Chaos

Manifold learning for organizing unstructured sets of process observations

Cite as: Chaos **30**, 043108 (2020); https://doi.org/10.1063/1.5133725
Submitted: 29 October 2019 . Accepted: 11 March 2020 . Published Online: 03 April 2020

Felix Dietrich, Mahdi Kooshkbaghi, Erik M. Bollt, and Ioannis G. Kevrekidis

Manifold learning for organizing unstructured sets of process observations

Cite as: Chaos **30**, 043108 (2020); doi: 10.1063/1.5133725
Submitted: 29 October 2019 · Accepted: 11 March 2020 ·
Published Online: 3 April 2020

Felix Dietrich,[1] Mahdi Kooshkbaghi,[2] Erik M. Bollt,[3] and Ioannis G. Kevrekidis[1,a)]

AFFILIATIONS

[1]Department of Chemical and Biomolecular Engineering and Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, Maryland 21218, USA
[2]Program in Applied and Computational Mathematics, Princeton University, Princeton, New Jersey 08544, USA
[3]Department of Electrical and Computer Engineering, Clarkson Center for Complex Systems Science, Clarkson University, Potsdam, New York 13699-5815, USA

[a)]**Author to whom correspondence should be addressed:** yannisk@jhu.edu

ABSTRACT

Data mining is routinely used to organize ensembles of short temporal observations so as to reconstruct useful, low-dimensional realizations of an underlying dynamical system. In this paper, we use manifold learning to organize unstructured ensembles of observations ("trials") of a system's response surface. We have no control over where every trial starts, and during each trial, operating conditions are varied by turning "agnostic" knobs, which change system parameters in a systematic, but unknown way. As one (or more) knobs "turn," we record (possibly partial) observations of the system response. We demonstrate how such partial and disorganized observation ensembles can be integrated into coherent response surfaces whose dimension and parametrization can be systematically recovered in a data-driven fashion. The approach can be justified through the Whitney and Takens embedding theorems, allowing reconstruction of manifolds/attractors through different types of observations. We demonstrate our approach by organizing unstructured observations of response surfaces, including the reconstruction of a cusp bifurcation surface for hydrogen combustion in a continuous stirred tank reactor. Finally, we demonstrate how this observation-based reconstruction naturally leads to informative transport maps between the input parameter space and output/state variable spaces.

This work presents a data-driven approach that leverages machine learning methods to organize fragments of parametrically observed responses of a dynamical system (e.g., its steady states) to assemble and learn the correct global response surface. Traditionally, given a mathematical model, such response surfaces are computed through established numerical continuation and bifurcation packages such as AUTO or MATCONT. In contrast to this "complete control, full model knowledge" scenario, we explore a scenario closer to what an experimentalist might observe when exploring a new, unknown, or only partially understood system as an ensemble of disorganized experimental "trials." We assume that we cannot measure all components of the system state and that the location of the starting point of each trial in the parameter and state space is also unknown (e.g., set by uncontrolled environmental conditions). We show that data mining the disorganized observations of the system response in such an uncontrolled setting can still lead to the construction of a meaningful realization of the correct response surface. This can subsequently be exploited to classify, analyze, and even predict the system response.

I. INTRODUCTION

When an accurate mathematical model of a dynamical system is available, one can systematically observe the dependence of its response (its long-term dynamics, for example, its steady states) on its parameters by computing the system bifurcation diagram/response surface through established numerical continuation and bifurcation packages such as AUTO[1,2] or MATCONT.[3] One starts from a well defined initial point on the response surface (e.g., a steady state at a particular set of parameter settings) and then builds the surface by systematically moving on it and recording observations. This exploration typically involves varying one parameter at a time (i.e., following one-dimensional curves on this surface, performing "one-parameter continuation"). It is also possible to explore

the response surface through simplicial continuation, systematically varying two (and possibly even more) parameters at a time. At every new point of a simplex in a parameter space, the algorithms return the steady state values of each and every model variable.

In contrast to this "complete control, full knowledge" scenario, we want to explore a scenario closer to what an experimentalist might observe when exploring a new, unknown, or only partially understood system. We may not be able to measure all components of the system state; during each "trial" (each sequence of experiments), we can vary experimental "knobs" that systematically change conditions but in a way unknown to us; and finally, the location of the starting point of each trial in the parameter and state space may also be unknown (e.g., set by uncontrolled environmental conditions).

In this paper, we will show that data mining the (partial) observations of the system response in such an uncontrolled (unstructured, agnostic) setting through manifold learning techniques such as diffusion maps[4] can facilitate the construction of a meaningful realization of the correct response surface. This can subsequently be exploited to classify, analyze, and even predict the system response to variations of our experimental "knobs." An instructive caricature of the procedure is illustrated in Fig. 1, in the form of the well-known cusp surface $0 = \mu + \lambda x - x^3$ where a single state variable x is depicted in relation to the two input parameters (λ, μ). If information about (λ, μ, x) in Fig. 1 at every point (or a fine enough grid of points) on the surface is available, it is straightforward to analyze and visualize the entire two-dimensional surface in three dimensions. A typical numerical computation would fix the value of all but one of the parameters. Here, we fix λ along, say, the blue or green curve and then continue the solution in the parameter μ along (yellow) one-parameter segments on these curves. In this context, direct observation of the surface would involve recording of x, μ, and λ triplets throughout the surface. The surface is the graph of the function $\mu = \mu(x, \lambda) := x^3 - \lambda x$, and manifold learning techniques such as diffusion maps can easily parametrize it, as we describe below.

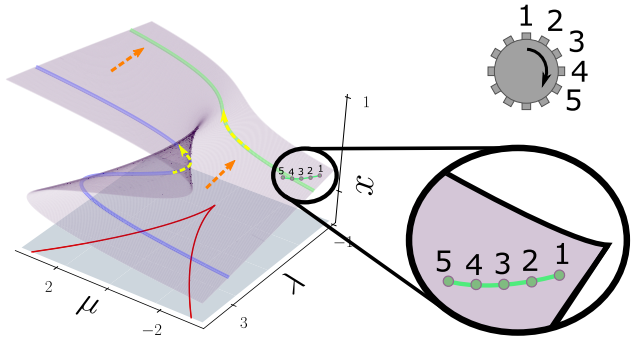

FIG. 1. The cusp surface embedded in parameters and state space $\mathbb{R}^2 \times \mathbb{R}$. The blue and green curves are observed at two constant λ values. Yellow and orange arrows indicate short observation segments for each of the two typical one-parameter continuation directions along the surface. For μ and λ values located between the two red "fold" lines, the system exhibits hysteresis. The inset, with the green curve and the "knob," illustrates our scenario in which only conditions along short curves on the surface can be visited in a systematic but agnostic manner.

In our more agnostic version of the exploration scenario, an experimenter cannot measure the parameter settings (λ, μ) of the system but is able to affect them through a single "knob" that changes them and through them changes the location of the state x (Fig. 1). Assume that the initial position of each trial, (marked as "1") in the figure inset, is determined randomly by the environment, also without the knowledge of the experimenter. After the trial is initialized, for each shown angle 1–5 of the knob, a corresponding point on the surface is visited, and the state value x is recorded along the green curve segment depicted. For simplicity, one can consider, as we do here, that turning the knob at a constant rate moves the point visited at constant speed along the response surface; note that the approach does not rely on this particular assumption. Repeating this tabulation with a large number of experiments from randomly distributed initial trials will result in a collection of recordings of five consecutive values of x [see Fig. (2), center] without knowing the corresponding λ, μ values.

First, we demonstrate how such partial and disorganized observations of a response surface can be integrated in a coherent surface whose topology (the right ordering of the trials) and parametrization can be systematically recovered in a data-driven fashion. Second, we extend the approach to different types of bifurcation observations (not just one-parameter continuation) and demonstrate it in a more applied scenario with a Continuous Stirred Tank Reaction (CSTR) combustion problem. Finally, we demonstrate how this observation-based reconstruction naturally leads to the construction of transport maps between the input (parameter space) and the output (state variable space) of the system or model.

Figure 2 illustrates the ensemble of *several short, disorganized, possibly partial* observation sequences one might obtain from a set of trials that densely sample the surface. Reconstructing the entire bifurcation surface from such a set connects with ongoing research in data-driven identification of dynamical systems from time series, e.g.,[5–7] where reconstructing useful phase space realizations and even dynamics from partial observations of time series has a long history.[8,9]

II. RECONSTRUCTION FROM PROCESS HISTORY

We first consider the case that λ is unknown but fixed (here, the blue curve in Fig. 1), and our observation mode leads to a large, disorganized collection of points in $\mathbb{R}^m$; here, $m = 5$ values of x along short segments of the blue curve; the parameters λ and μ are not observed. The central question becomes: To what extent can we reconstruct the blue curve from these trial records? This is where the embedding theorems by Whitney and Takens (see the Appendix) become relevant: we consider the values of x along the curve as analogous to time-delayed measurements along a temporal trajectory (where the angle of our parameter knob plays here the role of time); the blue curve is one-dimensional; therefore, following Whitney's theorem, five observations are sufficient ($m > 2d$, with $d = 1$ the intrinsic and m the embedding dimension) to reconstruct an embedding of it, *even though we have no direct observation of* μ; μ is inferred only implicitly.

The number of delays necessary for an embedding in Takens' theorem depends on the dimension of the manifold, which might not be known beforehand. Keeping more delays results in

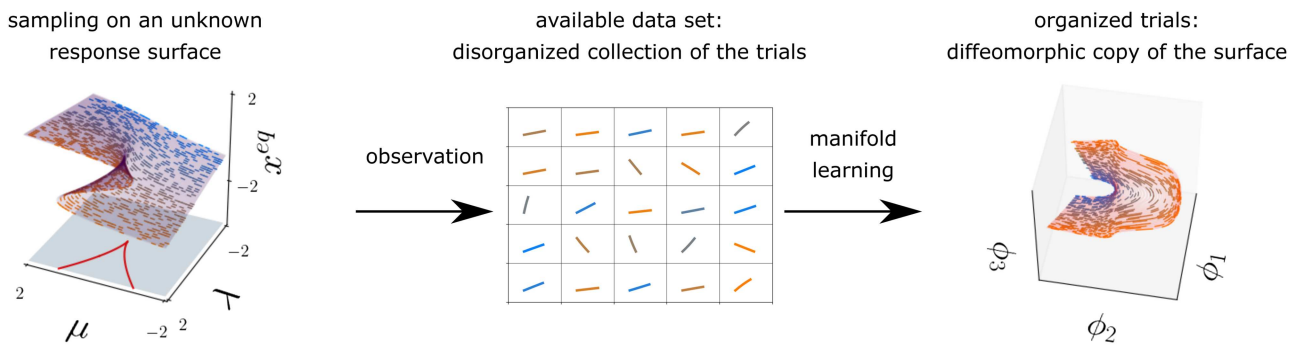

FIG. 2. Observations of short trials from an unknown response surface (e.g., x coordinates of the five points along each of the short lines, left) are available in a dataset with no information about the global position of the lines (they are randomly shuffled, center). Manifold learning will organize the dataset (right) into a diffeomorphic copy of the original surface. The paper explains and demonstrates that this approach can be employed for many different types of observations, as long as sufficiently rich local observation process histories are available.

an embedding in a Euclidean space that is higher-dimensional than necessary. To reduce this dimension, we use a non-linear manifold learning technique, Diffusion Maps (DMAPs).[4] Given N data points $\mathcal{D} = \{y_1, \ldots, y_N\}$ in an ambient Euclidean space $\mathbb{E} = \mathbb{R}^m$ close to a smooth manifold $\mathcal{M}$, we construct a graph between the points, where connectivity is based on a cutoff Gaussian kernel similarity measure with respect to the Euclidean distance in the ambient space $\mathbb{E}$: For a given scale parameter $\epsilon > 0$, the similarity between two distinct points y_i and y_j in $\mathbb{E}$ is defined through $K_{ij} = k(y_i, y_j) = \chi_{[0,\delta]}(r)\exp\left(-r^2/\epsilon\right)$, where $r := d(y_i, y_j)$ and $\chi_{[0,\delta]}$ is the indicator function on $[0, \delta]$, $\delta > 0$ such that the kernel has compact support. Appropriate choices of the parameters ϵ and δ depend on the data.[4,10] The DMAP algorithm is based on the convergence of the normalized graph Laplacian on the data to the Laplace–Beltrami operator on the manifold. If the data points $\mathcal{D}$ are not sampled uniformly in $\mathcal{M}$, the matrix K has to be normalized by an estimation of the density, $P_{ij} = \sum_{i=1}^{N} K_{ij}$, $\tilde{K} = P^{-\alpha} K P^{-\alpha}$ where $\alpha = 0$ (no normalization[11]) can be used in the case of uniform sampling and $\alpha = 1$ otherwise.[4] The kernel matrix $\tilde{K}$ is normalized by the diagonal matrix $D \in \mathbb{R}^{(N\times N)}$, where $D_{ii} = \Sigma_{j=1}^{N} \tilde{K}_{ij}$ for $i = 1, \ldots, N$. The non-linear parametrization (embedding) of the manifold is then given by a certain number L of eigenvectors of $A = D^{-1}\tilde{K} \in \mathbb{R}^{N\times N}$, scaled by their respective eigenvalue (and avoiding harmonics of previous eigenvectors[12]). The new embedding dimension L may be much smaller than the previous ambient space dimension m, in which case DMAP achieves dimensionality reduction.

A. Reconstructing a curve

Indeed, using DMAP on the dataset of vectors comprising five consecutive x observations along the blue curve in Fig. 1 uncovers the correct (one-dimensional) topology (the right relative ordering of the trials) and provides a consistent parametrization in terms of the first nontrivial DMAP coordinate. Therefore, x can be written as a function of the intrinsic variable ϕ_1 [see Fig. 3(a)].

Figure 3(b) illustrates a more challenging scenario, in which during a trial, we only observe sequences of values of μ (which is **not** one-to-one with the blue curve's arclength s). Still, sufficient (in the sense of Whitney) μ observations along each segment allow us to discover the intrinsic one-dimensionality, through the parametrization via the leading diffusion map eigenvector ϕ_1. Figure 3(b) confirms that ϕ_1 *is* one-to-one with arclength; the S-shape of the plot of the average of the μ values in each segment, $\bar{\mu}$, as a function of ϕ_1 provides a data-driven way to "discover" the system hysteresis from local partial observations, a key feature of the cusp surface. One-parameter continuation can be performed "to the left" or "to the right" of a given starting point; therefore, observing the values along a segment in reverse order is also possible. One can take this and other symmetries into account during data mining by constructing metrics that are invariant to them (e.g., Ref. 13).

B. Reconstructing a relation

Our first example was "easy," since an observable exists (here x) that is one-to-one with arclength along the overall sampled curve. Consider now a more interesting, multivalued relation

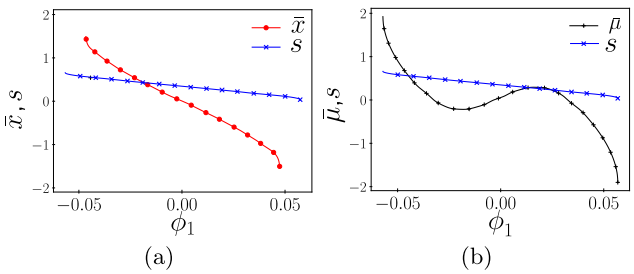

FIG. 3. Results from data mining 40 trial observations (each comprising five measurements) sampled along a single S-shaped curve. Each trial is separated from the next by 0.5% of the total arclength. (a) Recording five x values per segment; the plot shows the arclength s as a function of the first nontrivial DMAP coordinate ϕ_1, which is one-to-one with the arclength. The average observation value for each segment $\bar{x}$ is also plotted. (b) Recording only five μ values per segment; ϕ_1 again provides a consistent parametrization of the curve, even though μ is neither an injective function of ϕ_1 nor of the arclength.

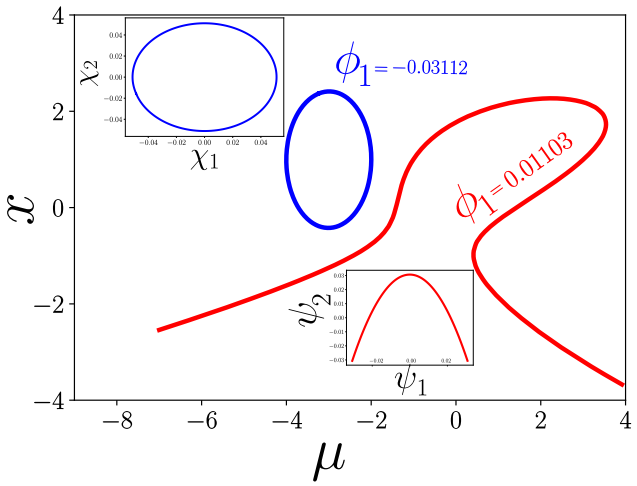

FIG. 4. Multi-valued input–output (parameter–state) relations, defined through $0 = \mu^2 - 2\mu x + x^3 - 2x - 2$ and $0 = (-\mu - 3)^2 - \exp(-0.005(2 - x)) + 0.5(3 - x)^2$. For a single value of the parameter μ, up to three distinct states can coexist. For the same state measurement, up to four different parameter values can be found. In this case, we only record three x values in each trial segment. The original dataset is colored by the first DMAP coordinate ϕ_1. This reveals two clusters with distinct values of $\phi_1 = 0.011\,03$ and $-0.031\,12$. The insets show DMAPS applied on each cluster separately: the isola can be embedded using its first two DMAP eigenvectors (χ_1, χ_2); the mushroom-like curve can be parameterized by its first eigenvector ψ_1. The remaining DMAP coordinates (including ψ_2) are harmonics of ψ_1.

between the input and output, such as the one shown in Fig. 4, where the projections to *both* observable axes are multivalued. Here, we record observations of the two components (red and blue) shown in Fig. 4 in terms of one-parameter arclength segments. The space of observations consisting of sequences of only x measurements along the curves contains a diffeomorphic copy of the two distinct components. Performing DMAP computations with a compactly supported kernel reveals these two disjoint sets since the Markov chain is reducible.[14,15] Separately parameterizing the red and blue components reveals the topology of each component—one is equivalent to a circle and the other to a line. Data mining is thus capable of learning *relations*, and not just functions, discovering various disconnected components, and then providing a useful parameterization of each one of them. We will show later in the paper that this can lead to a meaningful *transport map* between inputs and outputs that circumvents the complicating multivaluedness.

C. Reconstructing a response surface

The third example involves the data mining of short one-parameter segments varying in the μ direction, for a range of randomly distributed λ values, sampled across the entire surface. We only record x values; λ and μ are not explicitly measured but rather implicitly inferred. Figure 5 shows that DMAPS will recover a useful embedding of the entire two-dimensional response surface, one that organizes the unstructured observations in the correct topology.

Picking a different bifurcation observation (e.g., performing short one-parameter continuations in the λ direction for random initial values of x) would also allow us to reconstruct the surface

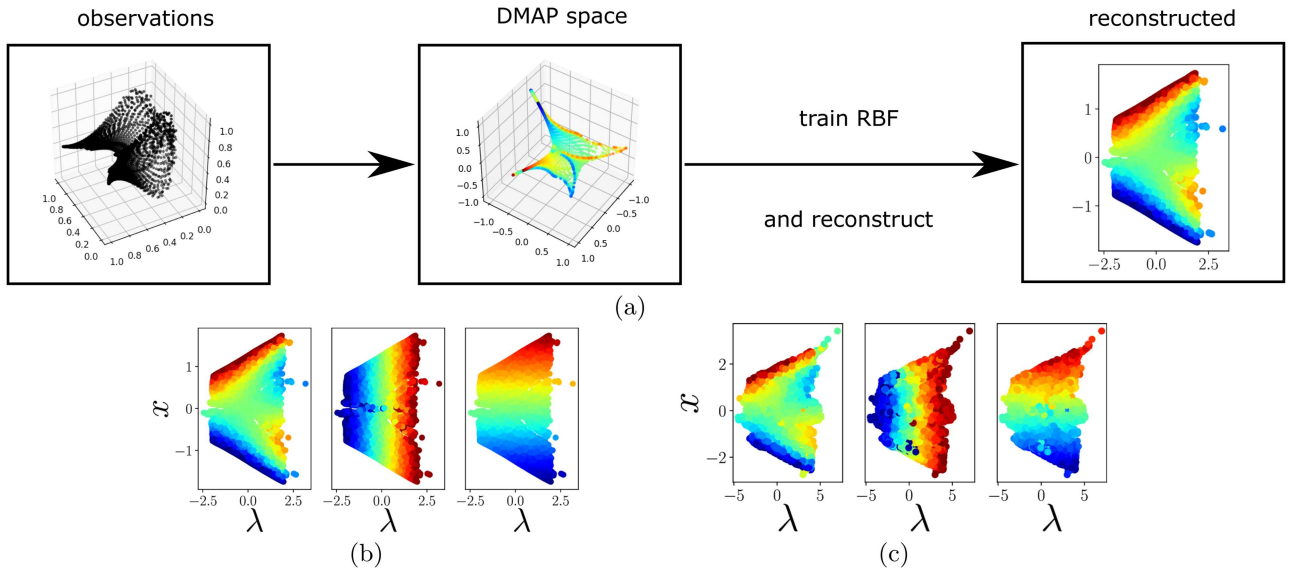

FIG. 5. (a) Recording only five x values from trials involving variation in the μ direction creates an embedding in the observation space $\mathbb{R}^5$ (a projection of the segments to $\mathbb{R}^3$ is shown in black). Parametrizing the data with DMAPS yields a lower-dimensional embedding. We then obtain (μ, λ, x) as functions on the DMAP space, here by using radial basis functions with 10% of the data used for their construction. The plots in (b) show the reconstruction of the remaining 90% of the data (not used for training). (c) Reconstructions of the coordinate functions (color) of the surface projected to the (x, λ) space by recording statistics of 200 μ values in each of 5535 discs (see the text).

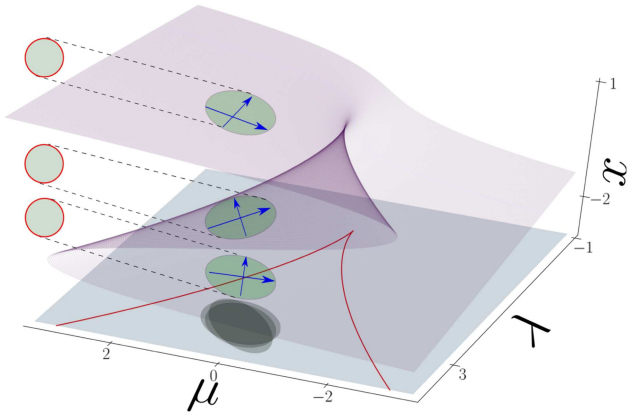

FIG. 6. Patches on the cusp surface projecting to discs in the (x, λ) plane and statistics of the μ values (denoted by arrows) collected in these patches. The parametrization and subsequent estimation of the coordinate functions are shown in Fig. 5(c).

(not shown); certainly, the dimension and the topology would be the same, even though the observation process is different.

An illustration of a *qualitatively different* observation process, one that is not constrained along a one-parameter path, is shown in Fig. 5(c). Now, we observe the response surfaces by selecting small patches around randomly sampled initial points (see Fig. 6). Here, we record *statistics of* μ *values* at points uniformly sampled in disc-shaped patches in the (x, λ) space. By statistics, we mean the first five moments of the distribution; other options are also possible (e.g., principal component analysis of histograms of μ values[16]). The correct dimensionality and topology and a useful parametrization of the surface are again recovered and will indeed be recovered for any sufficiently rich set of generic (alternatively, prevalent)[9,17] observables. In this particular example, the set of five moments of the measured distributions is demonstrated to be prevalent in Fig. 6(c): we can recover the correct topology on the moments, by approximating the three coordinate functions (x, μ, λ) on the embedding. If the moments would not be generic observables, they would collapse parts of the surface, and the coordinate functions would not be fully recoverable. The example demonstrates that we can recover the dimension and topology, and construct a useful geometry of the response surface, even in cases where only a few quantities can be measured. Different reconstructed surfaces, obtained, for example, from different types of observation of the same underlying surface, can be mapped to each other—we will return to this in the conclusions.

Takens embeddings[9] using temporal observations of the system state can be formulated in either discrete or continuous time: one can either use values of the observable at a discrete number of time points or use time derivatives of the observable at a given time point as the additional embedding dimensions. The analogous observables for response surfaces would be a state variable measurement for a single-parameter setting and then derivatives of the observable(s) with respect to the continuation parameter or, alternatively, with respect to the bifurcation curve arclength at that point. Indeed, k-jet extensions[18] lift mappings to higher-dimensional spaces by replacing the value of the mapping by its Taylor series expansion of degree k. This can be viewed as an approximation of using germs or, in our case, short continuation segments to represent the map. Singularity theory characterizes when these extensions have images that are smooth manifolds.

D. A chemical reactor example

Having introduced our main idea and demonstrated it through illustrative examples (with some of the relevant theory following in the Appendix), we now apply it in a hydrogen combustion model setting using a CSTR.

The reactor is a simplified model for the study of combustion dynamics, in which an extreme mixing assumption leads to a homogeneous mixture of reactive ideal gases. Due to the simplicity of the model, one can investigate the contribution of kinetics parameters on the observed dynamics. The dependence of the reactor state, e.g., the reactor temperature T, on the residence time in the reactor, τ, typically displays an S-shaped curve, connecting the weakly- and strongly burning steady state branches via an unstable steady state branch.[19] The turning points in the S-shaped curve correspond to ignition and extinction limits of the mixture. A system of n_s chemical species $M_i, i = 1, \ldots, n_s$ reacts according to n_r reversible elementary reactions,

$$\sum_{i=1}^{n_s} \nu'_{ik} M_i \rightleftharpoons \sum_{i=1}^{n_s} \nu''_{ik} M_i, \quad k = 1, \ldots, n_r. \tag{1}$$

Stoichiometry is defined by ν'_{ik} and ν''_{ik}, the stoichiometric coefficients of species i in reaction k for the reactants and products, respectively. The rate of the kth elementary reaction is

$$q_k = q_k^f - q_k^r = k_k^f \prod_{i=1}^{n_s} [X_i]^{\nu'_{ik}} - k_k^r \prod_{i=1}^{n_s} [X_i]^{\nu''_{ik}}, \quad k = 1, \ldots, n_r, \tag{2}$$

where $[X_i]$ denotes molar concentration of species i and k_k^f and k_k^r are the forward and reverse rate constants of reaction k. The production or consumption rate $\dot{\omega}_i$ of the ith species is the summation of the rates of all reactions involving the species i,

$$\dot{\omega}_i = \sum_{k=1}^{n_r} \nu_{ik} q_k, \tag{3}$$

where $\nu_{ik} = \nu''_{ik} - \nu'_{ik}$ is the net stoichiometric coefficient.

The temporal evolution of Y_i, the mass fraction of species i, and temperature T in a perfectly stirred reactor is described by a system of $(n_s + 1)$ ordinary differential equations (ODEs),[20]

$$\frac{dY_i}{dt} = \frac{1}{\tau}(Y_i^0 - Y_i) + \frac{\dot{\omega}_i W_i}{\rho}, \quad i = 1, \ldots, n_s, \tag{4}$$

$$\frac{dT}{dt} = \frac{1}{\bar{c}_p \tau} \sum_{i=1}^{n_s} (h_i^0 - h_i) Y_i^0 - \frac{1}{\rho \bar{c}_p} \sum_{i=1}^{n_s} h_i W_i \dot{\omega}_i,$$

where Y_i^0 and h_i^0 are the mass fraction and total enthalpy of species i at the inflow, W_i and h_i are the molecular weight and total enthalpy of species i, and $\bar{c}_p$ and ρ are the mixture heat capacity under

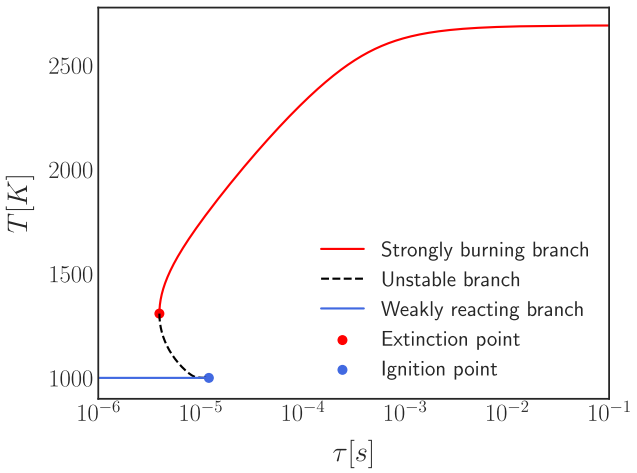

FIG. 7. The dependence of the reactor temperature on the residence time for a stoichiometric H_2/air mixture with the initial temperature $T_0 = 1000$ K under atmospheric pressure of 1 atm. The lower branch starts from a "frozen" state at very short residence time and stays "weakly reactive" up to the ignition point. The system state then jumps from an extinguished state to the "strongly burning branch." If the system is already ignited and the residence time is decreased gradually, the system will jump back to the weakly reactive state locus at the extinction limit.

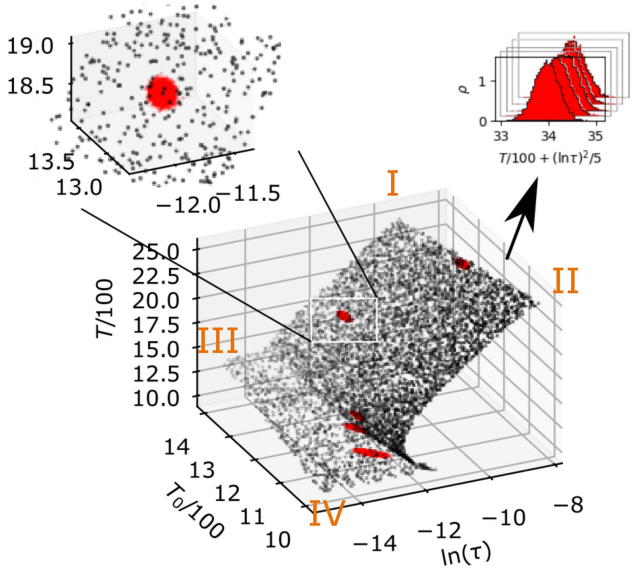

FIG. 8. Bifurcation surface of the CSTR, embedded in $(\ln\tau, T_0/100, T/100)$ space. The plot shows the surrounding neighborhoods (red) of five of the points on the surface, and the inset shows a zoomed-in version. The points in each red neighborhood are assembled into a histogram (top right), and we then compute the first four moments of the histogram—these four moments at each point on the surface will be the only data available for our data mining. The "corners" of the dataset are labeled (I–IV) for easier visual reference to Fig. 9.

constant pressure and density. Fixing the inflow mixture composition, the reactor temperature is a function of inlet temperature T_0, residence time τ, and the density, respectively.

In this study, we used a H_2/air ignition model in which the detailed kinetic mechanism includes 9 species ($n_s = 9$), participating in 21 elementary reversible reactions ($n_r = 21$).[21] AUTO-07p[1,2] code (a bifurcation analysis tool) is coupled with CHEMKIN[22] (a chemical kinetics database) for performing the numerical continuation computations.[20]

The dependence of the steady state reactor temperature T on the residence time for a stoichiometric H_2/air mixture (with pressure set to 1 atm, $T_0 = 1000$ K) in an adiabatic CSTR is shown in Fig. 7.

We assume that an experimenter is able to initialize the reactor at many different values of the initial temperature T_0 and residence time $\ln\tau$, but without necessarily having knowledge of the numerical value of these parameters. We also assume that the initialization at a fixed set of parameter values of $(T_0, \ln\tau)$ can be repeated, each time with a normally distributed inaccuracy δ in the position across the surface (see Fig. 8). After initialization, the experimenter then records real-valued observations. For the computational experiment, we rather arbitrarily choose $y = (\ln\tau)^2/5 + T$, a combination of τ and the temperature T at steady state as the observed/recorded quantity. Notably, the initial temperature T_0 is not part of the observation—in fact, by Theorem 4 (Appendix), almost any one-dimensional combination of the variables is admissible. We initialize 5535 points on the surface with rescaled coordinates $(\ln\tau, T_0/100, T/100) \in [-15, -8] \times [10, 14] \times [10, 24]$. Similar to the illustration in Fig. 6, at any given point p out of the 5535 points on the surface, we record values of $y = (\ln\tau)^2/5 + T$ for 10 000 points $(p + \delta)$, i.e., in a small neighborhood of p on the surface (red ellipses in Fig. 8). The first four moments (the mean, standard deviation, skew, and kurtosis) of the values of y are used as the measured features of the trial associated with the point p. This leads to a dataset $X \in \mathbb{R}^{5535\times4}$. We then apply diffusion maps with kernel bandwidth $\epsilon = 1.25$ (see Refs. 23 and 24 for a discussion of the choice of the kernel bandwidth) to parametrize the manifold embedded in this space. Figure 9 shows empirical evidence that a diffeomorphic copy of the two-dimensional response surface for this reactor can be recovered, even though we only had access to a few moments of the scalar observations y on the surface (a physically

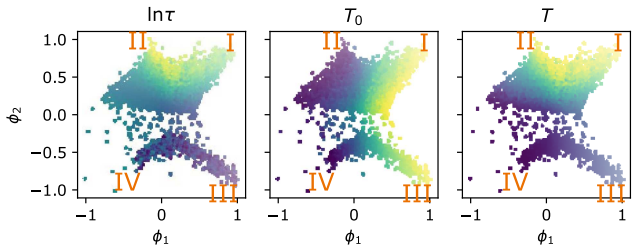

FIG. 9. Embedding of the moment data in the diffusion map space, colored by the log residence time $\ln\tau$, initial temperature T_0, and steady state temperature T. The two clusters in the point cloud belong to the regions before and after the fold on the diagram. The labels of the edges (I–IV) correspond to the labels in Fig. 8.

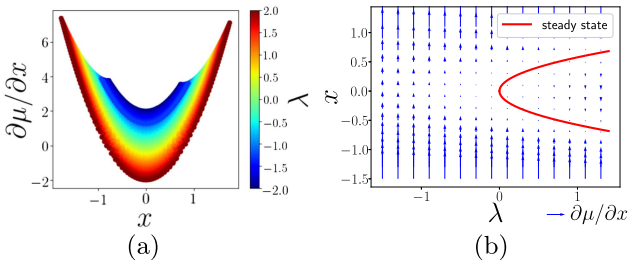

FIG. 10. (a) Embedding of the cusp surface in x, $\partial\mu/\partial x$ for each fixed λ. λ can be written as a function over this embedding (color). This shows that we can recover the second parameter without measuring it directly and just by local information attained by change in the first parameter. (b) A one-dimensional vector field over x, parametrized by λ. The vectors $\partial\mu/\partial x$ at each x (the arrows) are given by $3x^2 - \lambda$. Two steady state branches on which the vector field is zero are "born" at the turning point, corresponding to the original cusp point.

rather unusual observation choice!).

III. VECTOR FIELDS AND TRANSPORT MAPS

An embedding of the cusp surface in a two-dimensional $(x, \partial\mu/\partial x)$ space is shown in Fig. 10. It is maybe interesting that this response surface embedding can be alternatively interpreted as an observation of a vector field $\mathbf{V}$ over x, parametrized by λ, i.e., $\mathbf{V}(x;\lambda) := \frac{\partial\mu}{\partial x}(x;\lambda)$ or, equivalently, as the right-hand side of an ordinary differential equation, $\frac{d}{dt}x(t) = \mathbf{V}(x(t);\lambda)$. This ODE for $x(t)$ is not physically meaningful—it is just an alternative way to parametrize the surface (thus defining it implicitly), rather than storing it in its entirety as a geometric object. As shown in Fig. 10(b), the turning points of the cusp surface become steady state branches for the dynamics of this "surface-generating" vector field. The cusp point now describes a saddle-node bifurcation. Short continuations in the μ direction (for fixed λ) correspond to *time-series segments* for this constructed vector field. "Time" for this vector field goes to infinity at finite μ (at the turning points, which now become steady states).

Clearly, tools for signal-processing and identification of dynamical systems (e.g., Koopman operator approximations,[25–27] time-delay embedding techniques, information theory,[28,29] etc.) can be exploited to develop useful observables for continuation-based bifurcation surface reconstruction.

The process of observing response surfaces also has interesting implications in studying transport maps between inputs and outputs.[30] Consider the relation we explored in Fig. 4 and consider an observation process that samples the arclength of each of its two component curves at equal steps; this produces a constant density of observations along the arclength of each component. Projecting this uniform density onto the input and the output axes results in complicated density profiles, with the density approaching infinity at turning points of the curve (spikes in Fig. 11). Detection of the location of such singularities in a higher-dimensional setting is also possible through data-driven analysis of the graph Laplacian on the data.[31] Given the densities on input and output axes, it is natural to explore a transport map between them.

Instead of trying to construct such a transport map through an appropriately defined optimization process,[32,33] we can use data mining of the bifurcation observations themselves to provide a meaningful solution. In the example shown in Fig. 11, DMAP can uncover and parametrize the two disconnected components, allowing us to write the relation between input and output *in a parametric form in terms of DMAP coordinates*. Since the relation embodies the transport map, this helps us effectively recover it in the same parametric form. Projecting the density on the intrinsic parametrization separately to the input axis and then to the output axis demonstrates how the singularities in input and output densities form: the projection of the uniform density ρ_s along the arclength s results in the density ρ_y on the output (vertical, y) axis. The result ρ_y is seen, in this case, to be the sum of transports over piecewise invertible branch segments $g_k : \mathbb{R} \to \mathbb{R}$, $k = 1, \ldots, 4$ of the curves (differentiated by

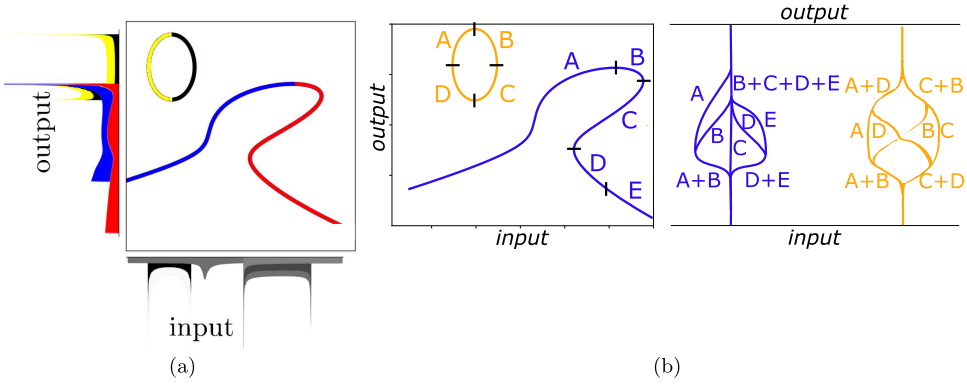

FIG. 11. (a) Uniform density on the arclength of each of two curves, transported through the relation to the two axes. The colors (yellow and black for the isola and red and blue for the "mushroom") show which part of the resulting density is caused by which part of the relation. The black density on the horizontal axis shows an equivalent transport (with the corresponding colored parts on the curve not indicated here). (b) One-to-one sub-branches of the relation, with branching during the transport. The colors on the right panel follow from the colors of the curves in the center panel. The curves are split into more segments than the four colored ones on the left panel because the branched transport can be considered in both the input–output and the output–input directions.

color in Fig. 11),

$$\rho_y(y) := \sum_{k=1}^{4} \left| \frac{\partial g_k}{\partial s}(g_k^{-1}(y)) \right|^{-1} \rho_s(g_k^{-1}(y)). \tag{5}$$

This formula encodes the action of a type of a transfer operator on the density ρ_s (see, e.g., Refs. 34–36).

Similar considerations apply to the projection on the "input" x axis. The book-keeping introduced through the intrinsic parametrization allows us to determine how the input and output densities are connected through a form of branched transport (see Fig. 11). On the plane, Steiner's minimal path problem[37] addresses finding an optimal connecting path that is embodying the transport plan between a finite number of points. The branched transport defined in the paper by Xia[38] greatly generalizes this to the decomposition of measures into atoms (and back), where the transport plan often involves several Y-shaped branches. Similarly, the decomposition of the bifurcation diagram *into invertible regions over state and parameter spaces* is visualized in Fig. 11(b). It decomposes the bifurcation diagrams into invertible functions over input and output spaces. Using this decomposition, we can specify how points are transported from input to output, reminiscent of branched transport. It is important to state that this parametrization of the transport map has required more than just single point observations. The additional information necessary to embed and parametrize the "relation manifold" is gained through *the observation process*, i.e., the "context" provided in each set of observations.

IV. CONCLUSIONS

In summary, we have demonstrated that observations of input–output relations, in the form of ensembles of disorganized, even partially recorded, short "trials" can be rearranged and combined to construct a realization (and parametrization) of the full response surface through manifold learning techniques. This relies crucially on the Whitney and Takens embedding theorems, typically used in time-series analysis of dynamical systems, now initially applied to one-parameter continuation for the construction of response surfaces. We also demonstrated that different types of observations (not just single-parameter continuation segments) can be used similarly. It is possible (though it is not shown here) to compute diffusion maps not with the Euclidean distance between observations, but with a Mahalanobis-like metric that takes into account local observation covariances.[39,40] This has the potential of fusing different observation sets (observations from one-parameter continuation runs in different directions, statistics from two-dimensional patches, etc.) in a single "master" observation surface but requires a consistent way to estimate covariance matrices for the given observation process. Our approach can then be useful in a domain-adaptation context by constructing meaningful realizations of the input and output ("source" and "target") domains. It would be interesting to compare the Mahalanobis-like metric driven observation fusion with other registration approaches developed in the domain-adaptation literature.[41,42]

There is a conceptual similarity to the dynamic Laplacian[43] in that, to gather covariance information, we have to start at several nearby initial trial points and then perform the trial associated with each one of them. In addition, we explored how the data-driven recovery of the dimensionality and topology/geometry of the input–output relation can lead to a useful transport map between parameter (input) and state (output) spaces. Current work explores extending the approach to other response surface types, as well as the relation of this "observation-based" transport map to those obtained through optimal transport considerations.[44]

ACKNOWLEDGMENTS

This work was partially funded by the National Science Foundation (NSF), the Defense Advanced Research Projects Agency (DARPA) (I.G.K. and F.D.), the SNSF (Grant No. P2EZP2_168833) (M.K.), and the Army Research Office (ARO) (I.G.K., E.M.B.) and Office of Naval Research (ONR) (E.M.B.). Discussions with Professor J. Guckenheimer are gratefully acknowledged.

APPENDIX: RELEVANT THEOREMS

Let $k \geq d \in \mathbb{N}$ and consider a manifold $\mathcal{M} \subset \mathbb{R}^k$ that is d-dimensional, compact, smooth, connected, oriented, and endowed with a Riemannian metric g induced by the embedding in a k-dimensional Euclidean space. Note that this setting is sufficient for our presentation but is more restrictive than allowed by the results cited below. Together with the results from Packard *et al.*[45] and Aeyels,[46] the definitions and theorems of Takens[9] describe the concept of observability of state spaces of nonlinear dynamical systems. A dynamical system is defined through its state space (here, the manifold $\mathcal{M}$) and a diffeomorphism $\phi : \mathcal{M} \to \mathcal{M}$.

Theorem 1 (Generic delay embeddings). *For pairs (ϕ, y), $\phi : \mathcal{M} \to \mathcal{M}$ a smooth diffeomorphism and $y : \mathcal{M} \to \mathbb{R}$ a smooth function, it is a generic property that the map $\Phi_{(\phi,y)} : \mathcal{M} \to \mathbb{R}^{2d+1}$, defined by*

$$\Phi_{(\phi,y)}(x) = \left(y(x), y(\phi(x)), \ldots, y(\underbrace{\phi \circ \cdots \circ \phi}_{2d\ times}(x)) \right), \tag{A1}$$

is an embedding of $\mathcal{M}$; here, "smooth" means at least C^2.

Genericity in this context is defined as "an open and dense set of pairs (ϕ, y)" in the C^2 function space. In Takens' paper, there is also an infinitesimal version of Theorem 1:

Theorem 2 (Generic differential embeddings). *For pairs (X, y), $X : \mathcal{M} \to T\mathcal{M}$ a smooth vector field with flow $\phi_t : \mathcal{M} \to \mathcal{M}$ and $y : \mathcal{M} \to \mathbb{R}$ a smooth function, it is a generic property that the map $\Phi_{(X,y)} : \mathcal{M} \to \mathbb{R}^{2d+1}$, defined by*

$$\Phi_{(X,y)}(x) = \left(y(x), \frac{d}{dt} y(\phi_t(x))\Big|_{t=0}, \ldots, \frac{d^{2d}}{dt^{2d}} y(\phi_t(x))\Big|_{t=0} \right), \tag{A2}$$

is an embedding of $\mathcal{M}$; here, "smooth" means at least C^{2d+1}.

For an extension of Takens' theorems to deterministically forced, input–output, irregularly sampled, and stochastic systems, we refer the reader to the results from Stark *et al.*[47–49]

Generic (open and dense) sets can have measure zero; therefore, Sauer *et al.*[17] refined the results significantly by introducing the concept of prevalence (a "probability one" analog in infinite dimensional spaces).

Definition 1. *A Borel subset S of a normed linear space V is* ***prevalent*** *if there is a finite-dimensional subspace E of V such that for each $v \in V$, $v + e$ belongs to S for (Lebesgue-) almost every e in E.*

Using this notion, one can strengthen the result from the original theorem of Whitney,[50] into the prevalence form:

Theorem 3 [Whitney (weak form)]. *The set $S \subset C^1$ of smooth maps $F : \mathbb{R}^k \to \mathbb{R}^{2d+1}$ that are embeddings of $\mathcal{M}$ is an open and dense set in the C^1-topology.*

Theorem 4 [Whitney (with prevalence)]. *The set $S \subset C^1$ of smooth maps $F : \mathbb{R}^k \to \mathbb{R}^{2d+1}$ that are embeddings of $\mathcal{M}$ is prevalent.*

In particular, given any smooth map F, not only are the maps arbitrarily near F that are embeddings (which is the notion of genericity from Takens), but "almost all" (in the sense of prevalence) of the maps near F are embeddings. The space E from the definition of prevalence in Theorem 4 is the $k(2d + 1)$-dimensional space of linear maps from $\mathbb{R}^k$ to $\mathbb{R}^{2d+1}$.

For completeness, we also add the statement of the following, a strong form of the Whitney theorem.

Theorem 5 [Whitney (strong form, existence)]. *For every d-dimensional manifold $\mathcal{M}$ of the form given above, there exists an embedding into $\mathbb{R}^{2d}$.*

Note that Theorem 5 is a stronger statement than Theorem 3 in terms of the dimension of the embedding space: ($\mathbb{R}^{2d}$ instead of $\mathbb{R}^{2d+1}$), yet is not as relevant in practice, since it does not have the same probabilistic notion of prevalence.

REFERENCES

[1] E. J. Doedel, Congr. Numer. **30**, 25 (1981).
[2] E. J. Doedel, T. F. Fairgrieve, B. Sandstede, A. R. Champneys, Y. A. Kuznetsov, and X. Wang (2007), see http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.423.2590.
[3] A. Dhooge, W. Govaerts, and Y. A. Kuznetsov, ACM Trans. Math. Softw. **29**, 141 (2003).
[4] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, Proc. Natl. Acad. Sci. U.S.A. **102**, 7426 (2005).
[5] T. Sauer, Phys. Rev. Lett. **72**, 3811 (1994).
[6] O. Yair, R. Talmon, R. R. Coifman, and I. G. Kevrekidis, Proc. Natl. Acad. Sci. U.S.A. **114**, E7865 (2017).
[7] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Proc. Natl. Acad. Sci. U.S.A. **113**, 3932 (2016).
[8] B. Moore, IEEE Trans. Automat. Contr. **26**, 17 (1981).
[9] F. Takens, in *Dynamical Systems and Turbulence, Warwick 1980* (Springer, 1981), pp. 366–381.
[10] T. Berry and T. Sauer, Appl. Comput. Harmon. Anal. **40**, 439–469 (2015).
[11] M. Belkin and P. Niyogi, Neural Comput. **15**, 1373 (2003).
[12] C. J. Dsilva, R. Talmon, R. R. Coifman, and I. G. Kevrekidis, Appl. Comput. Harmon. Anal. **44**, 759 (2018).
[13] T. Berry and J. Harlim, Appl. Comput. Harmon. Anal. **45**, 84 (2018).
[14] F. R. K. Chung, *Spectral Graph Theory* (American Mathematical Society, 1996).
[15] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis, Appl. Comput. Harmon. Anal. **21**, 113 (2006).
[16] T. Shnitzer, R. Talmon, and J.-J. Slotine, IEEE Trans. Signal Process. **65**, 904 (2017).
[17] T. Sauer, J. A. Yorke, and M. Casdagli, J. Stat. Phys. **65**, 579 (1991).
[18] M. Golubitsky and V. Guillemin, *Stable Mappings and Their Singularities* (Springer US, 1973).
[19] C. K. Law, *Combustion Physics* (Cambridge University Press, 2006).
[20] M. Kooshkbaghi, C. E. Frouzakis, K. Boulouchos, and I. V. Karlin, Combust. Flame **162**, 3166 (2015).
[21] M. Ó. Conaire, H. J. Curran, J. M. Simmie, W. J. Pitz, and C. K. Westbrook, Int. J. Chem. Kinet. **36**, 603 (2004).
[22] R. J. Kee, F. M. Rupley, E. Meeks, and J. A. Miller, "CHEMKIN-III: A FORTRAN chemical kinetics package for the analysis of gas-phase chemical and plasma kinetics," Technical Report, Sandia National Laboratories, Livermore, CA, 1996.
[23] T. Berry and J. Harlim, Appl. Comput. Harmon. Anal. **40**, 68 (2016).
[24] A. Singer, Appl. Comput. Harmon. Anal. **21**, 128 (2006).
[25] M. Budišić, R. Mohr, and I. Mezić, Chaos **22**, 047510 (2012).
[26] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, J. Nonlinear Sci. **25**, 1307 (2015).
[27] E. M. Bollt, Q. Li, F. Dietrich, and I. Kevrekidis, SIAM J. Appl. Dyn. Syst. **17**, 1925 (2018).
[28] C. Bandt and B. Pompe, Phys. Rev. Lett. **88**, 174102 (2002).
[29] R. Talmon, I. Cohen, S. Gannot, and R. R. Coifman, IEEE Signal Process. Mag. **30**, 75 (2013).
[30] L. C. Evans and W. Gangbo, *Differential Equations Methods for the Monge-Kantorevich Mass Transfer Problem* (American Mathematical Society, 1999).
[31] M. Belkin, Q. Que, Y. Wang, and X. Zhou, in *Proceedings of the 25th Annual Conference on Learning Theory*, Proceedings of Machine Learning Research, Vol. 23, edited by S. Mannor, N. Srebro, and R. C. Williamson (PMLR, Edinburgh, Scotland, 2012), pp. 36.1–36.26.
[32] L. N. Wasserstein, Probl. Inform. Transmission **5**, 47 (1969).
[33] C. Villani, *Optimal Transport* (Springer, Berlin, 2009).
[34] V. Baladi, *Positive Transfer Operators and Decay of Correlation* (World Scientific Publishing Co. Inc., 2000).
[35] D. Ruelle, *Thermodynamic Formalism* (Cambridge University Press, 2012).
[36] E. M. Bollt and N. Santitissadeekorn, *Applied and Computational Measurable Dynamics* (Society for Industrial and Applied Mathematics, 2013).
[37] E. N. Gilbert and H. O. Pollak, SIAM J. Appl. Math. **16**, 1 (1968).
[38] Q. Xia, Commun. Contemp. Math. **05**, 251 (2003).
[39] A. Singer and R. R. Coifman, Appl. Comput. Harmon. Anal. **25**, 226 (2008).
[40] C. J. Dsilva, R. Talmon, C. W. Gear, R. R. Coifman, and I. G. Kevrekidis, SIAM J. Appl. Dyn. Sys. **15**, 1327 (2016).
[41] N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy, IEEE Trans. Pattern Anal. Mach. Intell. **39**, 1853 (2017).
[42] O. Yair, M. Ben-Chen, and R. Talmon, IEEE Trans. Signal Process. **67**, 1797 (2019).
[43] G. Froyland and E. Kwok, J. Nonlinear Sci. (2017).
[44] C. Moosmüller, F. Dietrich, and I. G. Kevrekidis, arXiv:1907.08260v2 (2019).
[45] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw, Phys. Rev. Lett. **45**, 712 (1980).
[46] D. Aeyels, SIAM J. Control Optim. **19**, 595 (1981).
[47] J. Stark, D. Broomhead, M. Davies, and J. Huke, Nonlinear Anal. Theory Methods Appl. **30**, 5303 (1997).
[48] J. Stark, J. Nonlinear Sci. **9**, 255 (1999).
[49] J. Stark, D. Broomhead, M. Davies, and J. Huke, J. Nonlinear Sci. **13**, 519 (2003).
[50] H. Whitney, Ann. Math. **37**, 645 (1936).