynomials are dense in $L^1(\mathbb{T})$, then for $f_\ell(p) = \sum_{n \in \mathbb{Z}} \hat{f}_\ell(n) e^{i2\pi np} \in L^1(\mathbb{T})$,

$$[Uf_\ell](p) = \sum_{n \in \mathbb{Z}} \hat{f}_\ell(n) e^{i2\pi n\omega} \phi_n(p). \tag{32}$$

If $F = (f_1, \dots, f_K)^T$ is the vector-valued observable, then

$$[UF](p) = \sum_{n \in \mathbb{Z}} e^{i2\pi n\omega} \phi_n(p) \begin{bmatrix} \hat{f}_1(n) \\ \vdots \\ \hat{f}_K(n) \end{bmatrix}. \tag{33}$$

Hence the vectors of Fourier coefficients are the Koopman modes of the system. \square

Example 15 (Partial differential equations). We continue with the heat equation example from above (Example 4).

Recall that a map was defined on the space of Fourier coefficients by

$$(T(\mathbf{a}))_j = \exp(-4\pi^2 c^2 \|\mathbf{j}\|_2^2 h) a_j. \tag{34}$$

Note that the canonical coordinate projections, $\phi_j(\mathbf{a}) := a_j$, are eigenfunctions for the Koopman operator, with eigenvalues $\lambda_j = \exp(-4\pi^2 c^2 \|\mathbf{j}\|_2^2 h)$, by the computation,

$$\begin{aligned} [U\phi_j](\mathbf{a}) &= \phi_j(T(\mathbf{a})) = (T(\mathbf{a}))_j \\ &= \exp(-4\pi^2 c^2 \|\mathbf{j}\|_2^2 h) a_j \\ &= \exp(-4\pi^2 c^2 \|\mathbf{j}\|_2^2 h) \phi_j(\mathbf{a}). \end{aligned} \tag{35}$$

For the observable $f_x(\mathbf{a}) := \sum_{j \in \mathbb{Z}^2} a_j e^{i2\pi j \cdot x}$, we get

$$[U^m f_x](\mathbf{a}) = \sum_{j \in \mathbb{Z}^2} \lambda_j^m \phi_j(\mathbf{a}) e^{i2\pi j \cdot x}.$$

Suppose we can only measure the temperature at a finite number of locations $\{\mathbf{x}_1, \dots, \mathbf{x}_K\}$, then for $F = (f_{\mathbf{x}_1}, \dots, f_{\mathbf{x}_K})^T$,

$$[U^m F](\mathbf{a}) = \sum_{j \in \mathbb{Z}^2} \lambda_j^m \phi_j(\mathbf{a}) \begin{bmatrix} e^{i2\pi j \cdot \mathbf{x}_1} \\ \vdots \\ e^{i2\pi j \cdot \mathbf{x}_K} \end{bmatrix}. \tag{36}$$

In the expression above, each Koopman mode,

$$C_j(F) = \begin{bmatrix} e^{i2\pi j \cdot \mathbf{x}_1} \\ \vdots \\ e^{i2\pi j \cdot \mathbf{x}_K} \end{bmatrix}, \tag{37}$$

is just a “shape” function on the physical space $\mathbb{B} = [-1/2, 1/2] \times [-1/2, 1/2]$. This stresses the point that the eigenfunctions are defined on the state space M while the Koopman modes are functions in the output space. \square

The development thus far has only focused on the case when an observable is in the closed linear span of some set of eigenfunctions of the Koopman operator. No assumption was made on whether this set was a complete set for U or even if U possessed a complete set of eigenfunctions. One could ask what conditions we could impose on the system, which are sufficient to guarantee that U has a spectral decomposition. This is the case for measure-preserving dynamical systems, as we now explain.

Let $\mathcal{A} \subset M$ be the attractor of the dynamical system and μ the unique invariant measure supported on \mathcal{A} . Often, μ will be a so-called physical measure. These types of measures exhibit the important property,

$$\frac{1}{n} \sum_{k=0}^{n-1} [U^k f](p) \rightarrow \int_M f d\mu, \tag{38}$$

for any continuous observable $f : M \rightarrow \mathbb{C}$ and for Lebesgue-almost every $p \in M$ belonging to a positive Lebesgue measure set $V \subset M$ containing the attractor (see Young⁶¹ for a more detailed discussion). Such measures are important in applications since they guarantee the existence of well-defined time-averages

even when an experiment starts with initial conditions not on the attractor. In such cases, we can restrict our attention to the dynamics and observables on the attractor and recover all the asymptotic behavior of the system.

The situation on the attractor is quite nice when we consider the function space $\mathcal{F} = L^2(\mathcal{A}, \mu)$. The restriction of the dynamics to the attractor, $T|_{\mathcal{A}} : \mathcal{A} \rightarrow \mathcal{A}$, can be shown to be invertible μ -almost everywhere. The restriction of the Koopman operator to the attractor, $U|_{\mathcal{A}} : L^2(\mathcal{A}, \mu) \rightarrow L^2(\mathcal{A}, \mu)$, can then be defined by $U|_{\mathcal{A}} f = f \circ T|_{\mathcal{A}}$. In this case, the operator is unitary,^{22,40} implying that all of the eigenvalues lie on the unit circle and the eigenfunctions are orthogonal.^{34,36,40}

Since $U|_{\mathcal{A}}$ is unitary, there exists a spectral resolution⁴⁰

$$U|_{\mathcal{A}} f = \int_{S^1} \lambda dE(\lambda) f \tag{39}$$

where E is a projection-valued Borel measure on the unit circle; i.e., for any Borel set S in the unit circle, $E(S)$ is a projection operator. The measure E is supported on the spectrum of $U|_{\mathcal{A}}$. E can be decomposed into two measures, E_p and E_c , that are supported on the point spectrum and the continuous part of the spectrum, respectively. For any $f \in L^2(\mathcal{A}, \mu)$, the spectral resolution becomes

$$U|_{\mathcal{A}}^k f = \sum_j e^{i2\pi\omega_j k} P_j f + \int_0^1 e^{i2\pi\theta k} dE_c(\theta) f \tag{40}$$

where $P_j : L^2(\mathcal{A}, \mu) \rightarrow L^2(\mathcal{A}, \mu)$ is the orthogonal projection onto the eigenspace corresponding to the eigenvalue $\lambda_j = e^{i2\pi\omega_j}$ and E_c is the projection-valued measure corresponding to the continuous part of the spectrum. Either term on the right-hand side of Eq. (40) could be zero depending on whether the operator has no point spectrum or no continuous part of the spectrum. When the eigenvalues are simple, we get for a vector-valued observable $F = (f_1, \dots, f_K)^T \in \bigoplus_{i=1}^K L^2(\mathcal{A}, \mu)$,

$$[U|_{\mathcal{A}}^k F](p) = \sum_j e^{i2\pi\omega_j k} \phi_j(p) C_j(F) + \int_0^1 e^{i2\pi\theta k} dE_c(\theta) F(p). \tag{41}$$

where ϕ_j is the eigenfunction corresponding to the eigenvalue $\lambda_j = e^{i2\pi\omega_j k}$ and we have used $P_j F(p) = \phi_j(p) C_j(F)$ (when λ_i is non-simple, this identity does not necessarily hold). Finally, we note that the constant functions on the attractor are eigenfunctions of $U|_{\mathcal{A}}$ at eigenvalue 1 and that $P_0 f := \int_{\mathcal{A}} f d\mu$ defines a projection onto the constant functions. Then, Eq. (41) becomes

$$\begin{aligned} [U|_{\mathcal{A}}^k F](p) &= \int_{\mathcal{A}} F(p) d\mu(p) + \sum_{\{j: \omega_j \neq 0\}} e^{i2\pi\omega_j k} \phi_j(p) C_j(F) \\ &+ \int_0^1 e^{i2\pi\theta k} dE_c(\theta) F(p) \end{aligned} \tag{42}$$

(for further details on this decomposition consult Mezić³⁴).

Therefore, for the case of measure-preserving transformations or those systems possessing a physical measure, the asymptotic dynamics of the Koopman operator are given by

the contributions of three components: (1) the average value of the observable, (2) the portion admitting a Koopman mode expansion (the part corresponding to the point spectrum), and (3) the contribution of the continuous part of the spectrum. The third component, $dE_c(\theta)F(p)$, we refer to as the Koopman mode distribution, or the KM distribution for short. On the other hand, the Koopman mode expansion (the second component in the decomposition) is fairly well-understood with respect to its relation to the physics of a problem, the contribution of the KM distribution does not enjoy the same level of understanding and has received little attention, so far, in the literature.

B. Computation of Koopman modes

In the examples that were presented thus far, it was fairly easy to determine the eigenpairs of the Koopman operator and the corresponding Koopman modes. For a general system, however, things will not be so easy. This section will discuss a few methods to compute the projection of an observable onto the eigenspaces of the Koopman operator, from both the theoretical and numerical viewpoints.

1. Theoretical results

The first tool is given by the following theorem. It is a special case of a result found in Yosida.⁶⁰

Theorem 16. Let \mathcal{F} be a Banach space and $U : \mathcal{F} \rightarrow \mathcal{F}$. Assume $\|U\| \leq 1$. Let λ be an eigenvalue of U such that $|\lambda| = 1$. Let $\hat{U} = \lambda^{-1}U$ and define

$$A_K(\hat{U}) = \frac{1}{K} \sum_{k=0}^{K-1} \hat{U}^k.$$

Then A_K converges in the strong operator topology to the projection operator on the subspace of \hat{U} -invariant function; i.e., onto the eigenspace E_λ corresponding to λ . That is, for all, $f \in \mathcal{F}$

$$\lim_{K \rightarrow \infty} A_K f = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} \hat{U}^k f = P_\lambda f. \tag{43}$$

where $P_\lambda : \mathcal{F} \rightarrow E_\lambda$ is a projection operator.

Proof. See Yosida⁶⁰ or Krengel.²³

Consider the case when the eigenvalues are simple and $|\lambda_1| = \dots = |\lambda_\ell| = 1$ and $|\lambda_n| < 1$ for $n > \ell$. Then, $\lambda_j = e^{i2\pi\omega_j}$ for some real ω_j , when $j \leq \ell$. For vector-valued observables, the projections defined by Eq. (43) take the form,

$$\phi_j C_j(F) = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} e^{-i2\pi\omega_j k} [U^k F], \tag{44}$$

for $j = 1, \dots, \ell$. Hence, Theorem 16 reduces to Fourier analysis, as one might expect, for those eigenvalues on the unit circle and the projections can be computed with any implementation of a fast Fourier transform. For further discussion, consult Mezić and Banaszuk.³⁶ or Mezić.³⁴

When an observable is a linear combination of a finite collection of eigenfunctions corresponding to simple eigenvalues,

we get an extension of the above theorem to eigenvalues not having unit modulus.

Theorem 17 (Generalized Laplace analysis). Let $\{\lambda_1, \dots, \lambda_m\}$ be a (finite) set of simple eigenvalues for U ordered so that $|\lambda_1| \geq \dots \geq |\lambda_m|$ and let ϕ_i be an eigenfunction corresponding to λ_i . For each $n \in \{1, \dots, N\}$, assume $f_n : M \rightarrow \mathbb{C}$ and $f_n \in \text{span}\{\phi_1, \dots, \phi_m\}$. Define the vector-valued observable $F = (f_1, \dots, f_N)^T$.

Then the Koopman modes for F can be computed via

$$\phi_j C_j(F) = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} \lambda_j^{-k} \left[U^k F - \sum_{i=1}^{j-1} \lambda_i^k \phi_i C_i(F) \right]. \tag{45}$$

Proof. Since each ϕ_i is an eigenfunction, $\text{span}\{\phi_1, \dots, \phi_m\}$ is a U -invariant subspace, so the restriction of U to this subspace is a finite-dimensional linear operator and can be represented with a matrix. Any $f_n \in \text{span}\{\phi_1, \dots, \phi_m\}$ can be written as $f_n = \sum_{j=1}^m c_j(f_n) \phi_j$. Then $f_n - \sum_{i=1}^{j-1} c_i(f_n) \phi_i \in \text{span}\{\phi_j, \dots, \phi_m\}$. U restricted to $\text{span}\{\phi_j, \dots, \phi_m\}$ has eigenvalues $\{\lambda_j, \dots, \lambda_m\}$. Then $\lambda_j^{-1}U$ restricted to $\text{span}\{\phi_j, \dots, \phi_m\}$ has eigenvalues $\{1, \frac{\lambda_{j+1}}{\lambda_j}, \dots, \frac{\lambda_m}{\lambda_j}\}$. The modulus of any element of this set is ≤ 1 . Taking the average, as in Eq. (44), using the operator $\lambda_j^{-1}U$ restricted to $\text{span}\{\phi_j, \dots, \phi_m\}$ gives the projection onto $\lambda_j^{-1}U$ -invariant functions. The $\lambda_j^{-1}U$ -invariant functions are just elements of $\text{span}\{\phi_j\}$. Then

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} [\lambda_j^{-1}U]^k \left(f_n - \sum_{i=1}^{j-1} c_i(f_n) \phi_i \right)$$

is just $c_j(f_n) \phi_j$. This is equivalent to (45) when $F = f_n$. The extension a vector-valued F is obvious. \square

Remark 18. Theorem 17 is a simple consequence of Theorem 16. The case of F having elements in a generalized eigenspace is more difficult and is treated in forthcoming work by the authors. \square

Remark 19. Analogous expressions hold for continuous time, with Eqs. (44) and (45) replaced by

$$\phi_j C_j(F) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T e^{-i2\pi\omega_j t} [U^t F] dt. \tag{46}$$

and

$$\phi_j C_j(F) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T e^{-\lambda_j t} \left[U^t F - \sum_{i=1}^{j-1} e^{\lambda_i t} \phi_i C_i(F) \right] dt, \tag{47}$$

respectively. \square

The first thing to note is that for Theorem 17 and its continuous-time analogue a set of eigenvalues is needed; they are not computed as part of the theorem. To get the projections $\phi_j C_j(F)$, the more unstable modes must be subtracted off of the dynamics before the time-average is computed.

The most common case for which Theorems 16 and 17 (and their continuous-time analogues) are applied occurs when we restrict our attention to a compact invariant subset

\mathcal{C} of the basin of attraction for some attractor $\mathcal{A} \subset M$ and take for \mathcal{F} the product $L^2(\mathcal{A}, \mu) \times \mathcal{H}(\mathcal{C})$, where μ is the unique invariant measure supported on the attractor and $\mathcal{H}(\mathcal{C})$ is the space of analytic functions on \mathcal{C} . In this case, an eigenvalue for U satisfies $|\lambda| \leq 1$.

Example 20 (Harmonic oscillator). Consider the Harmonic oscillator

$$\begin{aligned} \dot{p}_1(t) &= p_2(t) \\ \dot{p}_2(t) &= -\omega^2 p_1(t) \end{aligned}$$

Letting $\mathbf{p}(t) = (p_1(t), p_2(t))^T$, the solution flow is

$$\mathbf{p}(t) = \Phi^t(\mathbf{p}(0)) = \begin{bmatrix} \cos \omega t & \frac{1}{\omega} \sin \omega t \\ -\omega \sin \omega t & \cos \omega t \end{bmatrix} \begin{bmatrix} p_1(0) \\ p_2(0) \end{bmatrix}$$

Note that this system is divergence-free and therefore preserves volume in the state space; all eigenvalues of U have modulus 1.

Let $F(\mathbf{p}(t)) = \mathbf{p}(t)$. Equation (46) is nonzero only for $\omega_1 = \omega$ and $\omega_2 = -\omega$. Then for $\lambda_1 = e^{i\omega}$ and $\lambda_2 = e^{-i\omega}$, we get

$$\begin{aligned} \phi_1(\mathbf{p}(0))C_1(F) &= \frac{1}{2} \left[\begin{pmatrix} p_1(0) \\ p_2(0) \end{pmatrix} - i \begin{pmatrix} p_2(0)/\omega \\ -\omega p_1(0) \end{pmatrix} \right] \\ &= \frac{1}{2} \left(p_1(0) - i \frac{p_2(0)}{\omega} \right) \begin{bmatrix} 1 \\ i\omega \end{bmatrix} \end{aligned}$$

and

$$\begin{aligned} \phi_2(\mathbf{p}(0))C_2(F) &= \frac{1}{2} \left[\begin{pmatrix} p_1(0) \\ p_2(0) \end{pmatrix} + i \begin{pmatrix} p_2(0)/\omega \\ -\omega p_1(0) \end{pmatrix} \right] \\ &= \frac{1}{2} \left(p_1(0) + i \frac{p_2(0)}{\omega} \right) \begin{bmatrix} 1 \\ -i\omega \end{bmatrix}, \end{aligned}$$

respectively. We can write

$$\begin{aligned} [U^t F](\mathbf{p}(0)) &= F(\mathbf{p}(t)) = \Phi^t(\mathbf{p}(0)) \\ &= e^{i\omega t} \phi_1(\mathbf{p}(0))C_1(F) \\ &\quad + e^{-i\omega t} \phi_2(\mathbf{p}(0))C_2(F), \end{aligned}$$

or more explicitly,

$$\begin{aligned} \begin{bmatrix} p_1(t) \\ p_2(t) \end{bmatrix} &= e^{i\omega t} \frac{1}{2} \left(p_1(0) - i \frac{p_2(0)}{\omega} \right) \begin{bmatrix} 1 \\ i\omega \end{bmatrix} \\ &\quad + e^{-i\omega t} \frac{1}{2} \left(p_1(0) + i \frac{p_2(0)}{\omega} \right) \begin{bmatrix} 1 \\ -i\omega \end{bmatrix}. \end{aligned}$$

We recognize the familiar normal mode expansion for the harmonic oscillator. The normal modes are the Koopman modes $C_1(F) = [1, i\omega]^T$ and $C_2(F) = [1, -i\omega]^T$, whereas the Koopman eigenfunctions $\phi_{1,2}(\mathbf{p}(0)) = \frac{1}{2}(p_1(0) \mp p_2(0)/\omega)$ are the terms of the normal mode expansion that are functions of only the initial conditions. \square

While in principle the projections onto the stable and unstable modes can be computed directly using Theorem 17, difficulties arise when we move past simple cases. If an explicit representation of U is not known for the observable

F , we would need to compute the projections numerically. For $|\lambda| \neq 1$, this would require a numerical implementation of a Laplace transform; these are generally unstable computations, precluding the direct numerical implementation of Theorem 17 when F has stable or unstable Koopman modes. Therefore, Theorem 17 is more suited as an analytical tool, rather than a numerical one.

2. A numerical algorithm: The dynamic mode decomposition (DMD)

Usually, we do not have access to an explicit representation of the Koopman operator. The behavior of the operator can only be ascertained by its action on an observable and usually at only a finite number of initial conditions. Thus, we are led to consider data-driven algorithms for computing the Koopman modes. By data, we mean a sequence of observations of a vector-valued observable along a trajectory $\{T^k p\}$. The following algorithms use these sequences of observations to approximate both the eigenvalues and the Koopman modes of U without having to numerically implement a Laplace transform. This is accomplished by finding the best approximation of U on some finite-dimensional subspace and computing eigenfunctions of the resulting finite-dimensional linear operator. The notion of *best approximation* will be made clear in what follows.

Fix a vector-valued observable $F : M \rightarrow \mathbb{C}^m$ and consider the cyclic subspace

$$\mathcal{K}_\infty = \text{span}\{U^k F\}_{k=0}^\infty; \tag{48}$$

that is, \mathcal{K}_∞ is the space of vector-valued observables in which finite linear combinations of elements from $\{U^k F\}_{k=0}^\infty$ are dense.

Fix an $r < \infty$ and consider the Krylov subspace

$$\mathcal{K}_r = \text{span}\{U^k F\}_{k=0}^{r-1}. \tag{49}$$

We will assume that $\{U^k F\}_{k=0}^{r-1}$ is a linearly independent set so that these functions form a basis for \mathcal{K}_r . Note that $U\mathcal{K}_r \subset \mathcal{K}_{r+1}$, so that, in general, \mathcal{K}_r is not U -invariant; it is only invariant if $U^r F \in \text{span}\{U^k F\}_{k=0}^{r-1}$.

Let $P_r : \mathcal{F}^m \rightarrow \mathcal{K}_r$ be a projection from the space of vector-valued observables onto \mathcal{K}_r . Then

$$P_r U|_{\mathcal{K}_r} : \mathcal{K}_r \rightarrow \mathcal{K}_r \tag{50}$$

is a finite-dimensional linear operator. This operator has a matrix representation, $\mathbf{A}_r : \mathbb{C}^r \rightarrow \mathbb{C}^r$, in the $\{U^k F\}_{k=0}^{r-1}$ -basis. Note that this matrix is dependent upon (1) the vector-valued observable, (2) the number of time-steps r used (the dimension of the Krylov subspace), and (3) the projection P_r used which is specified by the type of approximation we choose to make; in the following algorithms, this projection is the least-square approximation for evolution from a single point $p \in M$.

If (λ, \mathbf{v}) is an eigenpair for \mathbf{A}_r , where $\mathbf{v} = (v_0, \dots, v_{r-1})^T \in \mathbb{C}^r$, then $\phi = \sum_{j=0}^{r-1} v_j [U^j F]$ is an eigenfunction of $P_r U|_{\mathcal{K}_r}$. By restricting our attention to a fixed observable F and a Krylov subspace, we have reduced the problem of finding

eigenvalues and Koopman modes of the Koopman operator to finding eigenvalues and eigenvectors for a matrix A_r .

A standard method for computing eigenvalues of a matrix is the Arnoldi algorithm and its variants. These are iterative methods relying on Krylov subspaces. The basic idea behind these algorithms is to project the matrix onto a lower-dimensional subspace and to compute the eigenvalues of the lower-rank matrix. If the projection has a nice representation, then the eigenvalue problem for this lower-rank approximation can be efficiently solved.

The standard Arnoldi algorithm^{2,32,46} assumes we have a matrix $A : \mathbb{C}^m \rightarrow \mathbb{C}^m$ whose eigenvalues and eigenvectors we want to compute. Starting from a random vector $\mathbf{b} \in \mathbb{C}^m$ of unit norm, we form the Krylov subspace

$$\mathcal{K}_r := \text{span}\{\mathbf{b}, A\mathbf{b}, \dots, A^{r-1}\mathbf{b}\}.$$

Assuming full rank, an orthonormal basis $\{\mathbf{q}_j\}_1^r$ for \mathcal{K}_r can be found using a Gram-Schmidt procedure applied to $\{A^j\mathbf{b}\}_{j=0}^{r-1}$. Orthogonalization and renormalization is usually performed at each step j . Letting Q_r be the matrix formed from the orthonormal basis, we get the relation

$$H_r = Q_r^* A Q_r,$$

where H_r is of upper Hessenberg form and has the interpretation of the orthogonal projection of A onto \mathcal{K}_r . H_r can be diagonalized efficiently. The eigenvalues of H_r approximate the r eigenvalues of A of largest magnitude. Implementations of this algorithm use various additional methods to ensure numerical stability.

By applying the Arnoldi algorithm, we implicitly assume there exists a matrix A whose evolution $A^k\mathbf{b} \in \mathbb{C}^m$ matches the evolution $[U^k F](p) \in \mathbb{C}^m$ for $k = 0, \dots, r$. Unfortunately, since we do not have an explicit representation of the Koopman operator, we cannot use the standard Arnoldi algorithm. This is due to the orthogonalization and renormalization performed at each step, which is equivalent to changing the observable F at each step. If \mathbf{q}_k is the vector formed from normalizing the component of $[U^k F]$ orthogonal to $\text{span}\{[U^j F]\}_0^{k-1}$, then there is some $G : M \rightarrow \mathbb{C}^m$ such that $\mathbf{q}_k = G(p)$. Hence we are never looking at the action of the Koopman operator along a single trajectory and observable, precluding the use of the Arnoldi algorithm with data obtained from simulation.

A variant of the Arnoldi algorithm, utilizing companion matrices, was first described in Ruhe.⁴⁴ The algorithm was popularized in the Fluids community by Rowley *et al.*⁴³ and Schmid.⁴⁷ In Schmid,⁴⁷ the algorithm was dubbed the DMD, whereas Rowley *et al.*⁴³ related the algorithm to the approximation of Koopman modes.

The strength of the DMD algorithm is that it only requires a sequence of vectors $\{\mathbf{b}_k\}_{k=0}^r$, where

$$\mathbf{b}_k := U^k F(p) \in \mathbb{C}^m \tag{51}$$

for some fixed $F : M \rightarrow \mathbb{C}^m$ and fixed $p \in M$. The algorithm gives the best approximation, at the point $p \in M$, of the projection $\phi C(F)$ onto the eigenfunction ϕ (Eq. (26)). As will

be seen, this corresponds to a specific choice of the projection operator P_r appearing in Eq. (50).

To derive the DMD algorithm, let

$$K_r := [\mathbf{b}_0, \dots, \mathbf{b}_{r-1}].$$

The columns of K_r are the \mathbb{C}^m -vectors resulting from the point evaluations of the $\{U^k F\}$ -basis for the Krylov subspace \mathcal{K}_r at the point $p \in M$.

In general, \mathbf{b}_r will not be in the span of the columns of K_r . In this case, $\mathbf{b}_r = \sum_{j=0}^{r-1} c_j \mathbf{b}_j + \boldsymbol{\eta}_r$, where the c_j 's are chosen to minimize the \mathbb{C}^m -norm of the residual $\boldsymbol{\eta}_r$. This corresponds to choosing the projection operator P_r appearing in Eq. (50) so that $P_r U^r F$ is the least-squares approximation to $U^r F$ at the point $p \in M$ as measured by the \mathbb{C}^m -norm; i.e.,

$$\begin{aligned} \|[U^r F](p) - P_r[U^r F](p)\|_{\mathbb{C}^m} &= \left\| \mathbf{b}_r - \sum_{j=0}^{r-1} c_j \mathbf{b}_j \right\|_{\mathbb{C}^m} \\ &\leq \left\| \mathbf{b}_r - \sum_{j=0}^{r-1} d_j \mathbf{b}_j \right\|_{\mathbb{C}^m} \end{aligned}$$

for any other $\{d_0, \dots, d_{r-1}\}$.

Since $\mathbf{b}_r = K_r \mathbf{c} + \boldsymbol{\eta}_r$, where $\mathbf{c} = (c_0, \dots, c_{r-1})^T$, we get

$$U K_r = [\mathbf{b}_1, \dots, \mathbf{b}_r] = [\mathbf{b}_1, \dots, \mathbf{b}_{r-1}, K_r \mathbf{c} + \boldsymbol{\eta}_r],$$

or equivalently

$$U K_r = K_r A_r + \boldsymbol{\eta}_r \mathbf{e}^T \tag{52}$$

where $\mathbf{e} = (0, \dots, 0, 1)^T \in \mathbb{C}^m$ and

$$A_r = \begin{bmatrix} 0 & 0 & \cdots & 0 & c_0 \\ 1 & 0 & \cdots & 0 & c_1 \\ 0 & 1 & \cdots & 0 & c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & c_{r-1} \end{bmatrix} \tag{53}$$

is the $r \times r$ companion matrix; it is the matrix representation of $P_r U$ in the $\{U^k F\}_{k=0}^{r-1}$ -basis.

Diagonalize the companion matrix,

$$A_r = V^{-1} \Lambda V, \tag{54}$$

where Λ is the diagonal matrix of eigenvalues λ_i and the columns of V^{-1} are eigenvectors of A_r . Inserting the expression for A_r into Eq. (52) and multiplying on the right by V^{-1} gives

$$U K_r V^{-1} = K_r V^{-1} \Lambda + \boldsymbol{\eta}_r \mathbf{e}^T V^{-1}. \tag{55}$$

Define

$$E := K_r V^{-1}. \tag{56}$$

Then Eq. (55) becomes

$$U E = E \Lambda + \boldsymbol{\eta}_r \mathbf{e}^T V^{-1}. \tag{57}$$

For large enough m , it is hoped that $\|\eta_r e^T V^{-1}\|$ is small. If that is the case, then $UE \approx E\Lambda$ and the columns of E approximate some eigenvectors of U and the diagonal elements of Λ approximate some eigenvalues of U . Note that $\|\eta_r e^T V^{-1}\| = 0$ whenever $r > m$, since then the columns of K_r are linearly dependent which implies that $\eta_r = 0$.

Definition 21. Let Λ and E be defined as in Eqs. (55) and (57), respectively. Let w_i be the i^{th} column of E and λ_i be the i^{th} diagonal element of Λ . Then w_i is called an empirical Ritz vector and λ_i is called an empirical Ritz value.

Each empirical Ritz vector approximates $\phi_i(p)C_i(F)$, the projection of F onto some eigenvector ϕ_i , and the empirical Ritz values approximate the corresponding eigenvalues of U . For this reason, we will generally refer the empirical Ritz vectors and values as the Koopman modes and eigenvalues computed by the DMD algorithm although this is not strictly true and loosens the terminology.

Remark 22. The above algorithm is very much tied to the initial condition chosen. This dependence arises since the empirical Ritz values and vectors are formed using a Krylov subspace that is generated by a sequence of vector-valued observations along a finite trajectory having initial condition $p \in M$. A given initial condition may not reveal the full spectrum and different initial conditions can reveal different parts of the spectrum. For example consider the dynamical system,

$$p_{k+1} = \begin{cases} \lambda_1 p_k, & p_k \leq 0 \\ \lambda_2 p_k, & p_k > 0 \end{cases}, \quad (58)$$

where $0 < \lambda_1 < 1 < \lambda_2$. Let $\phi_1(p) = \min\{p, 0\}$ and $\phi_2(p) = \max\{0, p\}$. These are eigenfunctions of the Koopman operator at eigenvalues λ_1 and λ_2 , respectively. Let $F(p) = p \equiv \phi_1(p) + \phi_2(p)$. Analytically, we can decompose the observable into a sum of projections onto eigenspaces: $F(p) = P_1 F(p) + P_2 F(p)$, where P_i is the projection onto ϕ_i . Choosing an initial condition $p < 0$ and applying the DMD algorithm only computes the projection onto the stable mode; the DMD algorithm only reveals $P_1 F(p)$. Similarly, choosing $p > 0$ only reveals the unstable mode, $P_2 F(p)$. Therefore, the DMD algorithm may only reveal a subset of the spectrum of the Koopman operator and the corresponding Koopman modes.

It should also be remarked that if $F \notin \text{span}\{\phi_i\}$ for some eigenfunction ϕ_i , then the DMD algorithm will not reveal that mode. This is often the case for natural choices for a set of observables, as was the case in the linear system example above (see Example 13, p. 21). \square

The version of the DMD algorithm described tends to be numerically ill-conditioned. This is due to $A^k b_0$ converging to the eigenspaces corresponding to the largest magnitude eigenvalues, resulting in the columns of K_r becoming nearly linearly dependent. A robust version of the algorithm has been described in Schmid.⁴⁷ It amounts to first computing a singular value decomposition (SVD) of K_r and projecting A onto the Krylov subspace using the SVD basis (for details, consult Schmid.⁴⁷). Chen *et al.*⁸ discuss variants of the dynamic mode decomposition, relates it to discrete Fourier

transforms, and introduce an ‘‘optimized’’ DMD algorithm that computes an arbitrary number of modes from data.

C. Applications of Koopman modes

The theory of Koopman modes has led to a number applications in the literature. Broadly, the uses of Koopman modes can be classified under two headings: model reduction and coherency. Model reduction deals with extracting the spatial features of just a few Koopman modes and attempting to understand the physics of the system just based on those, neglecting the details contained in other Koopman modes. The notion of coherency, on the other hand, deals with how observables relate dynamically with respect to a Koopman mode. Coherency is always defined for a (not necessarily proper) subset of observables and an eigenvalue λ . The subset of observables is coherent for λ if the dynamics are identical. This reduces to checking if the initial magnitudes and phases of $C_\lambda(F)$ are the same for each observable in the subset. The following definition, first appearing in Susuki and Mezić,⁵³ makes this precise.

Definition 23 (Coherency between Koopman modes). Consider a vector-valued observable $F : M \rightarrow \mathbb{C}^m$, where $F = (f_1, \dots, f_m)^T$ and $f_j : M \rightarrow \mathbb{C}$ for $j = 1, \dots, m$. Let $\{C_1(F), \dots, C_\ell(F)\}$, $\ell < \infty$ be a collection of Koopman modes of interest. Note that $C_i(F) = (c_{i,1}, \dots, c_{i,m})^T \in \mathbb{C}^m$. Fix $\epsilon_1, \epsilon_2 > 0$ and consider $j_1, j_2 \in \{1, \dots, m\}$. Then, f_{j_1} and f_{j_2} are (ϵ_1, ϵ_2) -coherent (with respect to the chosen Koopman modes) if

- (i) $\left| |c_{i,j_1}| - |c_{i,j_2}| \right| < \epsilon_1$ and
- (ii) $|\angle c_{i,j_1} - \angle c_{i,j_2}| < \epsilon_2$,

for all $i = 1, \dots, \ell$. \square

The choice of two epsilon values in the above definition allows the practitioner to set the tolerances of the magnitude and phase independently. This is useful when one can tolerate more variation in either the modulus or the phase and still call the modes coherent. Thus c_{i,j_2} is coherent with c_{i,j_1} , if the complex number c_{i,j_2} is contained inside some rectangle centered at c_{i,j_1} .

The examples we present are necessarily a subset of those that exist in the literature.

1. Power systems

Koopman mode analysis has seen application in the analysis of power systems. In Susuki and Mezić,^{53,54} the authors used Koopman mode analysis to identify coherency in the short-term swing dynamics of multi-machine power systems, with the New England Test System and IEEE Reliability Test System-1996 being used as test cases for the methodology.^{53,54} In Susuki and Mezić,⁵⁶ the authors used Koopman mode analysis to identify precursors to the so-called coherent swing instability of power systems where a group of generators synchronously loses coherency with the rest of the system after a local disturbance. We will focus only upon the identification of coherency, as pursued in Susuki and Mezić,^{53,54} since it underlies the work in Susuki and Mezić⁵⁶ as well.

The New England system is a 39-bus system having 10 synchronous generators, while the IEEE systems has 73 buses and 99 synchronous generators.^{53,54} As the analysis of

the two systems is the same, we focus on the simpler New England system.

The swing dynamics of the New England system were given by the following systems of differential equations:^{53,54}

$$\begin{aligned} \frac{d\delta_i}{dt} &= \omega_i \\ \frac{H_i}{\pi f_b} \frac{d\omega_i}{dt} &= -D_i \omega_i + P_{mi} - G_{ii} E_i^2 \\ &\quad - \sum_{\substack{j=1 \\ j \neq i}}^{10} E_i E_j C_{ij}, \end{aligned} \quad (59)$$

where C_{ij} is the coupling term

$$C_{ij} = G_{ij} \cos(\delta_i - \delta_j) + B_{ij} \sin(\delta_i - \delta_j).$$

In this model, $i = 2, \dots, 10$ indexed the generators, δ_i was the angular position of the rotor of generator i relative to bus 1, and ω_i was the rotor speed of generator i relative to that of bus 1; D_i was the damping coefficient of generator i , E_i was the voltage of the generator, and P_{mi} the mechanical input power; G_{ii} was the internal conductance of generator i , while $G_{ij} + \sqrt{-1}B_{ij}$ was the transfer impedance between generators i and j ; H_i was a per unit time inertia constant and f_b a frequency.^{53,54} The variables H_i , E_i , D_i , f_b , and power loads were specified during simulations.^{53,54} The variables G_{ii} , G_{ij} , and B_{ij} were computed using power flow computations (see Susuki and Mezić,^{53,54} and the references therein, for full numerical details). The state space for each generator ($i = 2, \dots, 10$) was the cylinder $\mathcal{C} = [-\pi, \pi] \times \mathbb{R}$ and the full state space M for this system was $M = \mathcal{C} \times \dots \times \mathcal{C} = \mathcal{C}^9$. Each generator exhibited a stable equilibrium at $(\delta_i^*, \omega_i^* = 0)$, for some δ_i^* , computed using a power flow computation.

Let an observable $f_i : M \rightarrow \mathbb{R}$ be given by $f_i(\boldsymbol{\delta}, \boldsymbol{\omega}) = \omega_i$, where $\boldsymbol{\delta} = (\delta_2, \dots, \delta_{10})$ and $\boldsymbol{\omega} = (\omega_2, \dots, \omega_{10})$. The vector-valued observable chosen for Koopman mode analysis was $F = (f_2, \dots, f_{10})^T$, so that

$$F(\boldsymbol{\delta}, \boldsymbol{\omega}) = \begin{bmatrix} \omega_2 \\ \vdots \\ \omega_{10} \end{bmatrix}.$$

This was a physically relevant observable since in practice one measures the rotor speeds for each generation plant.^{53,54}

The system was evolved for a short time period with a disturbance from the equilibrium state localized at rotor 8. Study of the resulting trajectories showed that generators 2, 3, 6, and 7 were a coherent group;^{53,54} the four generators exhibited responses in angular frequencies ω_i having the same amplitude and phase.

The DMD algorithm was applied to the same simulation data. The Koopman modes of interest were those that had the largest norms and corresponding growth rates $|\lambda_j|$. Figure 1 shows a plot of the magnitude and phase of each component of the three Koopman modes with the largest growth rates and magnitudes. These were labeled as modes 7, 8, and 9 and corresponded to frequencies of 1.3078, 1.0962, and 0.3727 Hz,

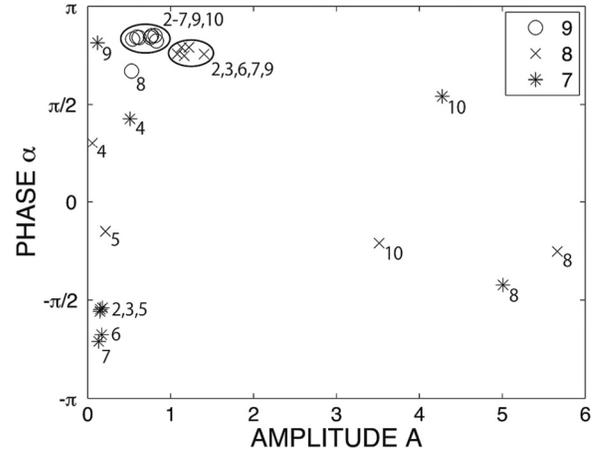


FIG. 1. Power Systems—Amplitudes and phases for the components of the three Koopman modes having largest growth rates and norms. Koopman modes and eigenvalues were computed using the DMD algorithm. The largest modes are labeled as modes 7, 8, and 9. The amplitudes and phases of the components of mode 9 are plotted with the symbol \circ . The numbers in the plot correspond to component of the mode. Modes 7 and 8 are plotted similarly, expect with symbols $*$ and \times , respectively. Generators 2, 3, 6, 7, and 9 are coherent with respect to mode 8. All generators, except generator 8, are coherent with respect to mode 9.

respectively. For each mode $j = 7, 8, 9$, there were amplitudes A_{ji} and phases α_{ji} for the generators $i = 2, \dots, 10$. In the figure, the amplitudes and phases for modes 7, 8, 9 are plotted with symbols $*$, \times , and \circ , respectively. Number labels within the plots specify the generator. It is seen that generators 2, 3, 6, 7, and 9 are coherent with respect to mode 8, while all but generator 9 are coherent with respect to mode 9. Therefore, generators 2, 3, 6, 7, and 9 were coherent with respect to both modes 8 and 9, as one found with visual inspection of the trajectories. While not done in this paper, the amplitudes and phases represented in this way allow using a number of clustering algorithms to automatically identify coherency.

2. Jet in crossflow

One of the earliest applications of Koopman modes was to the study of nonlinear fluid flows. Rowley *et al.*⁴³ introduced the concepts to the fluids community and demonstrated the methodology by computing a subset of the Koopman modes, using the DMD algorithm, for a jet in a crossflow. The jet in a crossflow configuration is a common way of mixing the jet fluid with a uniform crossflow. The crossflow moves parallel to a flat plate and the jet is injected through an orifice in the plate. It is known that such a system can exhibit self-sustained oscillations,⁴³ and Koopman modes were used to automatically identify the relevant frequencies and corresponding three-dimensional flow structures.

The flow field was studied in a $(L_x, L_y, L_z) = (75, 20, 30)\delta_0^*$ computational box, where δ_0^* was the displacement thickness at the crossflow inlet. The incompressible Navier-Stokes equations over a flat plate were solved using a Fourier-Chebyshev spectral method with a grid resolution of $256 \times 201 \times 144$ (see the reference for all the simulation details). Therefore, each point in the state space M is a sequence of Fourier-Chebyshev coefficients. The vector-valued observable, $F : M \rightarrow \mathbb{R}^m$, was chosen as the velocity measurements of the flow field at the

grid points. Hence, $m = 3(256 \times 201 \times 144) \approx 2.2 \times 10^7$ (three velocity components at each grid point). The situation is similar to the heat equation example above (Example 4, p. 14) where the state space was a space of Fourier coefficients and the observable was a heat distribution on a square.

Remark 24. The choice of state space M and a basis for the observables are not unique. For a fluid experiment on some physical domain \mathbb{B} , we assume that the solutions exist in some function space. Usually this is the space of finite energy flows so that a velocity profile $\mathbf{u}(\mathbf{x})$ exists in $L^2(\mathbb{B}, d\mathbf{x})$. Depending on the boundary conditions, many different bases for $L^2(\mathbb{B}, d\mathbf{x})$ exist. The particular basis chosen depends on the computational method used to solve the Navier-Stokes equations. As discussed, this method was a Fourier-Chebyshev spectral method; the resulting basis for $L^2(\mathbb{B}, d\mathbf{x})$ consisted of trigonometric and Chebyshev polynomials. The state space M was the sequence space of Fourier-Chebyshev coefficients corresponding to $L^2(\mathbb{B}, d\mathbf{x})$ functions. However, if periodic boundary conditions for the fluid flow are assumed, the trigonometric polynomials could be used to represent $L^2(\mathbb{B}, d\mathbf{x})$ and M would be the space of Fourier coefficients. Therefore, the choice of solution method for the Navier-Stokes equations determines the basis functions and implicitly defines the state space M that represents the system. \square

An initial velocity profile and boundary conditions were specified for simulations and this configuration corresponded to a fixed $p \in M$. Letting $\mathbf{b}_k := [U^k F](p)$, the DMD algorithm was performed on the sequence of observables $\{\mathbf{b}_{200}, \mathbf{b}_{202}, \dots, \mathbf{b}_{700}\}$. The reason for the delay in the start time of the sequence was to neglect the transient terms.

The top row of Figure 2 shows the time signals of the streamwise velocity recorded by a sensor probe close to the wall, just downstream from the jet orifice and a probe located downstream and on the jet shear layer. In the bottom row, the spectral content of the time signals are shown in black, whereas the spectral peaks identified by the DMD algorithm

are shown in red. Note that for the probe near the wall (left column of Figure 2) the signal contains only low-frequency components, whereas the signal near the jet contains both low- and high-frequencies. The probes are local measurements and only pick up a subset of the full spectrum for the fluid flow. The Koopman modes are global objects and, as shown in the figure, the DMD algorithm identifies both the low- and high-frequency components of the flow field.

Figure 3 shows the spectrum of the Koopman operator as computed by the DMD algorithm. Most eigenvalues lie on the unit circle implying that the flow field is near an attractor. The time-averaged flow (steady-state component) corresponds to $\lambda = 1$ and is indicated in blue in the left image of Figure 3. The rest of the eigenvalues have colors smoothly varying from red to white with the colors corresponding to the magnitude of the associated global mode. Red corresponds to large magnitudes for the Koopman modes, white to low magnitude. The magnitudes are given by the total energy of the mode (2-norm). The right image of Figure 3 shows the magnitudes of the Koopman modes at each frequency. The color scheme is the same as for the left image.

If we order the Koopman modes in order of decreasing magnitude, mode 1 corresponds to the time-average flow and the rest come in complex-conjugate pairs; modes 2 and 3 correspond to complex-conjugate eigenvalues and have the same magnitude. Figure 4, shows the streamwise velocity components of mode 2 (left) and mode 6 (right). Each mode oscillates at a single frequency, with mode 2 corresponding to a high-frequency ($St = 0.141$) and mode 6 to a low-frequency ($St = 0.0175$); T is the Strouhal number. These correspond to the tallest red line and the left most red line in the bottom row of Figure 2, respectively. In both images of Figure 4, the red surfaces correspond to positive streamwise velocities and blue surfaces to negative streamwise velocities. Mode 2 is associated with shear layer vortices with additional vortices extending toward the wall. Mode 6 has

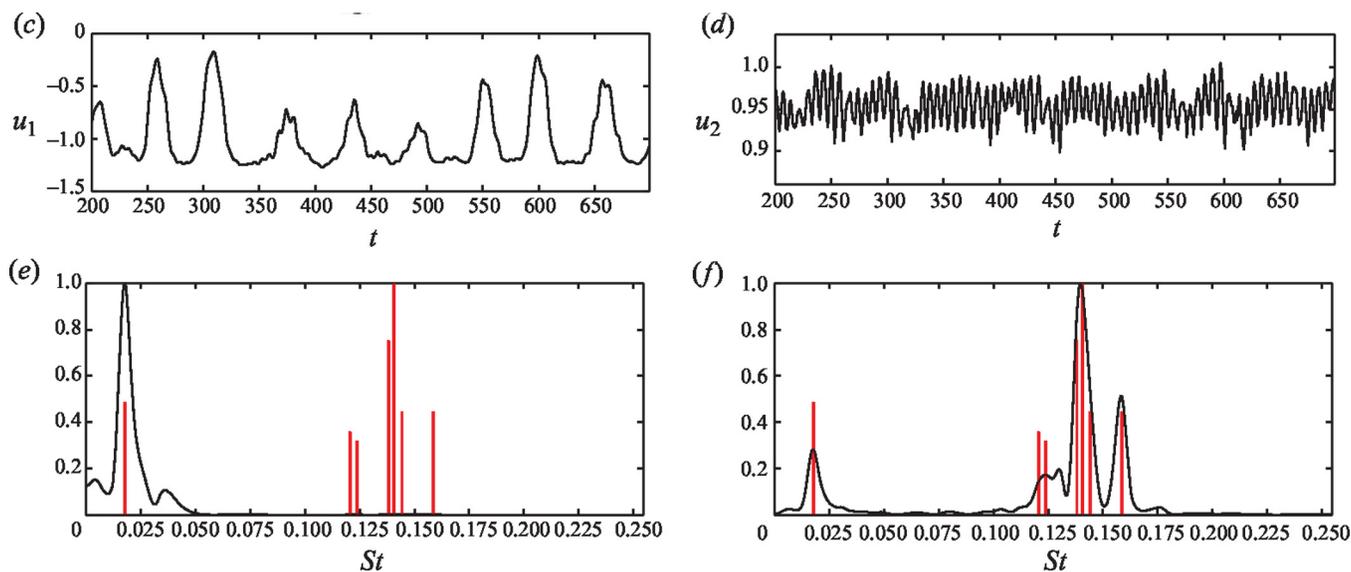


FIG. 2. Jet in Crossflow—The top row is the time signal of the streamwise velocity for a probe near the wall, downstream of the jet orifice (left) and for a probe downstream, in the jet trajectory (right). In the bottom row, the spectral content of the corresponding probe signals are shown in black. In red, the part of the spectrum of the Koopman operator captured by DMD algorithm. Only the positive frequencies are shown since eigenvalues occur in complex-conjugate pairs. Reprinted with permission from Rowley *et al.*, *J. Fluid Mech.* 641, 115–127 (2009). Copyright 2009 Cambridge University Press.

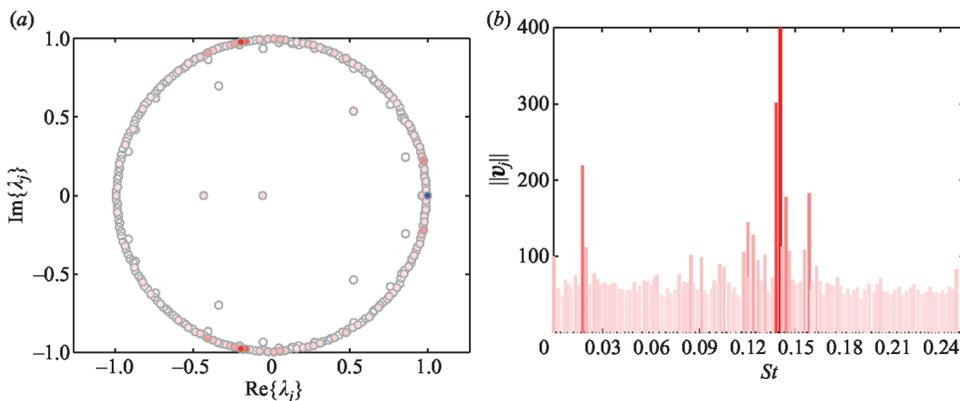


FIG. 3. Jet in Crossflow—(left) The spectrum of Koopman operator as identified with the DMD algorithm. Most of the Koopman eigenvalues are on the unit circle. The mode corresponding to the time-averaged flow (corresponding to $\lambda = 1$) is indicated in blue. The other eigenvalues are colored from red to white based on the total energy of the associated Koopman mode. Red corresponds to high-energy modes, white to low-energy modes. (right) The magnitudes of the Koopman modes at each frequency. The color scheme is the same as for the image on the left. Reprinted with permission from Rowley *et al.*, *J. Fluid Mech.* **641**, 115–127 (2009). Copyright 2009 Cambridge University Press.

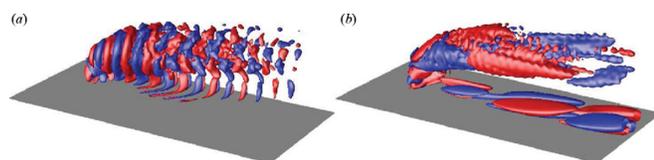


FIG. 4. Jet in Crossflow—Koopman modes 2 (left) and 6 (right) corresponding to high- ($St_2 = 0.141$) and low- ($St_6 = 0.0175$) frequencies, respectively. T is the Strouhal number. Red contours correspond to positive streamwise velocities and blue contours to negative streamwise velocities. Reprinted with permission from Rowley *et al.*, *J. Fluid Mech.* **641**, 115–127 (2009). Copyright 2009 Cambridge University Press.

large structures along the wall associated with shedding of the wall vortices. This shedding of wall vortices is coupled to the main jet body as indicated by the mode having structure along the jet body.

3. Self-sustained oscillations in a turbulent cavity

Seena and Sung⁴⁹ investigated the causes of self-sustained pressure oscillations in fluid flow over a cavity (see Figure 6 for a profile of the experiment geometry) by using the dynamic mode decomposition algorithm. We denote this domain by \mathbb{B} . The goals of the study were to identify the vortical structures that drove the hydrodynamic oscillations and

obtain dynamical information about those structures.⁴⁹ Both thin and thick incoming boundary layers were studied (Reynold’s numbers at the cavity of 12 000 and 3000, respectively); only the turbulent case ($Re = 12\ 000$) is covered here.

As with the case of the jet in a crossflow example, the fluid evolution was governed by the incompressible Navier-Stokes equations corresponding to the specified cavity geometry (see Seena and Sung⁴⁹ for details on the computational domain). Abstractly, the state space M was a sequence space of coefficients for basis functions on \mathbb{B} . Since the system was not periodic in all dimensions, the basis functions were not the 3-dimensional trigonometric polynomials (see Remark 24). A family of observables, parameterized by points in \mathbb{B} , was given by functions mapping a point in M (a sequence of coefficients) to the fluid pressure at a point $x \in \mathbb{B}$.

The solutions of the Navier-Stokes equations were computed using the Crank-Nicolson method using a second-order central difference scheme in space.⁴⁹ Cavity flows at the high Reynold’s number were simulated with a large Eddy simulation (LES).⁴⁹ Broadly speaking, LES filters the governing equations by removing the small-scale structures and replacing them with models. Refer Seena and Sung⁴⁹ for

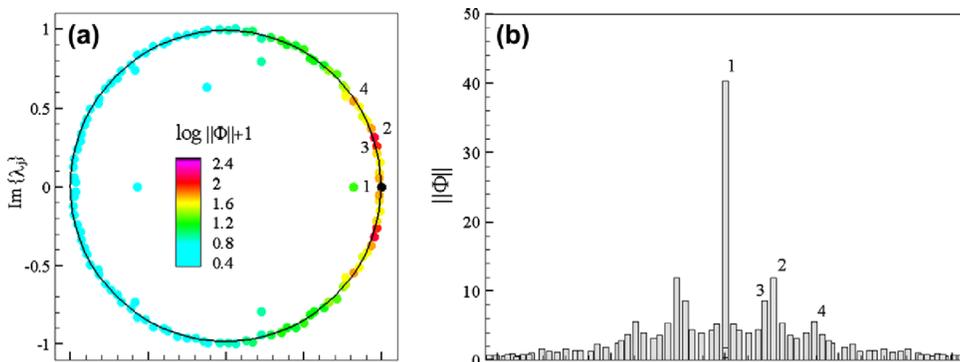


FIG. 5. Turbulent cavity flow—(left) Koopman eigenvalues computed using the DMD algorithm for turbulent flow inside the cavity at $Re = 12\ 000$. The eigenvalues are colored based on the energy of the corresponding Koopman mode. (right) The energy of the Koopman mode at each frequency, ω (rad/s). Reproduced by permission from A. Seena, *Int. J. Heat Fluid Flow* **32**(6), 13 (2011). Copyright 2011 by Institution of Mechanical Engineers.

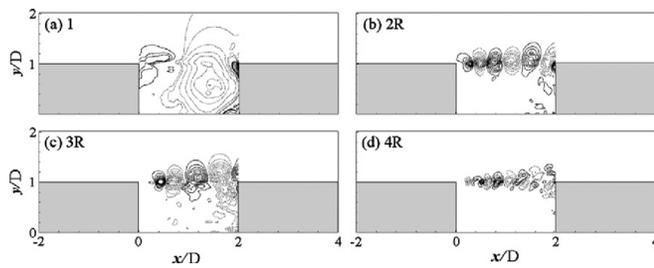


FIG. 6. DMD modes inside the cavity at $Re = 12\,000$ for modes labeled 1, 2, 3, and 4 in Figure 5. Reproduced by permission from A. Seena, *Int. J. Heat Fluid Flow* 32(6), 13 (2011). Copyright 2011 by Institution of Mechanical Engineers.

full details of the simulation parameters and Germano *et al.*¹⁹ for a discussion of LES.

The vector-valued observable chosen for analysis was the fluid pressure at each of the computational grid points in and above the cavity. The DMD algorithm was applied to a sequence of 124 flow snapshots recorded after allowing transients to decay.⁴⁹ Figure 5 shows the Koopman eigenvalues resulting from the DMD algorithm. On the left, almost all of the eigenvalues are seen to be on the unit circle, implying that the flow field is on or near an attractor. Colors in the plot correspond to the total energy of the corresponding Koopman mode. On the right, the total energies of the Koopman modes at each frequency are shown. The four dominant peaks are labeled. Only positive frequencies are marked since eigenvalues occur in complex-conjugate pairs. The mode labeled as 1 occurs at $\omega = 0$ and corresponds to the steady component of the flow. The corresponding Koopman mode is shown in Figure 6(a). Solid, black lines correspond to high-pressure regions, whereas broken, gray lines correspond to low-pressure regions. Koopman modes 2, 3, and 4, corresponding to the labeled peaks in Figure 5(b), are shown in Figures 6(b)–6(d). Modes 2 and 3 correspond to $\omega = 4.6$ rad/s and 3.5 rad/s, respectively. The exact frequency corresponding to mode 4 was not reported, but was less than 10 rad/s. The low- and high-pressure regions of the modes oscillate at a fixed frequency and suggest self-sustained oscillations in the cavity.⁴⁹ These results were consistent with results reported in literature.⁴⁹ The study successively used the DMD algorithm to identify the large vortices in the cavity driving the self-sustained oscillations.

Remark 25. Figures 3 and 5 show that the Koopman spectrum, as computed by the DMD algorithm, has some eigenvalues inside the unit circle, even though an initial block of data points was discarded before computation so that transients could die out. They are most likely spurious eigenvalues resulting from the finite truncation of the data and the DMD algorithm and not due to slowly decaying modes contained in the data. Unfortunately, no detailed investigation has been done on such eigenvalues; most attention has been focused on the modes corresponding to eigenvalues on the unit circle and those inside the unit circle have been ignored. \square

4. Energy efficiency in buildings

Energy use and efficiency in buildings has received much attention in recent years and Koopman mode analysis

has recently seen application in this field, as well. A researcher is generally interested in temperature distributions in buildings, as this can tell much about heating, ventilation and air conditioning (HVAC) and controller performance or if the building is operating near its design point.^{16,18} Such measurements can also be used for model validation.^{16,18}

Analysis requires understanding of the heat flow in the building subject to forcing from weather or even human traffic between areas of the building. In the most general case, heat transfer equations must be solved for a complicated domain with varying boundary conditions. The interesting observables are the temperatures at each point in the building, \mathbb{B} . The temperature distributions in the building are assumed to exist in some function space ($L^2(\mathbb{B}, dx)$, for example). A temperature distribution can be represented in a basis $\{b_1, b_2, \dots\}$ for the function space. The state space M is the sequence space of coefficients for these basis functions; i.e., a point $p \in M$ is $p = (c_1, c_2, \dots)$, for some coefficients $c_k \in \mathbb{C}$. An observable maps a point in M to the temperature at a point $x \in \mathbb{B}$. When collecting real data, the temperature is measured at a finite number of points, usually dictated by the placement of the temperature sensors by the building designers. The vector-valued observable of interest is the vector of the temperatures recorded by the building sensors.

In the most general case, the full model for the building is too complicated to solve. A variety of simplifying assumptions and a numerical package is needed to compute solutions to the building system. EnergyPlus¹¹ is a simulation package for modeling energy and water use in buildings. It is a free program offered by the U.S. Department of Energy and is widely used.³⁸ At a high level, an EnergyPlus building model consists of specifying the location of every surface in the building. A list of rooms is then created and the interaction of the rooms (between shared surfaces) is specified. For example, heat conduction through a certain type of material may be specified for one wall, whereas radiative heat transfer is important for windows. Models can also include weather data, HVAC usage, room occupancy, and building water usage. The software makes the assumption that rooms are well-mixed so that no temperature gradients exist in the room. Under the assumptions of the software, heat transfer in the building is approximated by a system of coupled ordinary differential equations.

In Eisenhower *et al.*,¹⁶ Koopman mode analysis was used to decompose the temperature evolution in a building into purely periodic global modes. The building investigated was the Y2E2 building at Stanford University and was modeled and simulated in the EnergyPlus software package. The physical building has 2370 HVAC sensors recording data throughout the building.¹⁶ The vector-valued observable chosen was the vector consisting of temperature readings taken at the sensors on the second floor.¹⁶ Koopman mode analysis was used to validate the EnergyPlus model by comparing the most energetic Koopman modes for the simulation and the raw data.

Figure 7 shows the spectral content of the Koopman modes computed using the DMD algorithm with the sensor data on top and the EnergyPlus data on the bottom. The

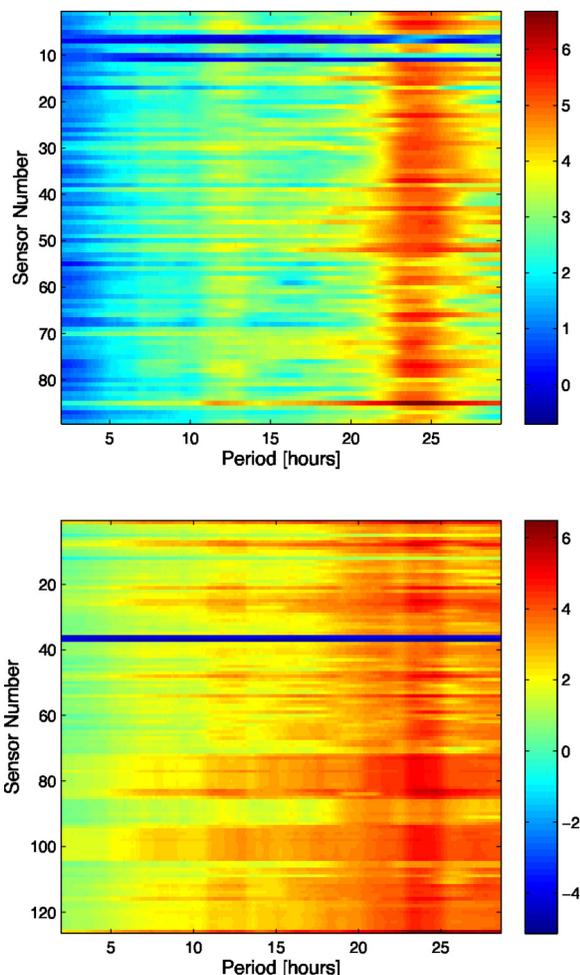


FIG. 7. Y2E2 building decomposition—Koopman spectra for the sensor data (top) and the EnergyPlus model (bottom). A large spectral content is seen at the 24 h period due to outdoor forcing conditions. Horizontal blue lines correspond to near constant temperatures at those sensors. These sensors correspond to rooms that are served by their own fan units allowing a tight control of temperature. Reproduced by permission from Eisenhower *et al.*, “Decomposing building system data for model validation and analysis using the koopman operator,” in SimBuild 2010—Fourth National Conference of IBPSA, USA, August 2010.

sensors are along the vertical axis with the period of the Koopman mode along the horizontal axis. A vertical streak in Figure 7 corresponds to a “global” mode; i.e., most sensors are affected by the Koopman mode at that frequency. In both images, a strong streak is seen at the 24 h period which is due to the daily forcing from the weather. Horizontal blue lines indicate near constant temperatures at those sensors and correspond to rooms served by their own fan units.¹⁶

Figure 8 shows the 24 h Koopman mode for both the sensor data and the EnergyPlus model. The top row corresponds to the sensor data, the bottom row to the simulation. The magnitude of the mode is shown on the left with the phase on the right. Note that the magnitude and phase of the Koopman mode is reported in reference to the outside air temperature (OAT). The relative magnitude of a mode is given in decibels (dB)

$$|KM_i| = 20 \log_{10} \left| \frac{C_i(F)}{C_{OAT}(F)} \right|,$$

and the phase is given in degrees

$$\angle KM_i = \angle C_i(F) - \angle C_{OAT}(F),$$

where $C_i(F)$ and $C_{OAT}(F)$ are the Koopman modes as computed by the DMD algorithm. For the sensor data, the magnitude of the 24 h mode relative to the external temperature was about -6 dB throughout the building. This implied that the peak temperature oscillation inside the building was smaller than the magnitude of the temperature fluctuations outside. For the EnergyPlus mode, the temperature magnitude inside was much closer to the external fluctuations since the relative magnitude of the mode was about 0 dB. Significant deviations in the relative phases were also seen between the sensor data and the model. These discrepancies implied that the parameters used in the simulation model were incorrect. The analysis lead to suggestions on the model parameters to modify.¹⁶

In practice, even under the simplifying assumptions utilized in programs such as EnergyPlus, a detailed model of the building can be prohibitively expensive to simulate over the time scales needed (e.g., hours up to years). Methods for model reduction are in order. Zoning approximations are one approach. In a detailed EnergyPlus model, each room is treated as a unique thermal zone. Each room is assumed well-mixed so that the thermal properties are uniform in the room. The idea of zoning is to lump adjacent rooms together in such a way that the thermal properties can then be assumed uniform across those rooms. This procedure results in less regions that need to be simulated in the model. Usually, zoning approximations are performed heuristically.¹⁸ However, model accuracy is quite sensitive to the zoning approximation used.¹⁸

Georgescu *et al.*¹⁸ used the notion of coherency between Koopman modes (Definition 23 above) to create zoning approximations for buildings and studied the Engineering Sciences Building (ESB) at the University of California, Santa Barbara as a test case for the methodology. The particular observable chosen was $F = (f_1, \dots, f_m)^T$, where each $f_j : M \rightarrow \mathbb{R}$, $j = 1, \dots, m$, represented the temperature of a room in the ESB. In the detailed EnergyPlus model, there were $m = 191$ zones. An additional physical assumption was imposed so that rooms were grouped into a zone if they were both on the same floor and adjacent as well as ϵ -coherent.

The modes of interest were those having periods of one year, 24 h, 12 h, 8 h, and 6 h. These corresponded to the most energetic modes. Figure 9 shows the Koopman modes of the ESB EnergyPlus simulation for the three most energetic modes. The magnitudes of the Koopman modes are given in the top row. Units are degrees Celsius relative to the average temperature in the room. The average temperatures were not reported. The bottom row corresponds to the phase in radians of the Koopman mode. Consider the block of four rooms in the center of the right hand side of the building (the ones that are dark blue in the magnitude plot of the year long mode). By visual inspection, the middle two rooms on this block would be lumped into a single zone. While the colors of those two rooms vary between each of the six images in Figure 9, within the same image those two rooms have the same

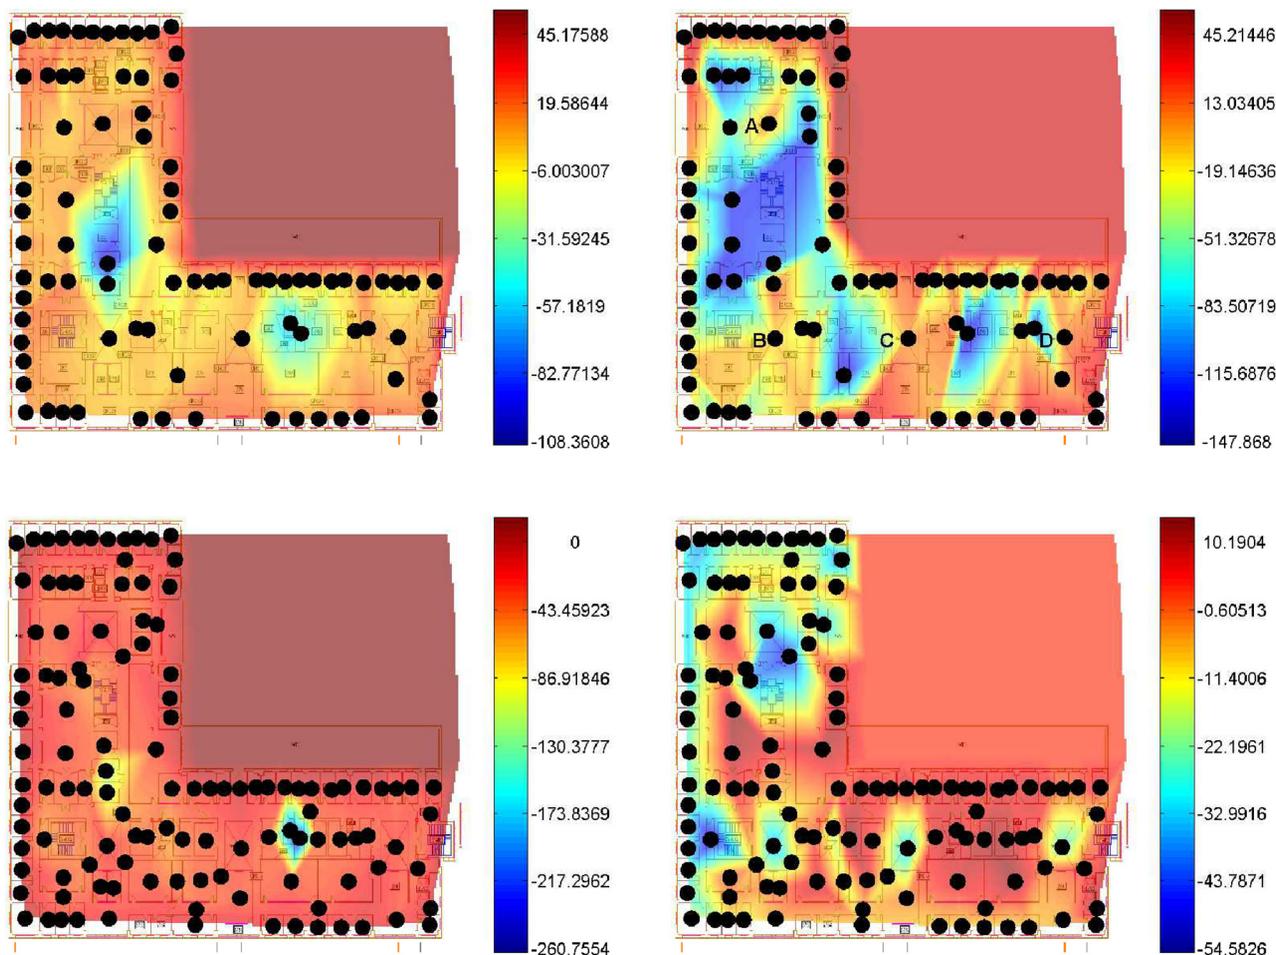


FIG. 8. Y2E2 building decomposition—Koopman mode for the 24h period. (top row) The magnitude and phase, respectively, of the raw data’s Koopman mode. (bottom row) The magnitude and phase of the EnergyPlus model’s Koopman mode. The magnitude units are given in decibels and phase in degrees. Both are relative to the external 24h mode. The discrepancy between scales shows the model mismatch with the real data. Reproduced from Eisenhower *et al.*, “Decomposing building system data for model validation and analysis using the koopman operator,” in SimBuild 2010—Fourth National Conference of IBPSA, USA, August 2010.

color, and hence almost identical values. These two rooms satisfy the ϵ -coherency definition for some small $\epsilon > 0$.

IV. ANALYSIS OF STATE SPACES USING EIGENQUOTIENTS

Sections I–III discussed the Koopman eigenfunctions in the context of spectral decomposition of the Koopman operator, and we did not pay much attention to values that eigenfunctions take on the state space. In this section, we demonstrate that eigenfunctions carry information about transport between parts of the state space and provide a way to identify invariant sets. Going further, we will endow the collections of invariant sets with their own metric topology, which allows for associating smaller invariant sets into larger coherent structures.

The material in this section appeared originally in a string of papers^{6,7,26,36,37,52} with the technical details given also in dissertations of two of the authors.^{5,33}

A. State-space analysis of Koopman eigenfunctions

The level sets of Koopman eigenfunctions at eigenvalues $|\lambda| = 1$ form invariant sets, and periodic and wandering chains of sets in the state space. The eigenfunction level-set

partitions depend not only on the choice of the eigenspace E_λ from which eigenfunctions are taken, but also on the choice of a particular function in the eigenspace. However, an exhaustive procedure for analysis of level-set partitions can be devised, based on partition products, such that the result in the limit does not depend on the particular choice of eigenfunctions studied, but rather only on the eigenvalue λ , which determines the dynamics of level sets, e.g., periodic or invariant. In this section, we first focus on the ergodic partition, which corresponds to analysis of the invariant eigenspace of the Koopman operator, i.e., $\lambda = 1$. The generalization to “periodic” eigenspaces, i.e., those corresponding to $\lambda = e^{i2\pi\omega}$ for $\omega \in \mathbb{Q}$, is straightforward, and we present it at the end of this section.

From this point onward, we will assume the measurable setup of the Koopman operator: $T : M \rightarrow M$ will be a measurable map between measure spaces (M, \mathfrak{B}) , with at least one invariant measure μ . The observables \mathcal{F} are, at the very least, a subset of complex-valued measurable functions. Given an invariant function ϕ , i.e., a function in invariant eigenspace of the Koopman operator $\phi \in E_1(U)$, we can form its level set partition $\zeta(\phi) := \{S_z : \forall z \in \mathbb{C}\}$, where the level sets are $S_z := \{x \in M : \phi(x) = z\}$. Since ϕ is a

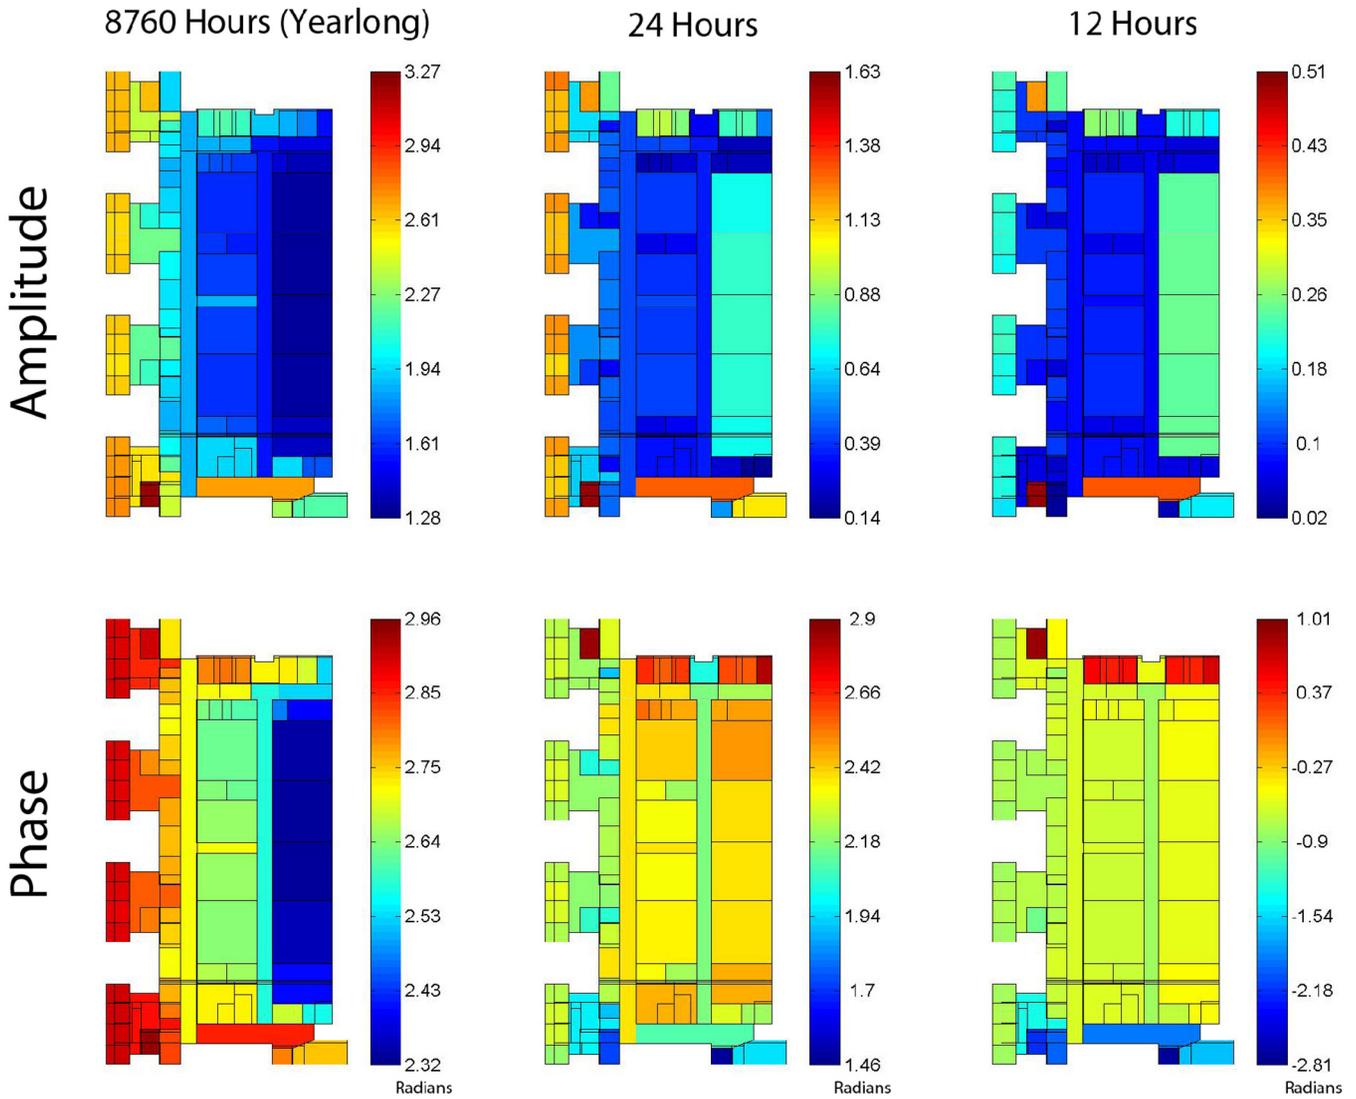


FIG. 9. Zoning approximations of buildings—Koopman mode magnitude and phase for the modes having the largest magnitudes. Only one floor of the ESB is shown. (Top) The magnitudes of the Koopman modes. Units are in degrees Celsius relative to the mean temperature of the room. (Bottom) The phase of the Koopman mode in radians. Reproduced from Georgescu *et al.*, “Creating zoning approximations to building energy models using the Koopman operator,” in SimBuild 2012—Fifth National Conference of IBPSA, USA, August 2012.

measurable function, its level sets S_z are measurable sets. Additionally, they are invariant sets, by virtue that each S_z collects *all* state-space points on which ϕ takes the value z , i.e., if it contains x , it will also contain $T(x), T^2(x), \dots$, since ϕ is constant along trajectories. A partition of the state space into invariant sets is called a *stationary partition*.

Unless the system is ergodic, the choice of $\phi \in E_1$ is not unique (see Remark 12): often there will be at least two functions $\phi, \psi \in E_1$, which are not linearly related and yield two different partitions $\zeta(\phi), \zeta(\psi)$. The *product of two partitions* $\zeta(\phi) \vee \zeta(\psi)$, which contains intersections of sets in each of the partitions (see Fig. 10), is again a stationary partition. The product is *finer* than either of the original partitions, as sets in $\zeta(\phi)$ or $\zeta(\psi)$ can be recovered as unions of sets in $\zeta(\phi) \vee \zeta(\psi)$. Naturally, given a finite set of functions $\phi_k, i = 1, \dots, K$, the product $\bigvee_{k=1}^K \zeta(\phi_k)$ of the associated partitions can be formed incrementally.

Given a partition ζ of the space and some set A , the *quotient map* $\pi_\zeta : M \rightarrow A$ is any map that separates the partition

$\zeta : \pi_\zeta(x) = \pi_\zeta(y)$ iff $\exists S \in \zeta$ such that $x, y \in S$. In other words, all the points in any $S \in \zeta$ will be mapped to a single point $p \in A$. The image set $\xi := \pi(M)$ is termed the *quotient* ξ of the partition ζ . If we are given just the map π_ζ , we could reconstruct the partition ζ from level sets of π_ζ , in terms of values of the quotient map codomain A .

For a product partition $\zeta_K = \bigvee_{k=1}^K \zeta(\phi_k)$, the most straightforward way to construct a quotient map is by arranging

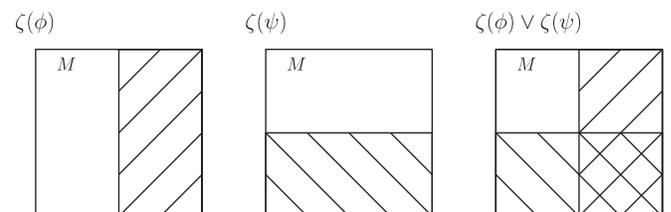


FIG. 10. Partitions of the state space M into $\zeta(\phi)$, $\zeta(\psi)$ and their product $\zeta(\phi) \vee \zeta(\psi)$.

eigenfunctions in a vector-valued function to define $\pi_K : M \rightarrow \mathbb{C}^K$, $\pi_K(x) := (\phi_1(x), \phi_2(x), \dots, \phi_K(x))$. Notice, however, that the quotient map is not unique: any bijection $h : \mathbb{C}^K \rightarrow A$, e.g., a coordinate transform, can be used to form another quotient map through composition $h \circ \pi_K$. While the construction of π_K might appear academic, it is significant in that it translates the problem of constructing partition ζ_K from the topological setting, where intersections of sets were used to check whether two points x, y are in the same product set, to an analytical setting, where the equality $\pi_K(x) = \pi_K(y)$ can be checked, either numerically or analytically. We will make use of this fact heavily in Sec. IV B.

Evaluation of eigenfunctions ϕ_k , which feature in construction of quotient maps, is easy when we are given a method for evaluating the projection of the space of observables onto an eigenspace $P_\lambda : \mathcal{F} \rightarrow E_\lambda$. In this case, we can compute a dense set of eigenfunctions in E_λ by evaluating $P_\lambda f$ for a dense set $f_k \in \mathcal{F}$. Moreover, when \mathcal{F} is a Hilbert space, we can compute the projection on an orthogonal basis for \mathcal{F} , obtaining a (not necessarily orthogonal) basis for E_λ . A practical way to numerically evaluate P_λ for $|\lambda| = 1$ using trajectory averages is explained in Sec. IV C.

For $\lambda = 1$, projecting functions f_k using P_1 results in an infinity of invariant functions that we can attempt to use in forming incremental level-set partitions. It is not trivial that the partition intersection can be extended to the case of an infinity of U -invariant functions ϕ_k . The first result of this sort was obtained by Sine⁵¹ for continuous observables $\mathcal{F} = C(M)$. Such setting is fairly restrictive for deterministic dynamics. Therefore, we focus on a more general result^{36,37} that holds for $\mathcal{F} = L^1(M, \mu)$. Assume we are given a countable, bounded family $\phi_k \in E_1 \subset L^1(M, \mu)$, for $k = 1, 2, \dots$, whose span is dense in the invariant eigenspace E_1 of the Koopman operator U . The incremental partition into level sets

$$\zeta_e = \bigvee_{k=1}^{\infty} \zeta(\phi_k)$$

exists as the level set partition of the map $\pi_e : M \rightarrow \ell^\infty$, defined as $\pi_e := (\phi_1, \phi_2, \dots)$, where ℓ^∞ is the space of bounded sequences. The map π_e is termed the *ergodic quotient map*, with *ergodic quotient* $\zeta_e := \pi_e(\zeta)$ as its image set.

While the concept of the ergodic partition was known since at least Rokhlin,^{41,42} this computationally feasible construction is fairly recent.^{36,37} The justification for the name of the ergodic partition comes from analyzing restrictions of T to invariant sets $S \in \zeta_e$. Each S carries a measure μ_S such that the restricted dynamics $T : S \rightarrow S$ is *ergodic with respect to μ_S* . By one of the several equivalent definitions of ergodicity,⁴⁰ this means that for any $f \in L^1(S, \mu_S)$

$$\int_S f \mu_S = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} f \circ T^n(x), \tag{60}$$

at almost every $x \in S$, with respect to μ_S . The measures μ_S are called *ergodic measures* and are extreme points of the space of invariant measures (by the Ergodic Decomposition Theorem).²¹ Moreover, integrals against ergodic measures

define the projection onto the invariant eigenspace $P_\lambda f(x) = \int_{S(x)} f d\mu_{S(x)}$, where $S(x)$ denotes the ergodic set containing point x . This is a consequence of the equality (Eq. (60)) along with Mean Ergodic Theorems,⁵ e.g., von Neumann theorem for $\mathcal{F} = L^2(M, \mu)$ or Yosida ergodic theorem for general Banach spaces, used in earlier sections as Theorem 16.

The ergodic partition is unique up to μ -measure zero sets, i.e., given two stationary partitions ζ_1 and ζ_2 that satisfy the above properties, for any $S_1 \in \zeta_1$ there exists $S_2 \in \zeta_2$, such that $\mu(S_1 \Delta S_2) = 0$. Consequently, the partition ζ_e does not depend on the choice of functions ϕ_k used to construct it: different choices of functions just yield different *representations* of the quotient map π_e ; however, the limit partition will be the same.

When there is only one T -invariant measure μ , up to multiplication by a scalar, the system is *uniquely ergodic*. In that case, the eigenspace E_1 is one-dimensional, containing only a.e. constant functions. In turn, the ergodic partition contains a single set, and the ergodic quotient map maps almost all points (with respect to μ), into a single point in the sequence space ℓ^∞ . While uniquely ergodic systems might be of interest in general, they are not the target of this approach. The ergodic quotient analysis is intended to help the analysis of state spaces which contain a lot of ergodic measures, e.g., systems with families of periodic or quasi-periodic orbits and systems possessing interspersed regions of regular and irregular dynamics.

For most conservative regular systems, ergodic partitions are similar to partitions of the state space into orbits. However, in zones of irregular dynamics, e.g., on strange attractors and in zones of strong mixing, it is difficult to gain intuition about behavior of the systems simply from orbits. Orbits are inherently zero-dimensional (for maps) or one-dimensional sets (for flows), but in chaotic regimes, every one of them densely fills a bigger set, possibly even a set of a positive invariant measure, e.g., positive volume. Moreover, in a mixing zone any two trajectories look nothing alike, yet they are contained in the same set and have the same statistical behavior. Through such reasoning, we might be interested not in description of trajectories themselves, but rather in description of minimal invariant sets containing individual trajectories. Precisely, the ergodic partition ζ is a *measurable hull* of the decomposition of the space into orbits $T^n(x)$, i.e., a partition of the space into minimal measurable invariant sets that contain orbits.⁴² In this sense, ζ is the appropriate counterpart of the state space portrait in the context of measurable dynamical systems.

The understanding of the ergodic partition, derived from functions in the invariant eigenspace E_1 , can be extended to eigenspaces E_λ for which $|\lambda| = 1$. When the eigenfunctions are taken from a “periodic” eigenspace E_λ , for an eigenvalue with property $\lambda^b = 1$, for some $b \in \mathbb{N}$ (period), the extension is straightforward. Level sets of $\phi \in E_\lambda$ are then periodic sets, which are arranged in *b-periodic chains*, $\{S_a, T(S_a), T^2(S_a), \dots, T^{b-1}(S_a)\}$. Similarly, if $\lambda = e^{i2\pi\omega}$ is such that ω/\mathbb{Q} , the chains of level sets extend into infinity, forming *wandering chains*. The partition $\zeta(\phi)$ is again an invariant partition, in the sense that for each $S \in \zeta(\phi)$, $T(S) \in \zeta(\phi)$. However, $\zeta(\phi)$ is not stationary, as S are not invariant sets; it

is an invariant b -periodic partition or, when ω/\mathbb{Q} , an invariant wandering partition. The partition products also generalize: if $\phi, \psi \in E_\lambda$, $\zeta(\phi) \vee \zeta(\psi)$ will also be a b -periodic/wandering invariant partition.

The invariant measures can be generalized to the concept of complex eigenmeasures,³⁶ which satisfy

$$\mu(T^{-1}S) = \lambda^{-1}\mu(S).$$

Integrating against extrema of eigenmeasures enables us to evaluate P_λ away from $\lambda = 1$, and construct *eigenquotient* maps $\pi_\lambda(x) = (\dots, P_\lambda f_k, \dots)$, analogously to constructions of ergodic quotient maps from a basis f_k for \mathcal{F} . In this sense, the ergodic quotient map is the eigenquotient map at $\lambda = 1$.

The eigenquotient maps collect basis functions for the eigenspace E_λ . Note that this basis set might be overdetermined, since we do not know ahead of time what $\dim E_\lambda$ is. In Remark 12, we have indicated that the dimension of E_λ is bounded by the number of mutually singular components of measure μ which was used to define the space of observables as $\mathcal{F} = L^p(M, \mu)$. These components are precisely the ergodic measures, and there are as many of them as there are ergodic sets.

Given an eigenfunction $\phi \in E_\lambda$, for $|\lambda| = 1$, its modulus and complex angle functions are, respectively, $|\phi|$, $\angle\phi$, for which $\phi(x) = |\phi(x)|\exp[i2\pi\angle\phi(x)]$. The modulus function $|\phi|$ is an invariant function, since $U|\phi| = |\phi \circ T| = |\lambda\phi| = |\phi|$. On the other hand, $\angle\phi \in E_\lambda$ is a factor map: it conjugates the dynamics of T with a dynamical system $\theta \mapsto \theta + \omega$, evolving on a circle \mathbb{S}^1 , with ω as the angle of the eigenvalue $\lambda = re^{i2\pi\omega}$,^{36,58} since

$$\angle\phi(Tx) = \angle\phi(x) + \omega.$$

To illustrate what these functions might reveal about dynamics, we will take the Chirikov Standard Map

$$\begin{aligned} x_{n+1} &= x_n + p_n + \epsilon \sin(2\pi x_n) \\ p_{n+1} &= p_n + \epsilon \sin(2\pi x_n), \end{aligned} \tag{61}$$

with $\epsilon = 0.15$. The system preserves the area measure m ; therefore, the Koopman operator $U : L^2(M, dm) \rightarrow L^2(M, dm)$ is unitary and its eigenvalues all lie on the unit circle. Figure 11 shows level sets of an invariant eigenfunction, and modulus and angle of an eigenfunction at $e^{i2\pi\omega}$ for $\omega = 1/3$, demonstrating the distinction between $\phi \in E_1$ and the modulus of $\psi \in E_\lambda$.

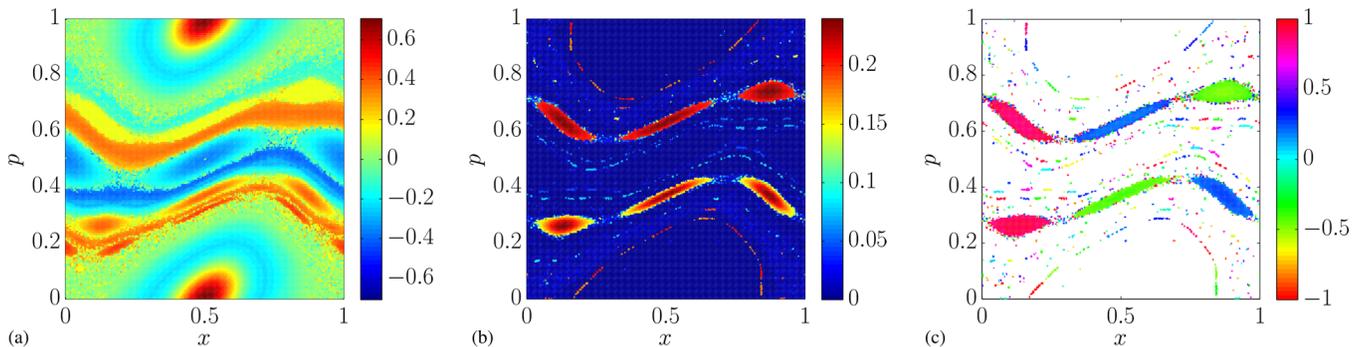


FIG. 11. Level sets of eigenfunctions $\phi \in E_1$ and $\psi \in E_\lambda$ for $\lambda = e^{i2\pi\omega}$ at $\omega = 1/3$ for the Chirikov Standard Map at $\epsilon = 0.15$. (Functions were evaluated by computing Fourier averages, see Sec. IV C). (a) Invariant eigenfunction $\phi \in E_1$. (b) Modulus $|\psi|$ of a 3-periodic eigenfunction ψ . Modulus of any eigenfunction for $|\lambda| = 1$ is an invariant function. (c) Angle $\angle\psi$ of a 3-periodic eigenfunction ψ . Angle of an eigenfunction at is a factor map for dynamics. (Only regions where angle is well defined, i.e., $\psi \neq 0$, are colored.)

Even though both $|\psi|$ and ϕ are invariant functions, $|\psi|$ is non-zero only over trajectories that are periodic with period 3. At $\epsilon = 0.15$, such trajectories are mostly concentrated in two large groups of periodic chains near the middle of the state space. On the other hand, invariant functions that are not formed by taking absolute values of periodic functions do not necessarily make a distinction between periodic and non-periodic dynamics. The angle function $\angle\psi$ can be used to infer the order in which sets are ordered in the associated periodic chains. Even such simple visualizations can be useful aids in quick assessment of a new dynamical system, by identifying regions that are not dynamically connected.

B. Geometry of eigenquotients

Section IV A, we explained how eigenfunctions of the Koopman operator connect to invariant and periodic structures in the state spaces of dynamical systems. We have also indicated that the finest invariant partitions ζ can be represented in sequence spaces using quotient maps $\pi : M \rightarrow \ell^\infty$ and their image sets, eigenquotients $\xi = \pi(M)$. In this section, we will be treating geometry of sets ξ , explaining it for the case of the ergodic quotient, which corresponds to invariant sets. As explained in Section IV A, a generalization to periodic sets is straightforward.

Identification of ergodic sets can be thought of as checking whether for two state space points $p_1, p_2 \in M$ it holds that $\pi(p_1) = \pi(p_2)$ or not; essentially, we are using a discrete topology on ergodic quotient ξ to compare points in the state space M . We can, however, use other metric topologies on the ergodic quotient to extract additional information about the state space, e.g., to obtain a low-dimensional representation of dynamics, or identify functions acting as integrals of motion over invariant sets, not necessarily entire state spaces. By a low-dimensional representation of dynamics, we mean that we are looking for a minimal set of directions in the state space, in which we can move a point across the boundary of an ergodic set S and land in another ergodic set S' where trajectories look similar, on average, to those in S .

As a motivation for the analysis of ergodic quotient we take the Morse theory analysis of Hamiltonian systems. Consider three state spaces sketched in the top row of Figure 12. Each state space contains a region where trajectories are

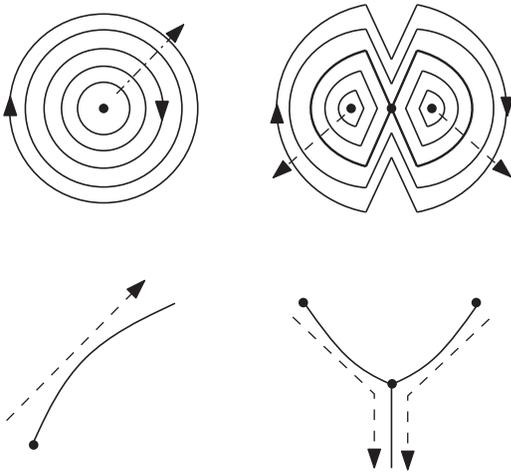


FIG. 12. Sketches of two Hamiltonian oscillators: a harmonic oscillator and a double-well oscillator. Top row: state space portrait. Bottom row: graphic representation of Fomenko-Reeb graphs for level-sets of Hamiltonian functions. Dashed arrows indicate how moving an initial condition across the level sets reflects on the Reeb graph. Reprinted with permission from M. Budišić, *Physica D Nonlinear Phenom.* 241(15), 1255–1269 (2012). Copyright 2012 North-Holland.

tightly layered together, and where they can be parametrized using a single continuous parameter: distance from elliptical fixed points. For Hamiltonian systems, these properties can be formally stated by inducing a topology on the state space based on the energy function. In other words, the “similarity” of neighboring ergodic sets that we are after is here represented by the similar values that the energy function attains on neighboring ergodic sets.

Level sets of any monotonic function of the Hamiltonian can be represented by *Reeb graphs* sketched next to the state spaces. Reeb graphs are a topological tool, featured in Morse theory, which analyzes manifolds through level sets of differentiable functions on them. The Reeb graphs provide a concise description of changes in topology of such level sets as the level, i.e., value of the function, is changed. In dynamical systems, they have been used to study level sets of integrals of motion for integrable Hamiltonian systems;^{4,15} in this context, they are known as Fomenko-Reeb graphs. Since integrals of motion are invariants of projections P_1 onto the $\lambda = 1$ eigenspace, in certain settings we could draw parallels between Fomenko-Reeb graphs and the ergodic quotient ξ .

The vertices in Fomenko-Reeb graphs correspond to those values of the Hamiltonian $H : M \rightarrow \mathbb{R}$ at which the topology of level sets of H changes. The edges connecting the vertices correspond to families of connected components of level sets that are homotopically related with respect to continuous variation of values of H . Figure 12 illustrates what the Fomenko-Reeb graphs look like for two simple Hamiltonian systems. Without going into details of their construction here, an important feature of these graphs is that continuous segments in them correspond to one-parameter families of periodic orbits. For the double-well oscillator, strands merge at the value of energy where two separate wells merge across the separatrix to form the outer layer of periodic orbits. Therefore, by looking at number of independent parameters

in a Fomenko-Reeb graph, we can deduce the number of families of periodic orbits in the state space.

Currently, Fomenko-Reeb graph analysis is formulated only for integrable Hamiltonian systems.⁴ By analyzing the ergodic quotient, we can establish an approach similar in spirit, even for systems that do not have an explicit energy function, but might contain families of invariant sets in the state space. The ergodic quotient map collects, in a sense, *all* possible invariant functions that can be thought as generalizations of energy. Unlike the case of Hamiltonian systems, where topology of the energy codomain was used to define topology on the state space, the ergodic quotient ξ is a subset of an infinite-dimensional sequence space ℓ^∞ , which can be endowed with many Euclidean-like topologies that are not all mutually equivalent. For this reason, we need to choose a metric structure (ξ, d) , where d is a distance function, so that we can obtain a context in which a generalization of Reeb analysis of Hamiltonian systems to non-Hamiltonian systems is possible.

The role played by the Hamiltonian function in Fomenko-Reeb graphs is taken up by the quotient map π . The image space of Hamiltonian was \mathbb{R} with its natural metric topology; the ℓ^∞ topology on the ergodic quotient ξ is too sensitive to establish analogous arguments. Instead, we will use a Sobolev space metric which requires interpretation of the image under the quotient map $\pi(x)$ as a spatial Fourier transform of ergodic measures μ_x . Within this setting, we seek to extract coherent onion layers in the state space by identifying connected components in (ξ, d) . The challenge lies in extracting a particularly appropriate low-dimensional parametrization of ξ even though it is embedded within the infinite-dimensional space ℓ^∞ .

First, as observables which will be projected onto E_1 we choose the normalized harmonic basis for \mathcal{F} , e.g., on $M \cong [0, 1]^D$

$$f_k(x) = e^{i2\pi k \cdot x},$$

where $k \in \mathbb{Z}^D$ and $k \cdot x = \sum_{j=1}^D k_j x_j$. This results in the representation of the ergodic quotient map π

$$\begin{aligned} \pi(p) &= (\dots, [P_1 f_k](p), \dots) \\ &= (\dots, \hat{\mu}_p(k), \dots) \end{aligned}$$

that maps the points of the state space p to Fourier coefficients $\hat{\mu}_p(k) := \int_M f_k d\mu_p$ of the associated ergodic measure μ_p . This opens up opportunities to use metrics available on the space of Fourier coefficients, e.g., weighted ℓ^2 metrics, to compare dynamics in different ergodic sets, and generalize the Fomenko-Reeb analysis from Hamiltonian systems to a more general class of systems. Note, however, that when we are interested only in identifying ergodic partitions using discrete topology on ξ , i.e., ask whether $x, y \in M$ are elements of the same ergodic set $S \in \xi$ by checking if $\pi(x) = \pi(y)$, any continuous basis of observables will be sufficient.³⁶ The choice of basis is largely a matter of convenience of constructing other metric structures on ξ , alternative to the discrete metric.

Due to boundedness of Fourier coefficients, the ergodic measures are elements of the Sobolev space $W^{2,-s}$, whose norm can be defined as a weighted Euclidean norm,

$$\|\mu\|_{2,-s}^2 := \sum_{k \in \mathbb{Z}^D} \frac{|\hat{\mu}(k)|^2}{[1 + (2k\pi)^2]^s}, \quad (62)$$

with the index determined as $s = (D + 1)/2$ where $D = \dim M$. We use $\ell^{2,-s}$ as the symbol for the associated Fourier coefficient space with the weighted Euclidean norm. An excellent introduction to the Sobolev space theory is the classic textbook by Adams and Fournier,¹ which contains all the material relevant for this analysis.⁵

To establish a metric structure on the ergodic quotient, we use the $\|\cdot\|_{2,-s}$ -induced metric, and analyze continuity of ergodic quotient maps π in it. The $\ell^{2,-s} \cong W^{2,-s}$ metric structure induces a distance-like function on the state space

$$d_s(p_1, p_2)^2 := \|\mu_{p_1} - \mu_{p_2}\|_{2,-s}^2. \quad (63)$$

This function is not a true distance but a pseudo-distance precisely because the points at zero distance from each other are those that lie in the same ergodic set. Clearly, when the system is uniquely ergodic with respect to μ , the function d_s evaluates to zero for μ -almost any pair of points p_1, p_2 .

Remark 26. The choice of the Sobolev space $W^{2,-s}$ has a justification in comparison of ergodic measures interpreted as measures of residence times of trajectories in measurable sets of a compact state space M .^{29,30} Let $B(x, r)$ denote Euclidean balls in M and $\chi_{x,r}$ their characteristic (indicator) functions. The quantity $\mu_p[B(x, r)]$ is the residence time of trajectory $T^n(p)$ in $B(x, r)$, due to ergodicity of the system $T : S \rightarrow S$, i.e., equality of integrals and time averages (60):

$$\mu_p[B(x, r)] = \int_S \chi_{x,r} d\mu_p = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} \chi_{x,r} \circ T^n(p),$$

where $S \subset M$ is the element of the ergodic partition containing the initial condition p .

The distance between trajectories originating at p_1, p_2 can then be formulated by integrating the difference in residence times over all spherical sets, with R chosen such that $M \subset B(x, R)$:

$$d(p_1, p_2)^2 := \int_0^R \int_M |\mu_{p_1}[B(x, r)] - \mu_{p_2}[B(x, r)]|^2 dx dr. \quad (64)$$

The choice for the index of Sobolev norm $s = (D + 1)/2$ results in the equivalence of norms, i.e., existence of constants $\alpha, \beta > 0$ such that for any $p_1, p_2 \in M$

$$\alpha d(p_1, p_2) \leq d_s = \|\mu_{p_1} - \mu_{p_2}\|_{2,-s} \leq \beta d(p_1, p_2).$$

This argument shows that (64) is again only a pseudo-distance, like d_s . We will return to this interpretation in Sec. V; the formulation (62) is more useful in analysis of the ergodic quotient ξ .

The space $(\xi, \|\cdot\|_{2,-s})$ can be analyzed computationally, allowing for extraction of coherent structures in state spaces of dynamical systems. We are looking to extract connected components of ξ , i.e., curves $\mathcal{C} \subset \ell^{2,-s}$. In our motivational Hamiltonian examples, regions mapping to curves in Fomenko-Reeb graphs (see Fig. 12) contained an onion layering of trajectories, which could be interpreted as regions of

uniform dynamical behavior. We can visualize such regions by coloring all the points $x \in \pi^{-1}(\mathcal{C})$ using the same color.

In a recent paper,⁷ we have presented an algorithm that performed the analysis described above. Using averaging of harmonic observables along trajectories (see Sec. IV C), it evaluates an approximation of the ergodic quotient map π on a set of initial conditions $\{p_n\}_{n=1}^N$ covering the region in state space that is to be analyzed. The connected components are extracted based on pairwise evaluations of the induced Sobolev metric $d(p_1, p_2) \sim \|\mu_{p_1} - \mu_{p_2}\|_{2,-s}$.

The Diffusion Maps algorithm^{9,10} was used to provide a coordinate change for the ergodic quotient ξ . The particular ergodic quotient map π used earlier was constructed using averaged harmonic functions as coordinates. This choice was driven by desire to efficiently evaluate the Sobolev metric as a weighted Euclidean metric. However, the harmonic basis set is chosen without any regard for intrinsic geometry of ξ , and therefore it might not be the most efficient way of analyzing geometry of ξ .

The Diffusion Maps algorithm interprets ξ as a heat-conductive object, and computes the modes of heat spread along it, resulting in a coordinate change $\Psi : \xi \rightarrow \ell^2$. The ℓ^2 distance between points $\Psi(\xi)$ corresponds to the diffusion distance: an intrinsic, coordinate-independent distance along the ergodic quotient. This reflects the fact that the Diffusion Maps algorithm obtains intrinsic geometry of ξ , regardless of the coordinate system that ξ was originally represented in. Components of Ψ are functions $\psi_k : \xi \rightarrow \mathbb{R}$, ordered with $k \in \mathbb{N}$ according to the spatial scales over which they vary. For example, if ξ is a simple line segment, as it is the case for the harmonic oscillator, ψ_k will just be Legendre polynomials: solutions of the heat equation with no-flux boundary condition. For more general ξ , the diffusion modes ψ_k become more complicated, however, they retain the scale-ordering of harmonic functions. The ψ_k are good candidates for the analysis similar to the Reeb graph analysis for manifolds, which was previously noticed in medical image analysis.⁵⁰

As mentioned earlier, a bijection composed with the ergodic quotient map π results in another ergodic quotient map. Therefore, by forming $\chi := \Psi \circ \pi$, with components $\chi_k = \psi_k \circ \pi$ we obtain the ergodic quotient represented in a ‘‘good’’ coordinate set, which can be efficiently truncated to obtain low-dimensional representations of the ergodic quotient. The Euclidean distance over a truncated set of diffusion coordinates approaches the diffusion distance limit exponentially fast, in number of diffusion modes retained.¹⁰ Since Euclidean distance is the most common distance used in applied problems, a host of off-the-shelf algorithms can be used to post-process the ergodic quotient, e.g., a k -means clustering or proximity graph analysis for extraction of connected components.

We use the double-well potential to illustrate this analysis, with results presented in Figure 13(b). The Hamiltonian function for this system, which serves as the basis for the well-known Duffing oscillator, is given by

$$H(q, p) = \frac{1}{2}p^2 - k\left(\frac{1}{2}q^2 - \frac{1}{4}bq^4\right); \quad (65)$$

in this analysis we chose $k = 1, b = 2$. From Figure 13(b) it is evident that the computations using metric $\|\cdot\|_{2,-s}$ retain the

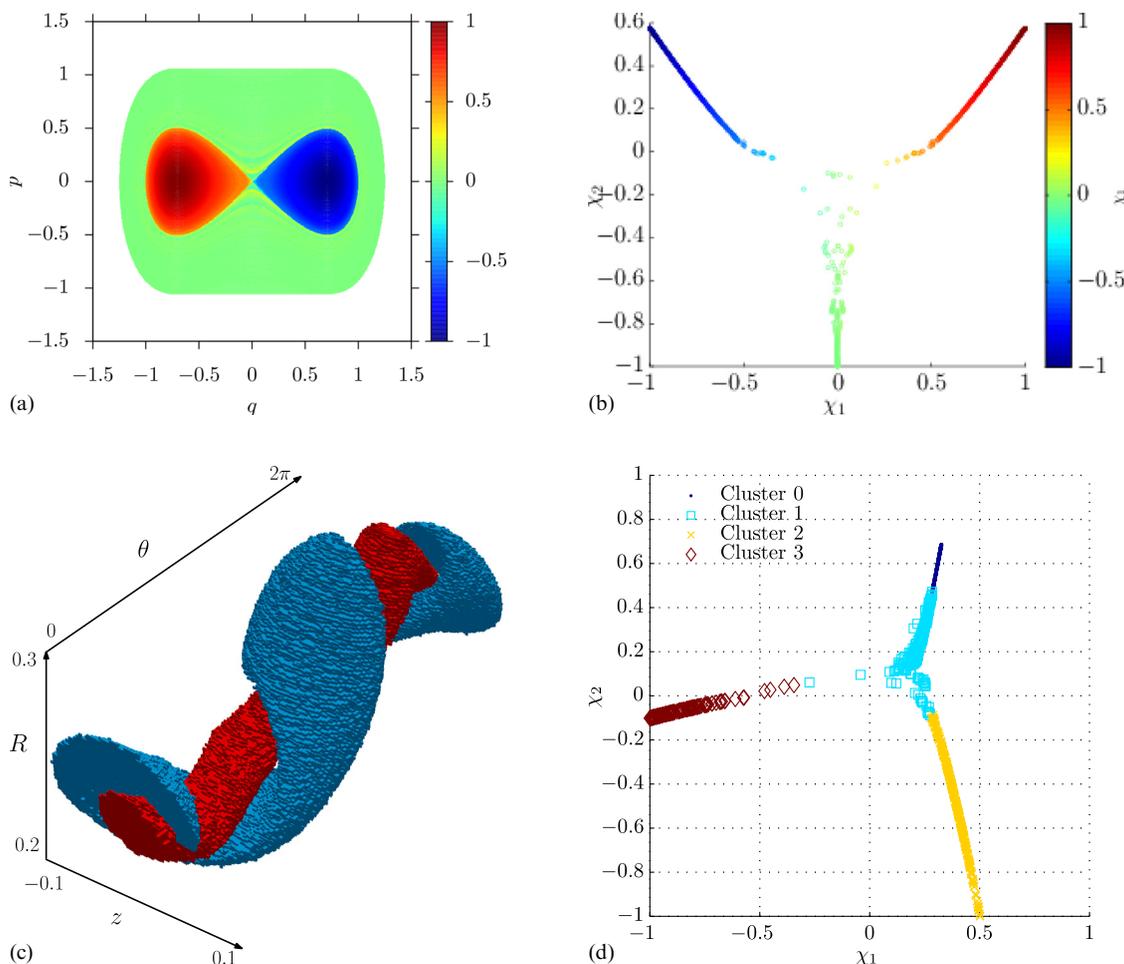


FIG. 13. Geometric analysis of the ergodic quotient for a planar Hamiltonian system with Hamiltonian (Eq. (65)) and a periodically forced 3d fluid-like flow defined by (Eq. (66)). Reprinted with permission from M. Budišić, *Physica D Nonlinear Phenom.* **241**(15), 1255–1269 (2012). Copyright 2012 North-Holland. (a) State space of a double-well potential. Color is the first diffusion coordinate χ_1 , corresponding to color in b. (b) Embedding of ξ for a double-well potential into first two diffusion coordinates. Colors are values of χ_1 for easier comparison with a (Cf. Fig. 12, center). (c) Two vortices in the Poincaré section of the periodic 3D Hill’s vortex flow extracted based on continuous segments in ξ . (d) Embedding of ξ for the periodic 3D Hill’s vortex flow into first two diffusion coordinates.

desired intuition established by Figure 12: the diffusion coordinate χ_2 approximates the energy function of the system, while coordinate χ_1 discriminates between wells of the potential. The gaps in the numerical result are due to finite number of initial conditions x used to evaluate the ergodic quotient map, and irregularities are due to the finite averaging process used to evaluate the projection P_1 ; the Diffusion Maps algorithm can be adaptively tuned to tolerate such errors. When the state space is visualized using pseudocoloring based on diffusion coordinates, the regions of uniform dynamical behavior are clearly distinguished.

The presented process goes beyond Hamiltonian flows: as an illustration we use a periodically forced 3d fluid flow based on the classical Hill’s vortex flow.⁷ It is a solution to an ODE system on $x = (R, z, \theta) \in \mathbb{R}^+ \times \mathbb{R} \times \mathbb{T}$

$$\begin{bmatrix} \dot{R} \\ \dot{z} \\ \dot{\theta} \end{bmatrix} = \underbrace{\begin{bmatrix} 2Rz \\ 1 - 4R - z^2 \\ c/2R \end{bmatrix}}_{A(x,t)} + \epsilon \underbrace{\begin{bmatrix} \sqrt{2R}\sin\theta \\ z(\sqrt{2R})^{-1}\sin\theta \\ 2\cos\theta \end{bmatrix}}_{\sin 2\pi t}, \quad (66)$$

where parameters c and ϵ are swirl and perturbation strengths, respectively. Figure 13(c) shows two vortices that

exist at $c = \epsilon = 0.3495$ colored based on the colors assigned to components of ergodic quotient shown in Figure 13(d).

These two examples demonstrate that functions $\chi_k \in \mathcal{F}$ are invariant eigenfunctions for the Koopman operator, but whose level sets resemble energy functions. In this sense, instead of finding a function that would provide a “good” labeling of ergodic sets, as we did with energy functions in motivational Hamiltonian examples, we constructed a set of such a function from data generated by analyzing a basis for the invariant eigenspace of the Koopman operator. The entire process is made computationally feasible by truncating the set χ_k , therefore discarding the fine-scale features in favor of computability, while retaining coarse-grained features in an organized manner.

C. Infinite-time averages as projections onto eigenspaces

In our presentation of analyses of the state space using eigenfunctions of the Koopman operator, we have assumed we can compute eigenfunctions $\phi^{(\lambda)}$ by projecting an observable f onto the eigenspace at λ using the projection operator P_λ . When $\lambda = e^{i2\pi\omega}$, i.e., for eigenvalues on the unit circle, we can evaluate the associated projection operator P_λ using

infinite-time averages. For this reason, we restrict ourselves to the eigenspaces E_λ for which $|\lambda| = 1$ in this section.

The focus of our interest will be the averages

$$\frac{1}{N} \sum_{n=0}^{N-1} e^{-i2\pi\omega n} f \circ T^n(x), \tag{67}$$

which we would want to extend pointwise into the limit as $N \rightarrow \infty$. To describe the set of observables \mathcal{F} which have well defined limits, we start by assuming that the map T conserves a measure μ on the Borel algebra \mathcal{A} in M . For μ -integrable observables f the finite averages of the form (Eq. (67)) can be extended into $N \rightarrow \infty$ to result in well defined *Fourier averages* $\tilde{f}^{(\omega)}$:

Theorem 27 (Wiener, Wintner). Let $T : M \rightarrow M$ preserve a measure μ on a measurable space M . Then the set of points $\Sigma(f) \subset M$ on which the limit

$$\tilde{f}^{(\omega)}(x) := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} e^{-i2\pi\omega n} f \circ T^n(x) \tag{68}$$

is well defined can be chosen for all $f \in L^1(M, \mu)$ and $\omega \in \mathbb{R}$, independently of ω , and such that $\mu(M/\Sigma) = 0$.

Proof. The original proof by Wiener and Wintner⁵⁹ was found to contain an error. Several correct proofs have been compiled by Assani.³ \square

Moreover, when $L^1(M, \mu)$ contains a dense countable set, the convergence set Σ can be chosen independently of f .³⁶ When frequency $\omega = 0$ is chosen, Fourier averages (Eq. (68)) are referred to as *ergodic averages*,

$$\tilde{f}(x) := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} f \circ T^n(x). \tag{69}$$

As mentioned before, the concept of eigenmeasures $\mu_p^{(\omega)}$ for map T provides means for evaluating the projection operator, to which we now add a trajectory-wise formulation,

$$P_\lambda f(x) = \tilde{f}^{(\omega)}(x) = \int_M f d\mu_x^{(\omega)}.$$

A recent monograph by Wichtrey⁵⁸ provides a detailed analysis of existence of Fourier averages and their applications in linear, nonlinear, and control systems.

While this presentation would perhaps suffice in the abstract, in practical settings we should take additional interest in

1. the influence of choice observable f on the information contained in eigenfunctions $P_\lambda f$,
2. the “size” of the convergence set Σ on which the averages converge, and
3. the rate or error in approximating $\tilde{f}^{(\omega)}$ using finite averages.

The influence of ω on the value of $\tilde{f}^{(\omega)}(x)$ is easily seen by fixing f and x and studying the sequence of (complex) numbers $a_n := f \circ T^n(x)$. The average

$$\sum_{n=0}^{N-1} e^{-i2\pi\omega n} a_n$$

is then just a Discrete Fourier Transform (DFT) of the sequence a_n .

From the basic knowledge of Fourier analysis, the values $\tilde{f}^{(\omega)}$ will be non-zero only for those ω that correspond to temporal modes present in the sequence a_n , which represent the trace of the observable $f \circ T^n(x)$. Whether a_n contains a mode at ω depends not only on trajectory $T^n(x)$, but also on the choice of the observable. This is easiest to see when f is chosen as a constant function (if constant functions are in \mathcal{F}). In that case, $\tilde{f}^{(\omega)} \equiv 0$ for all non-zero ω , and $\tilde{f}^{(0)} \equiv f$, regardless of the underlying dynamics T .

Remark 28. Several concepts in this paper are associated with Joseph Fourier’s name. This is understandable, as his ideas are at the root of all of them, but unfortunate as it might imply direct connections where there are none.

Several recent works^{36,58} use the name *harmonic averages* for $\tilde{f}^{(\omega)}$, in reference to harmonic analysis of dynamical flows. To avoid confusion with harmonic means of numbers, we decided to use *Fourier averages* instead, as do Wiener, Wintner, and Assani. This name is connected with the *temporal* Fourier transforms of sequences, as explained above, since for fixed $x \in M$ and $f \in \mathcal{F}$, the function $\omega \mapsto \tilde{f}^{(\omega)}(x)$ is the Discrete Fourier Transform of the sequence $n \mapsto f \circ T^n(x)$.

Only through ergodicity, i.e., equality of space and time averages $\int_M f d\mu = \tilde{f}$, do we obtain a connection with *spatial* Fourier transform. Choosing $f_k(x)$ as harmonic functions $e^{i2\pi k \cdot x}$, where k is now a wavenumber, we can interpret any sequence $k \mapsto \tilde{f}_k(x)$ in the ergodic quotient ξ as the spatial Fourier transform of the averaging ergodic measure μ_x . We refer to the image set of the spatial Fourier transform as the *spatial Fourier coefficients*.

The eigenquotient maps, i.e., $x \mapsto (\dots, P_\lambda f(x), \dots)$, could be justifiably named Fourier quotient maps, in analogy to ergodic/Fourier average dichotomy. Despite our earlier choices of terminology,⁷ we expect this connection to hold only for eigenvalues $\lambda = e^{i2\pi\omega}$, when $P_\lambda f_k(x) = \tilde{f}^{(\omega)}(x)$. As one could conceivably formulate an eigenquotient map for $|\lambda| \neq 1$, in this paper we chose to use the term eigenquotient instead.

Finally, the harmonic functions $f_k(x) = e^{i2\pi k \cdot x}$, $k \in \mathbb{Z}^d$, i.e., solutions to Laplace equation $\Delta f = 0$ on a torus \mathbb{T}^d , are sometimes called *Fourier functions* or Fourier harmonics. We use these functions as observables, i.e., functions that are acted on by the Koopman operator, to facilitate evaluation of Sobolev norms. Out of all, this connection in name is the least significant of all presented here, nevertheless, we mention it in this remark to clarify our terminology.

Remark 29. Starting from the same observable f , two different eigenfunctions at $\omega = 0$ can be computed as $\phi(x) := \tilde{f}^{(0)}(x)$ and for any $\omega \neq 0$, $\psi := |\tilde{f}^{(\omega)}(x)|$. An example can be found in Figure 11 where the first two images were computed by Fourier averages of $f(x, p) = \sin(\pi x - \pi/4) \cos(6\pi p)$ at $\omega = 0$ and at $\omega = 1/3$, after which the modulus was taken. Both observables are invariant with respect to dynamics.

For a generic observable f , the main difference is that $|\tilde{f}^{(\omega)}(x)|$ is sensitive to ω -periodic dynamics: even though the information about the phase of ω -periodic sets is removed by taking the modulus, $|\tilde{f}^{(\omega)}(x)|$ will be zero on all points that are not in any of ω -resonant chains, which can be

easily explained by DFT interpretation. There is no such restriction on $\tilde{f}^{(0)}$, whose value can vary between different invariant sets (see Fig. 11(a)). Conversely, $|\tilde{f}^{(1/3)}|$ is zero everywhere except on the period-3 island, where trajectories contain $\omega = 1/3$ frequency components (Fig. 11(b)).

The ω values corresponding to the $e^{i2\pi\omega}$ that are eigenvalues of the Koopman operator will result in non-zero eigenfunction $f^{(\omega)}$ for at least some observable f . The ω values that are not frequencies of any of the eigenvalues will result in $\tilde{f}^{(\omega)} \equiv 0$. The frequency $\omega = 0$, however, corresponds to eigenvalue $\lambda = 1$ of U , which is always in the spectrum of the Koopman operator. This frequency is of a particular interest, as the eigenfunctions $\tilde{f}^{(0)} \equiv \tilde{f}$ are invariant functions for the system.

To illustrate the difference in eigenfunctions obtained by starting from different observables, we again used the Chirikov Standard Map (61) at $\epsilon = 0.15$. The averages were computed for $\omega = 0$, which projects observables to the invariant eigenspace of U . As Figure 14 shows, the detail that level sets of averaged observables reveal about the state space is highly dependent on the starting observable, however, when a “good” observable is chosen, it can reveal a lot about the state space. The lack of intuition about how to select such a “good” starting observable led to development of the ergodic quotient analysis, presented in Sec. IV B. The diffusion coordinates χ_k used therein can be interpreted as constructed observables that reveal detailed information about the state space.

The convergence set Σ was established to be of full measure μ by Theorem 27. The measure μ is a measure preserved by the system, which is also used to define the space of integrable observables $L^1(M, \mu)$ for which the Wiener-Wintner theorem holds. Depending on the analyzed system, the existence of an invariant measure can be inferred through different arguments. For example, a system $\dot{x} = F(x)$ governed by divergence-free vector field $\nabla \cdot F \equiv 0$ conserves

the volume-measure on the state space. In dissipative dynamics, the systems may conserve Sinai-Ruelle-Bowen measures on chaotic sets.⁶¹ Finally, on a compact state space M the only necessary assumption for existence of an invariant μ is continuity of T , by the theorem of Krylov-Bogolyubov.²¹

The Wiener-Wintner theorem holds for *any* invariant measure μ . Therefore, by choosing the measure used to formulate the space of observables $\mathcal{F} = L^1(M, \mu)$, we influence the amount of information we can collect about the system by evaluating Fourier averages and the size of the set Σ . In applied contexts, we would often want to include as many open sets as possible in the support of measure μ chosen: for volume-preserving systems, volume of the state space is often a good choice. Caution is still needed, as there could be a μ -zero, yet dense, non-convergence set Σ^c . Nevertheless, in our experience, simulations of dynamical systems that model physical phenomena do not contain such extreme pathological cases.

For dissipative systems, one often requires a that the set of Birkhoff-regular initial conditions,¹³ for which Fourier averages of continuous observables converge, is of positive volume, i.e., that the system preserves a *physical measure*⁶¹ supported on the attractor. In those cases, computing quotient maps using Fourier averages will identify basins of attractions of the attractors, instead of attractors themselves, since the points in the basin will have the same averages as the points on the attractor.

A well-known example¹⁷ of a dissipative system for which trajectory averages do not converge for almost all open sets of initial conditions contains an attractive heteroclinic cycle (see Fig. 15). Ergodic measures are δ -measures supported on each of the fixed points, consequently, any invariant measure will be singular with respect to Lebesgue. Trajectories that approach heteroclinics spend longer and longer times along each of the exterior fixed points, possibly in such a manner that the finite time averages do not

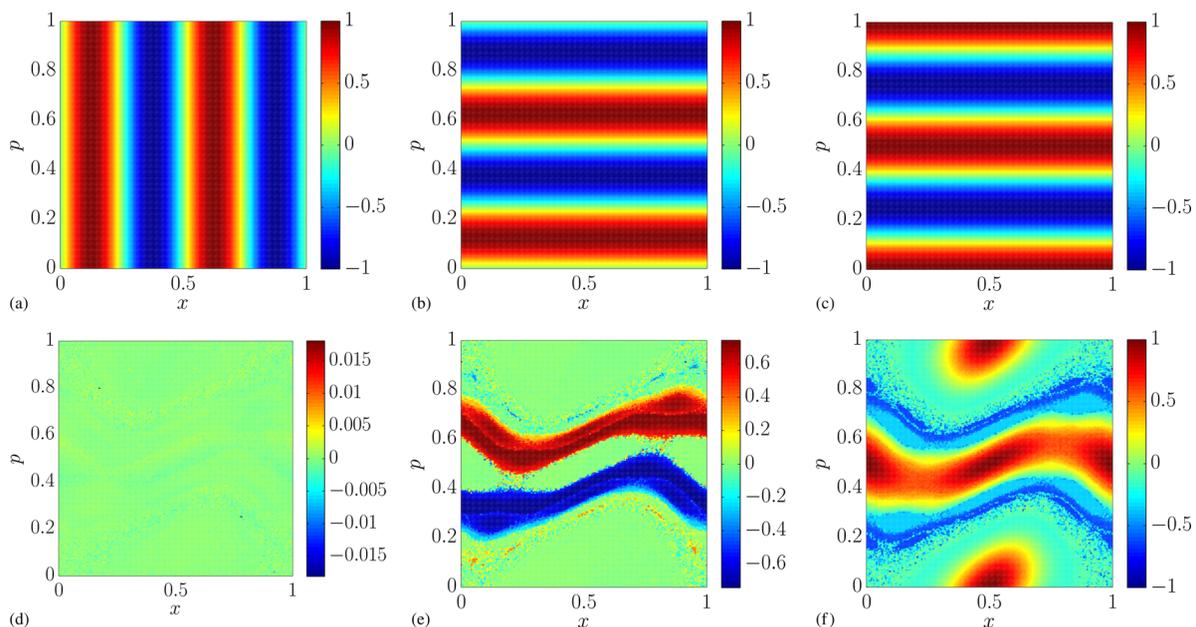


FIG. 14. Three different observables f_j (top row) and associated invariant functions \tilde{f}_j (bottom row), obtained by ergodic averages along trajectories of the Chirikov Standard Map (IV A) for $\epsilon = 0.15$. The pseudocolor is the value of each function. (a) Level sets of $f_1(x, p) = \sin(4\pi x)$. (b) Level sets of $f_2(x, p) = \sin(4\pi p)$. (c) Level sets of $f_3(x, p) = \cos(4\pi p)$. (d) Level sets of $\tilde{f}_1(x, p)$ (e) Level sets of $\tilde{f}_2(x, p)$. (f) Level sets of $\tilde{f}_3(x, p)$.

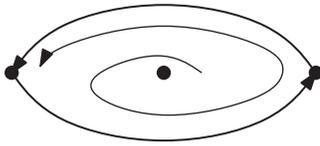


FIG. 15. Attracting heteroclinic cycles may prevent convergence of ergodic averages.

converge. Notice that, while the conditions of Theorem 27 are formally satisfied, they are almost vacuous for practical purposes, as only the subsets of heteroclinic orbits have well-defined averages, yet the convergence set is still of full measure, due to singular nature of the measure conserved.

From the authors’ experience in applying the averaging technique to dynamical systems that model physical phenomena for practical analysis, finding an appropriate convergence set Σ is rarely a problem. A more practically significant issue comes from the different rates of convergence of finite to infinite averages for points *within* Σ .

The rate of convergence of finite limits

$$A_N f(x) := \frac{1}{N} \sum_{n=0}^{N-1} f \circ T^n(x),$$

and their Fourier counterparts (67) on the set Σ is not uniform. It is a classical result that the rates of convergence over periods of length N can range from exponentially fast, i.e., $\mathcal{O}(e^{-\lambda N})$ for $\lambda > 0$, for trajectories approaching exponentially stable fixed points, to algebraic convergence $N^{-\alpha}$, where $\alpha > 0$ is arbitrarily small, see, e.g., Petersen, Sec. III A.⁴⁰ From a practical perspective, in a lot of cases the situation does not look so bleak. Trajectories on periodic orbits and in strongly mixing regions achieve the rates of convergence of $\mathcal{O}(N^{-1})$ and $\mathcal{O}(N^{-1/2})$, respectively. The slopes $0 < \alpha < 1/2$ are to be expected near homoclinic and

heteroclinic orbits, especially if such orbits are embedded within zones of intermittency, where trajectories get entrained around marginally stable fixed points for long times, before eventually moving away from them.¹² Such zones appear, for example, in perturbed hamiltonian and volume-preserving systems.³⁹ Studies of volume-preserving systems, in theoretical⁵⁷ and computational⁷ contexts, have shown that such regions are small in area.

To illustrate, we plotted convergence errors for the Chirikov Standard Map (61), simulated for $\epsilon = 0.18$ where the state space contains both regular and mixing regions. This brief analysis is similar to those in literature.^{7,25,26} As an indicator of convergence speed, we compared the ergodic average, i.e., Fourier average with $\omega = 0$, after $N_\infty = 10^6$ iterates with the average after (significantly) shorter number $N \ll N_\infty$ of iterates. To avoid making conclusions based on a single observable, a truncated set of Fourier harmonics

$$f_k(x, p) = \frac{1}{2\pi} \exp[i2\pi(k_x x + k_p p)],$$

for $k = (k_x, k_p) \in [-5, 5]^2 \subset \mathbb{Z}^2$, was used; this is the type of set used to practically approximate the quotient maps described in Sec. IV B and for each trajectory the largest absolute difference in averages over that set was taken as an indication of the convergence error.

Figure 16(a) shows the error in averages after a fixed time $N = 10^5$ iterates, while Figures 16(b)–16(d) show the time required for the error to stay within 10^{-3} for 100 iterates. It is clear that the speed of convergence varies across the state space, with convergence in chaotic regions slower than in regular regions. We proposed⁷ that the simulated time is varied for individual trajectory based on the relative convergence error, which resulted in efficient simulation runs with only a small subset of initial conditions, in regions of intermittency, requiring long simulation times.

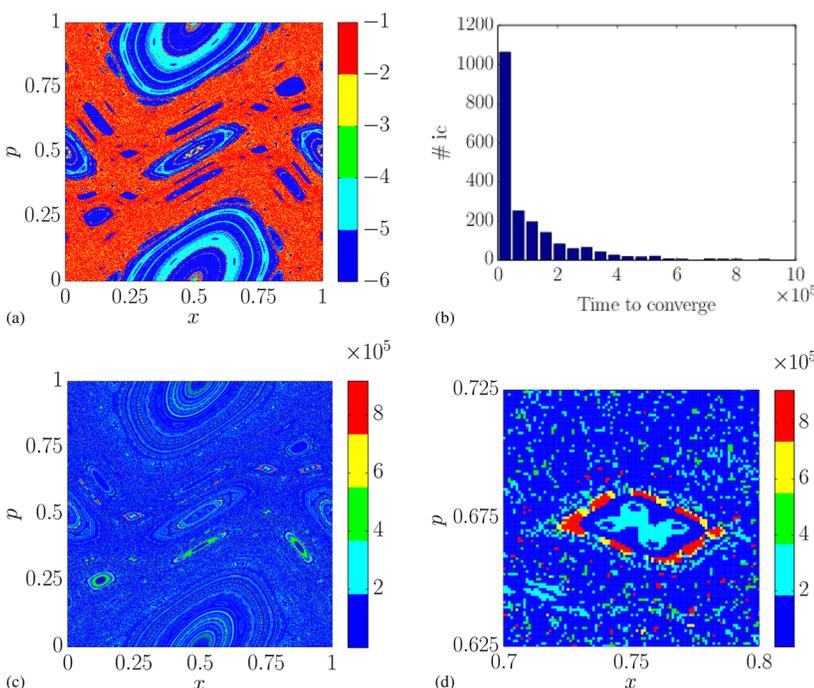


FIG. 16. Convergence of ergodic averages. The difference in averages of Fourier harmonics for trajectory starting at $x \in M$ is computed as $\max_k |A_N f_k(x) - A_{N_\infty} f_k(x)|$ over $k \in [-5, 5]^2 \subset \mathbb{Z}^2$. Step $N_\infty = 10^6$ was taken as “true” infinity. A total of 2025 initial conditions, seeded from a uniform rectangular grid, was used. Colors were interpolated from the first 500 iterates of each trajectory. (a) Difference in averages after $N=10^5$ and $N_\infty=10^6$ iterates (\log_{10} -scale). The bottom of the color scale includes all the values below -6 (b) Distribution of number of initial conditions over iterations required for the difference in averages compared to $N_\infty=10^6$ iterates to reach 10^{-3} . (c) Number of steps required for the difference in averages compared to 10^6 iterates to reach 10^{-3} . (d) Number of steps required for the difference in averages compared to 10^6 iterates to reach 10^{-3} (enlarged).

This section dealt with the problem of extending finite-time averages into the infinite limit, to be able to approximate the limiting measures of empirical distributions. It might be surprising that even explicitly finite time averages find their use in practical applications. In Sec. V, we demonstrate how to formulate continuous indicators of mixing and ergodicity, and use them in design of feedback control for technical systems.

V. CONTINUOUS INDICATORS OF ERGODICITY AND MIXING

When dynamical systems are analyzed as measure-preserving transformations, *ergodicity* and *mixing property* are among the first concepts discussed in introductory textbooks.^{24,27} Let measure μ be preserved by the system. In plain language, ergodicity with respect to μ means that sets left invariant by the transformation/flow are either full measure μ , or μ -negligible. The mixing property, which implies ergodicity, means that any set of positive μ -measure will be distributed by the flow according to the mixing measure μ .

While the standard introductions to these properties⁴⁰ often include equivalence theorems that provide alternative formulations of ergodicity and mixing, cf. Eqs. (60) and (71), all the definitions treat them as binary indicators: either the system is ergodic/mixing, or it is not. From the perspective of applied dynamical systems, especially in the context of design of dynamics, having a binary indicator of a desired property is insufficient. Designers prefer to work with continuous indicators, e.g., an ergodicity indicator that takes values in $[0, 1]$ where the extremum 0 would imply classical ergodicity, and the higher values would indicate how far, in some sense, the system is from being ergodic. The continuous indicators of ergodicity and mixing can be constructed from averages of functions along trajectories.^{29,48} The averaging process corresponds to an invariant measure, which is then compared to the *a priori* measure μ to infer how close the system is to being μ -ergodic or μ -mixing.

Designing systems to be ergodic or mixing has a number of practical applications. Consider the case of micro-mixers, devices whose task is to mix two fluids, reactants, such that when the reaction is initiated, there are no “pockets” or unused reactants, and the reaction occurs uniformly in the vessel. At usual macro-scales, the fluids are easily mixed just by shaking them; on micro- and nano-scales this is not possible. Instead, micro-mixers use dynamical, advective transport, to reach the mixed state. An example where ergodicity is desirable comes from search-and-rescue missions, where helicopters or airplanes are used to scan a large area for survivors of airplane crashes and capsized boats. It is not trivial to design a flight path that ensures that the entire area is covered and that particular zones are searched more often or more thoroughly: such task can be phrased as design of an ergodic flight path. Both of these problems can be addressed by casting them into a framework of dynamical systems and requiring that the trajectories are mixing or ergodic.

While trajectory averages were treated as a technique for computation of projections onto eigenspaces of the Koopman operator, which required their extension into

infinity, in this section we will show what can be extracted from finite-time averages (Eq. (67)), with $\omega = 0$, i.e., the finite-time version of ergodic averages. The average of a bounded observable $f : M \rightarrow \mathbb{C}$ over a finite-time trajectory segment, $f \mapsto \frac{1}{N} \sum_{n=0}^N f \circ T^n(x)$, is a linear continuous functional. By the Riesz representation theorem, the finite-average functionals are represented by *empirical measures*, whose distributions, formally defined as

$$c_{x,N}(p) := \frac{1}{N} \sum_{n=0}^{N-1} \delta[p - T^n(x)], \tag{70}$$

can be thought of as strings (or ribbons) of δ -distributions, supported along the orbits (see Fig. 17).

The consequence of defining the empirical distribution using the Riesz representation theorem is the well-known equality between spatial and time averages. With a bit of breadth in notation, the expression $\langle f, g \rangle = \int_M f(p)g(p)dp$ can be used to couple functions f and measures $g(p)dp$. In the case of measures with density functions, i.e., absolutely continuous measures, both f and g can be taken as functions; when $g(p)dp = d\mu(p)$ corresponds to a more general measure μ , f is taken as a test function, and g is in the class of distributions, or generalized functions, e.g., for a point mass μ , g is a Dirichlet δ distribution. Coupling the empirical measure (Eq. (70)) with a test function f yields

$$\begin{aligned} \langle c_{x,N}, f \rangle &= \int_M \left\{ \frac{1}{N} \sum_{n=0}^{N-1} \delta[p - T^n(x)] \right\} f(p) dp \\ &= \frac{1}{N} \sum_{n=0}^{N-1} \int_M \delta[p - T^n(x)] f(p) dp \\ &= \frac{1}{N} \sum_{n=0}^{N-1} f[T^n(x)]. \end{aligned}$$

If, as an observable, we take the characteristic function of a measurable set A , $r_{x,N}(A) = \langle c_{x,N}, \chi_A \rangle$ is the residence time: the fraction of time $[0, N]$ that trajectory $T^n(x)$ spends inside set A . Therefore, the empirical measures $\langle c_{x,N}, \cdot \rangle$ induced by averaging functionals are probability measures, for any initial condition $x \in M$ and $N \geq 0$. For any finite time, empirical distributions are not absolutely continuous with respect to the Lebesgue measure. In the limit $N \rightarrow \infty$, however, they converge in weak-* topology to ergodic measures;⁶¹ their limits might be point masses, on equilibria, singular

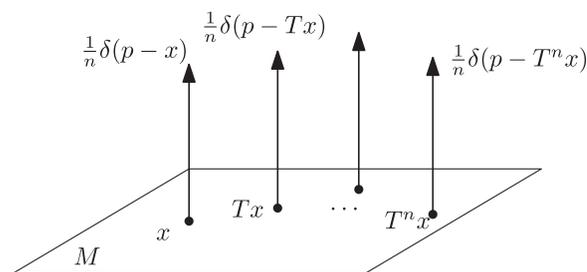


FIG. 17. For finite time N , averaging functionals can be represented by empirical probability measures, whose distributions $c_{x,N}$ can be thought of as sequences of δ distributions supported on trajectories.

measures, e.g., on periodic orbits, or absolutely continuous measures, e.g., in mixing regions.

Ergodicity of the flow T^n with respect to a flow-invariant measure μ can be stated as the condition that for $[\mu]$ -a.e. $x \in M$,

$$\lim_{N \rightarrow \infty} \langle c_{x,N}, \chi_A \rangle = \mu(A), \quad (71)$$

for all measurable sets A . The state-space points $x \in M$ whose trajectories satisfy the ergodicity condition are called μ -generic. The first step towards a continuous indicator of ergodicity was presented by Scott *et al.*,⁴⁸ who used a Haar wavelet basis as observables: these wavelets directly relate to indicator functions of a basis for open sets on a rectangular state space. Since wavelet bases are parametrized by scale at which sets are resolved, the constructed ergodic indicator, termed *ergodicity defect*, acted as a coarse-grained version of ergodicity, reflecting an engineering perspective: instruments have finite precisions, therefore, if the equalities (Eq. (71)) are satisfied for all sets coarser than precision of the instrument, the system can be thought of as ergodic for the particular application.

Based on a similar idea, Mathew and Mezić²⁹ constructed a different continuous ergodicity indicator. For declaring ergodicity, it is sufficient to verify the condition (Eq. (71)) on a basis for the Borel algebra, e.g., Euclidean balls $B(x, r) = \{p \in M : \|p - x\|_2 < r\}$, whose indicator functions we label $\chi_{(p,r)}$. Integrating the deviations between trajectory averages and measures of sets, we obtain the empirical ergodicity $E_x(n)$

$$E_x(n)^2 := \int_0^R \int_M |\langle c_{x,n}, \chi_{(p,r)} \rangle - \mu[B(p, r)]|^2 dp dr, \quad (72)$$

where R is selected such that the largest ball includes the entire state space M (cf. pseudodistance (Eq. (64))). It is almost immediate that the ergodicity is equivalent to $\lim_{n \rightarrow \infty} E_x(n) = 0$ if x is selected as a μ -generic point. Strictly speaking, E_x compares the $N \rightarrow \infty$ limit of empirical measures to the prior μ ; to assure that the same behavior occurs almost everywhere, E_x can be integrated along the state space.

The practical value of $E_x(n)$ is in how its constituents are computed: $\langle c_{x,n}, \chi_{(p,r)} \rangle$ is computed as a time-average of an indicator function, which is very easily computed as the system is being simulated from an ODE. When μ is the volume measure, quantities $\mu[B(x, r)]$ are just volumes of Euclidean balls, and are computed using well-known formulas. Consequently, evaluation of $E_x(n)$ is simple, as it requires only evaluation of finite-time averages.

On the other hand, μ can be a more detailed measure, as in the mentioned case of probabilities of target detection.^{28,29} In those cases, evaluating (Eq. (72)) might be a more complicated effort, requiring careful spatial gridding to control errors in the integral. Instead of measures of spherical sets, the metric $E_x(n)$ can be expressed using Fourier coefficients of μ , facilitating the control of spatial scale resolution. The integrand in Eq. (72) can be interpreted as a difference between generalized expansion coefficients, where the basis

set is the set of characteristic functions, whose supports are a base for Borel sets

$$\mu[B(p, r)] = \langle \partial\mu, \chi_{(p,r)} \rangle.$$

The $\partial\mu$ stands for the (formal) density of the prior measure μ .

If the basis $\chi_{(p,r)}$ is replaced by a harmonic basis f_k , then $\hat{\mu}(k) := \langle \partial\mu, f_k \rangle$ are just spatial Fourier coefficients, easily computed to very high orders by fast Fourier transform if we know the density $\partial\mu$ explicitly. At the same time, practicality of evaluation of $\langle c_{x,n}, f_k \rangle$ is not sacrificed. The resulting empirical ergodicity is given by a metric induced by the negative-index Sobolev norm $\|\cdot\|_{2,-s}$ on the space of distributions $W^{2,-s}$ (cf. Eq. (62)), if both $\partial\mu$ and $c_{x,n}$ are in it.

$$\|c_{x,n} - \partial\mu\|_{2,-s} = \sum_{k \in \mathbb{Z}^D} \frac{|\langle c_{x,n}, f_k \rangle - \langle \partial\mu, f_k \rangle|^2}{[1 + (2\pi\|k\|_2)^2]^s}, \quad (73)$$

where the state space is, for simplicity, $M \simeq \mathbb{T}^D$, $k \in \mathbb{Z}^D$. Wavevectors $k \in \mathbb{Z}^D$ parametrize harmonic functions $f_k(x) = (2\pi)^{D/2} e^{i2\pi k \cdot x}$.

When the order of the Sobolev space s is chosen as $s = (D+1)/2$, $\|c_{x,n} - \partial\mu\|_{2,-s}$ and $E_x(n)$ are equivalent:²⁹ there exist constants $\alpha > 0$ and $\beta > 0$ such that

$$\alpha \|c_{x,n} - \partial\mu\|_{2,-s} \leq E_x(n) \leq \beta \|c_{x,n} - \partial\mu\|_{2,-s},$$

at all n . As a consequence, decay of $\|c_{x,n} - \partial\mu\|_{2,-s}$ can be used as the proxy for computing decay of $E_x(n)$ to detect ergodicity, since $\|c_{x,n} - \partial\mu\|_{2,-s}$ is easier to numerically evaluate for most measures μ .

A practical application of $E_x(n)$ can be seen on a model search-and-rescue problem, where an UAV, e.g., a small helicopter, would search an area containing a target whose position is estimated by a probability density.²⁹ The search path for the vehicle can be planned using a dynamical system: the searcher is treated as a passive particle, possibly with a non-zero inertia, in a fluid-like flow. The probability density is a product of Gaussians, modeling *a priori* estimates of the target, and discontinuous indicator functions, modeling foliage on the ground where the UAV has no visibility. The goal is to design the dynamics of the flow, such that the trajectory traced out by a particle avoids foliage and explores the feasible area according to probability of finding the target.

The devised control algorithm used $E_x(t)$ (in continuous-time setting) as the optimization function in a Hamilton-Jacobi-Bellman framework. Instead of optimizing for decay of $E_x(t)$ over a finite horizon $t \in [0, T]$, a greedy approach is chosen, where the optimization horizon is shrunk to the instance by taking $T \rightarrow 0$, resulting in a closed-form expression feedback law which relies on finite-time averages of harmonic observables, computed along the path that the searcher traveled. As a result, trajectories designed for a team of UAVs searched the area efficiently, with minimal crossings over the zones with high foliage (Fig. 18), with further details presented in Refs. 28 and 29.

In addition to quantifying ergodicity, the $W^{2,-s}$ norm can be used to quantify mixing with respect to a target density $\partial\mu$,

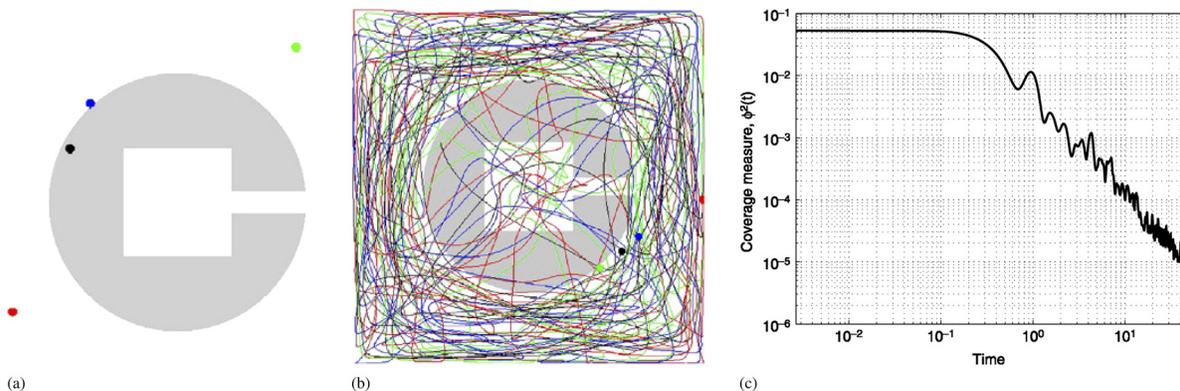


FIG. 18. Searching for a target by optimizing the decay of empirical ergodicity $E_x(n)$, through computation of $\|c_{x,n} - \partial\mu\|_{2,-s}$ for several trajectories. Reprinted with permission from G. Mathew, *Physica D Nonlinear Phenom.* **240**(4–5), 432–442 (2011). Copyright 2011 North-Holland. (a) Target probability distribution and initial searcher positions. Probability density is positive and constant on white region, zero on gray. (b) Paths of searchers, showing the searchers sampling the area with positive density. (c) Decay of $\|c_{x,n} - \partial\mu\|_{2,-s}^2$ (labeled by Φ in the original paper).

in which context it was originally developed.^{30,31} In the final design of a micro-mixer, an initial concentration of fluid reactants ρ was advected by a dynamical system, whose effect was captured by the Perron-Frobenius operator $P^n\rho$. To quantify mixing, the metric $M(n) := \|P^n\rho - \partial\mu\|_{2,-s}$ was computed with the order $s = 1/2$. By optimizing the dynamics of the fluid flow, a quick mixing of reactants was successfully achieved (Fig. 19), both in numerical and experimental settings.

The main difference between $M(n)$ and $E_x(n)$ is that the mixing metric compares *instantaneous* advected density $P^n\rho$ to the prior μ , which, in the dual Koopman framework, is equivalent to using $U^n f$. Recall that the ergodicity metric used the temporal averages $\frac{1}{N} \sum_{n=0}^{N-1} [U^n f]$ (cf. Eq. (70)). The other difference between mixing and ergodicity indicators can be interpreted as the requirement on the scales at which measure μ is sampled: the ergodicity indicator $E_x(n)$ gives more weight to larger spherical sets, while in $M(n)$, differences over all sets are weighted the same, regardless of the set

size, due to rescaling of the measure $dxdr$ by the measure of the spherical set to $dxdr/\mu(x,r)$. Therefore, those trajectories that lead to uniform decay of $E_x(n)$ will first appear well distributed in μ over larger sets, and only then will they sample μ on smaller scales. Conversely, $M(n)$ does not make such a distinction, requiring that the trajectories distribute on all scales equally fast.

In practice, the choice between $E_x(n)$ and $M(n)$ is driven by the application: for mixing of fluids, it was important that two reactants were in contact at the instant of reaction initiation, e.g., initiation of burning in the chamber. In the search-and-rescue application, instantaneous behavior was not of the essence, as the goal was to find the target over the entire course of the search mission, not necessarily to have the same (high) probability to locate the target at any particular time instance. In future applications, a similar analysis of the problem would decide which property, ergodicity $E_x(n)$ or mixing $M(n)$, would be a more appropriate design criterion.

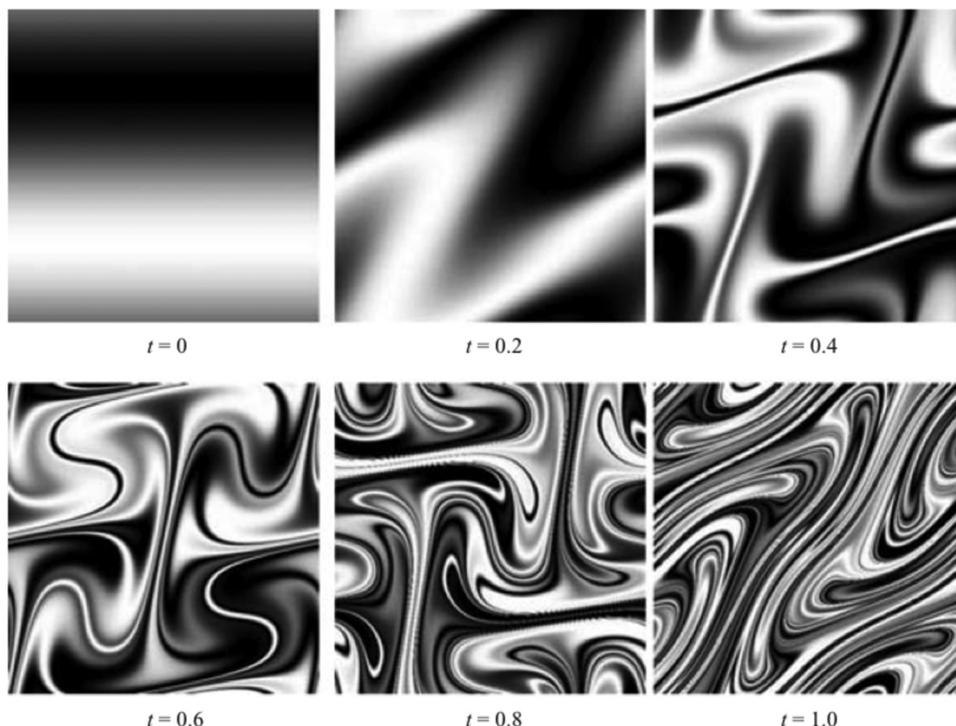


FIG. 19. Simulation of mixing of reactants in a micro-chamber using dynamics designed by optimizing a $W^{-s,2}$ distance between advected density and uniform target density. Reprinted with permission from G. Mathew *et al.*, *J. Fluid Mech.* **580**, 261 (2007). Copyright 2007 Cambridge University Press.

The indicators for ergodicity and mixing presented in this section are based on representations of ergodic measures in the space of Fourier coefficients. They, therefore, represent ergodic measures as accurately as it is possible through their Fourier coefficients: numerically, approximations will converge quicker the smoother the ergodic measure studied. However, it is well known that the shape of the ergodic measure's density does not determine uniquely the dynamical system. In this sense, the indicators are not intrinsic to the dynamical system studied. For example, if one would use a unique ergodic measure μ for a map T_1 and seek to drive the ergodic measure of another system T_2 to match μ , there would be no guarantee that trajectories of T_1 and T_2 would be conjugate, without further restrictions on properties of maps T_1 and T_2 .

VI. CONCLUSIONS

We have reviewed the theoretical aspects and applications of the spectral theory of the Koopman operator in dynamical systems. The use of these concepts holds promise to provide a theory that extends and complements tools from geometrical dynamical systems theory that enabled so much development in science and technology over the last century. The presented material described three parallel branches of Koopman operator analysis

1. Koopman mode analysis,
2. eigenquotient analysis, and
3. indicators of ergodicity and mixing.

The Koopman modes generalize the notion of linear eigenmodes, known from, e.g., linear mechanical vibration theory, to the nonlinear context, without linearizing the dynamics first. It is interesting that projections onto eigenspaces of the Koopman operator are achieved via an extension of Laplace and Fourier analysis, and that such extension works in the nonlinear case. The Koopman mode analysis has proved especially useful in analysis of dynamics in infinite-dimensional state spaces, which were observed using a high-dimensional measurements, e.g., thermal dynamics of a building system, with a distributed measurement of temperatures or fluids experiments.

The eigenquotient analysis generalizes the analysis of smooth integrals of motion, which is central in theoretical mechanics, to systems which do not have any smooth invariants. The non-smooth eigenfunctions are constructed using infinite-time averaging along trajectories of the dynamical system. The eigenfunctions are used to construct a geometry for families of invariant sets which can be used to extract invariant regions that locally resemble phase portraits of integrable Hamiltonian systems. The same analysis can not only be applied to the study of invariant sets, but also periodic and wandering sets.

Finally, we showed how the trajectory averages can be used even when only finite-time data is available. The values of finite-time averages of functions along trajectories can be used to evaluate quantitative indicators of how closely the system is to accurately sampling a prior measure. Such indicators can be used as optimization criteria in

a feedback algorithm for trajectory planning of unmanned aerial vehicles.

The theory reviewed here was introduced and developed mostly for measure-preserving deterministic systems, including dynamics on the attractor of a dissipative system. Much more work remains to be done in the case of dissipative systems and extensions to non-smooth and hybrid (deterministic/stochastic) case.

ACKNOWLEDGMENTS

The authors would like to thank the Editor, Dr. Holmes, and anonymous reviewers for their extensive comments that helped improve the manuscript. The authors were funded by the following Grant Nos. ARO W911NF-11-1-0511, AFOSR FA9550-10-1-0143, AFOSR FA9550-09-1-014, ONR MURI N00014-11-1-0087, ONR N00014-07-1-0587, and ONR N00014-10-1-0611.

- ¹R. A. Adams and J. J. F. Fournier, "Sobolev spaces," *Pure and Applied Mathematics (Amsterdam)*, 2nd ed. (Elsevier and Academic, Amsterdam, 2003), Vol. 140.
- ²W. E. Arnoldi, "The principle of minimized iteration in the solution of the matrix eigenvalue problem," *Q. Appl. Math.* **9**(1), 17–29 (1951).
- ³I. Assani, *Wiener Wintner Ergodic Theorems* (World Scientific Publishing Co., Inc., River Edge, NJ, 2003).
- ⁴A. V. Bolsinov and A. T. Fomenko, *Integrable Hamiltonian Systems* (CRC, 2004).
- ⁵M. Budišić, "Ergodic quotients in analysis of dynamical systems," Ph.D. dissertation (UC Santa Barbara, 2012).
- ⁶M. Budišić and I. Mezić, "An approximate parametrization of the ergodic partition using time averaged observables," in *Proceedings of the 48th IEEE Conference on Decision and Control, held jointly with the 28th Chinese Control Conference*, CDC/CCC (2009), pp. 3162–3168.
- ⁷M. Budišić and I. Mezić, "Geometry of the ergodic quotient reveals coherent structures in flows," *Physica D* **241**(15), 1255–1269 (2012).
- ⁸K. K. Chen, J. H. Tu, and C. W. Rowley, "Variants of dynamic mode decomposition: Boundary condition, Koopman, and Fourier analyses," *J. Nonlinear Sci.* **22**(6), 887–915 (2012).
- ⁹R. Coifman and S. Lafon, "Diffusion maps," *Appl. Comput. Harmon. Anal.* **21**, 5 (2006).
- ¹⁰R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, "Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps," *Proc. Natl. Acad. Sci. U.S.A.* **102**(21), 7426–7431 (2005).
- ¹¹D. B. Crawley, C. O. Pedersen, L. K. Lawrie, and F. C. Winkelmann, "Energy plus: Energy simulation program," *ASHRAE J.* **42**, 49–56 (2000), <http://bookstore.ashrae.biz/journal/journalissue.php?pubmonth=April&pubyear=2000>.
- ¹²P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay, *Chaos: Classical and Quantum*, stable 13th ed. (Niels Bohr Institute, Copenhagen, 2010).
- ¹³D. D'Alessandro, I. Mezić, and M. Dahleh, "Statistical properties of controlled fluid flows with applications to control of mixing," *Syst. Control Lett.* **45**(4), 249–256 (2002).
- ¹⁴M. Dellnitz and O. Junge, "On the approximation of complicated dynamical behaviour," *SIAM J. Numer. Anal.* **36**(2), 491–515 (1999).
- ¹⁵V. Dragović and M. Radnović, "Bifurcations of Liouville tori in elliptical billiards," *Regular Chaotic Dyn. Int. Sci. J.* **14**(4-5), 479–494 (2009).
- ¹⁶B. Eisenhower, T. Maile, M. Fischer, and I. Mezić, "Decomposing building system data for model validation and analysis using the koopman operator," in *SimBuild 2010—Fourth National Conference of IBPSA, USA*, August 2010.
- ¹⁷A. Gounersdorfer, "Time averages for heteroclinic attractors," *SIAM J. Appl. Math.* **52**(5), 1476–1489 (1992).
- ¹⁸M. Georgescu, B. Eisenhower, and I. Mezić, "Creating zoning approximations to building energy models using the Koopman operator," in *SimBuild 2012—Fifth National Conference of IBPSA, USA*, August 2012.

- ¹⁹M. Germano, U. Piomelli, P. Moin, and W. H. Cabot, "A dynamic subgrid-scale eddy viscosity model," *Phys. Fluids A: Fluid Dyn.* **3**, 6 (1991).
- ²⁰C. K. R. T. Jones, "Whither applied nonlinear dynamics?," *Mathematics Unlimited—2001 and Beyond*, edited by B. Engquist and W. Schmid (Springer, Berlin, 2001), pp. 631–645.
- ²¹A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems, Encyclopedia of Mathematics and its Applications Vol. 54* (Cambridge University Press, Cambridge, 1995).
- ²²B. O. Koopman, "Hamiltonian systems and transformations in Hilbert space," *Proc. Natl. Acad. Sci. U.S.A.* **17**(5), 315–318 (1931).
- ²³U. Krengel, "Ergodic theorems," in *de Gruyter Studies in Mathematics* (Walter de Gruyter and Co., Berlin, 1985), Vol. 6.
- ²⁴A. Lasota and M. C. Mackey, "Chaos, fractals, and noise: Stochastic aspects of dynamics," in *Applied Mathematical Sciences*, 2nd ed. (Springer-Verlag, New York, 1994), Vol. 97.
- ²⁵Z. Levnajić and I. Mezić, "Ergodic theory and visualization II: Visualization of resonances and periodic sets," e-print arXiv:1206.3164.
- ²⁶Z. Levnajić and I. Mezić, "Ergodic theory and visualization. I. Mesochronic plots for visualization of ergodic partition and invariant sets," *Chaos* **20**(3), 033114 (2010).
- ²⁷R. Mañé, "Ergodic theory and differentiable dynamics," in *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]* (Springer-Verlag, Berlin, 1987), Vol. 8.
- ²⁸G. Mathew and I. Mezić, "Spectral multiscale coverage: A uniform coverage algorithm for mobile sensor networks," in *Proceedings of the 48th IEEE Conference on Decision and Control, 2009 held jointly with the 28th CCC/CCC* (2009), pp. 7872–7877.
- ²⁹G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D* **240**(4–5), 432–442 (2011).
- ³⁰G. Mathew, I. Mezić, and L. Petzold, "A multiscale measure for mixing," *Physica D* **211**(1–2), 23–46 (2005).
- ³¹G. Mathew, I. Mezić, S. Grivopoulos, U. Vaidya, and L. Petzold, "Optimal control of mixing in Stokes fluid flows," *J. Fluid Mech.* **580**, 261 (2007).
- ³²C. Meyer, *Matrix Analysis and Applied Linear Algebra* (SIAM, Philadelphia, PA, 2000).
- ³³I. Mezić, "On geometrical and statistical properties of dynamical systems: Theory and applications," Ph.D. dissertation (California Institute of Technology, 1994), <http://resolver.caltech.edu/CaltechETD:etd-07212005-131406>.
- ³⁴I. Mezić, "Spectral properties of dynamical systems, model reduction, and decompositions," *Nonlinear Dyn.* **41**(1–3), 309–325 (2005).
- ³⁵I. Mezić and A. Banaszuk, "Comparison of systems with complex behavior: Spectral methods," in *Proceedings of the 39th IEEE Conference on Decision and Control* (2000), p. 8.
- ³⁶I. Mezić and A. Banaszuk, "Comparison of systems with complex behavior," *Physica D* **197**(1–2), 101–133 (2004).
- ³⁷I. Mezić and S. Wiggins, "A method for visualization of invariant sets of dynamical systems based on the ergodic partition," *Chaos* **9**(1), 213–218 (1999).
- ³⁸See <http://apps1.eere.energy.gov/buildings/energyplus/>.
- ³⁹A. D. Perry and S. Wiggins, "KAM tori are very sticky: Rigorous lower bounds on the time to move away from an invariant Lagrangian torus with linear flow," *Physica D* **71**(1–2), 102–121 (1994).
- ⁴⁰K. Petersen, "Ergodic theory," *Cambridge Studies in Advanced Mathematics* (Cambridge University Press, Cambridge, UK, 1989), Vol. 2.
- ⁴¹V. A. Rokhlin, "Selected topics from the metric theory of dynamical systems," *Usp. Mat. Nauk* **4**(2(30)), 57–128 (1949).
- ⁴²V. A. Rokhlin, "Selected topics from the metric theory of dynamical systems," *Ten Papers on Functional Analysis and Measure Theory*, American Mathematical Society Translations, Series 2 Vol. 49 (American Mathematical Society, 1966), pp. 171–240.
- ⁴³C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, "Spectral analysis of nonlinear flows," *J. Fluid Mech.* **641**, 115–127 (2009).
- ⁴⁴A. Ruhe, "Rational Krylov sequence methods for eigenvalue computation," *Numer. Linear Algebra Appl.* **58**, 391–405 (1984).
- ⁴⁵I. Rypina, S. E. Scott, L. J. Pratt, and M. G. Brown, "Investigating the connection between complexity of isolated trajectories and Lagrangian coherent structures," *Nonlinear Processes Geophys.* **18**(6), 977–987 (2011).
- ⁴⁶Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM J. Sci. Stat. Comput.* **7**(3), 14 (1986).
- ⁴⁷P. J. Schmid, "Dynamic mode decomposition of numerical and experimental data," *J. Fluid Mech.* **656**, 24 (2010).
- ⁴⁸S. E. Scott, T. C. Redd, L. Kuznetsov, I. Mezić, and C. K. R. T. Jones, "Capturing deviation from ergodicity at different scales," *Physica D* **238**(16), 1668–1679 (2009).
- ⁴⁹A. Seena and H. J. Sung, "Dynamic mode decomposition of turbulent cavity flows for self-sustained oscillations," *Int. J. Heat Fluid Flow* **32**(6), 13 (2011).
- ⁵⁰Y. Shi, R. Lai, K. Kern, N. Sicotte, I. Dinov, and A. Toga, "Harmonic surface mapping with Laplace-Beltrami eigenmaps," *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2008*, edited by D. Metaxas, L. Axel, G. Fichtinger, and G. Székely (Springer, Berlin/Heidelberg, 2008), pp. 147–154.
- ⁵¹R. Sine, "Geometric theory of a single Markov operator," *Pac. J. Math.* **27**, 155–166 (1968).
- ⁵²Y. Susuki and I. Mezić, "Ergodic partition of phase space in continuous dynamical systems," in *Proceedings of the 48th IEEE Conference on Decision and Control held jointly with the 28th Chinese Control Conference, CDC/CCC* (2009), pp. 7497–7502.
- ⁵³Y. Susuki and I. Mezić, "Nonlinear Koopman modes of coupled swing dynamics and coherency identification," in *Power and Energy Society General Meeting* (IEEE, 2010), pp. 1–8.
- ⁵⁴Y. Susuki and I. Mezić, "Nonlinear Koopman modes and coherency identification of coupled swing dynamics," *IEEE Trans. Power Syst.* **26**(4), 1894–1904 (2011).
- ⁵⁵Y. Susuki and I. Mezić, "Correction to "Nonlinear Koopman modes and coherency identification of coupled swing dynamics" [Nov 11 1894–1904]," *IEEE Trans. Power Syst.* **26**(4), 2584 (2011).
- ⁵⁶Y. Susuki and I. Mezić, "Nonlinear Koopman modes and a precursor to power system swing instabilities," *IEEE Trans. Power Syst.* **27**(3), 1182 (2012).
- ⁵⁷D. Treschev, "Width of stochastic layers in near-integrable two-dimensional symplectic maps," *Physica D* **116**(1–2), 21–43 (1998).
- ⁵⁸T. Wichtrey, "Harmonic limits of dynamical and control systems," Ph.D. dissertation (Logos Verlag, Berlin, 2010).
- ⁵⁹N. Wiener and A. Wintner, "Harmonic analysis and ergodic theory," *Am. J. Math.* **63**, 415–426 (1941).
- ⁶⁰K. Yosida, *Functional Analysis* (Springer-Verlag, 1995).
- ⁶¹L.-S. Young, "What are SRB measures, and which dynamical systems have them?," *J. Stat. Phys.* **108**(5–6), 733–754 (2002).