

Identifying finite-time coherent sets from limited quantities of Lagrangian data

Matthew O. Williams,¹ Irina I. Rypina,² and Clarence W. Rowley³

¹*Program in Applied and Computational Mathematics, Princeton University, New Jersey 08544, USA*

²*Department of Physical Oceanography, Woods Hole Oceanographic Institute, Massachusetts 02543, USA*

³*Department of Mechanical and Aerospace Engineering, Princeton University, New Jersey 08544, USA*

(Received 21 December 2014; accepted 14 July 2015; published online 29 July 2015)

A data-driven procedure for identifying the dominant transport barriers in a time-varying flow from limited quantities of Lagrangian data is presented. Our approach partitions state space into coherent pairs, which are sets of initial conditions chosen to minimize the number of trajectories that “leak” from one set to the other under the influence of a stochastic flow field during a pre-specified interval in time. In practice, this partition is computed by solving an optimization problem to obtain a pair of functions whose signs determine set membership. From prior experience with synthetic, “data rich” test problems, and conceptually related methods based on approximations of the Perron-Frobenius operator, we observe that the functions of interest typically appear to be smooth. We exploit this property by using the basis sets associated with spectral or “mesh-free” methods, and as a result, our approach has the potential to more accurately approximate these functions given a fixed amount of data. In practice, this could enable better approximations of the coherent pairs in problems with relatively limited quantities of Lagrangian data, which is usually the case with experimental geophysical data. We apply this method to three examples of increasing complexity: The first is the double gyre, the second is the Bickley Jet, and the third is data from numerically simulated drifters in the Sulu Sea. © 2015 AIP Publishing LLC.

[<http://dx.doi.org/10.1063/1.4927424>]

Transport barriers separate a fluid flow into regions with qualitatively different Lagrangian behaviors and are important for understanding transport and stirring processes in geophysical flows. We present a method for identifying these barriers by partitioning the state space of the system into coherent sets that are chosen to minimize the number of trajectories that “switch sets” in a given time interval. There are many conceptual similarities between our approach and probabilistic methods, but our approach is tailored to problems with limited quantities of Lagrangian data, which is often the case when the data come from real instruments such as “drifters” released into the ocean. In particular, we exploit the apparent smoothness of the functions of interest by employing basis functions associated with spectral or “mesh-free” methods, which can converge more rapidly than indicator functions in this regime. As a result, useful (although not fully converged) approximations of the coherent sets can be obtained from fewer Lagrangian trajectories compared to other methods. This approach is applied to identify coherent sets in three fluid flows: The first is the double gyre, which is commonly used as a benchmark for different methods, the second is the Bickley Jet, which is an idealized model for stratospheric flow, and the third is a realistic numerically generated near-surface flow in the Sulu Sea.

often chaotic dynamics. The locations (or absence) of these barriers help to determine the mixing properties of the underlying flow^{1–3} and have practical implications in a number of engineering contexts including chemical reactors and combustion¹ as well as ecological applications, such as predicting the extent of oil spills.^{4,5} As a result, a number of effective yet conceptually different approaches for extracting these structures have been developed. Geometric methods focus on the identification of invariant manifolds and finite-time hyperbolic material lines^{6,7} and include methods based on Finite Time Lyapunov Exponents (FTLEs) or Finite Scale Lyapunov Exponents^{8–10} and the associated Lagrangian Coherent Structures^{11–13} and are perhaps the most widely used set of approaches at the current time. However, there are alternative techniques including variational methods,¹⁴ ergodic quotient partitions,¹⁵ trajectory complexity measures,¹² and Lagrangian descriptors.¹⁶

In recent years, probabilistic methods, which use a different definition of coherence, have proved to be a useful alternative to geometric methods,^{17–21} and although they have been applied to general flows, have the advantage of identifying minimally dispersive regions if the flow happens to be autonomous or time-periodic.²² Many of these approaches define coherent sets based on the spectral properties of the Perron-Frobenius operator, which is also referred to as the transfer operator.^{17–21} In practice, this information is often obtained by constructing a finite-dimensional approximation of this operator using the Ulam-Galerkin method,^{17,18,21,23} which has been implemented efficiently in software packages such as GAIO.²³

I. INTRODUCTION

The identification of transport barriers is an important step in understanding fluid flows that have complex and

Many of these methods also assume that the velocity field that defines the motion of fluid parcels or drifters is available. In problems where this field is unknown, it can often be estimated from data using tools such as optical flow.^{24,25} Although these approaches could, in principle, be applied directly to Lagrangian data, the amount of data required for an accurate approximation is often too large to be practical in an experimental setting. Our ambition in this article is to demonstrate that effective approximations of coherent sets can be obtained with limited quantities of Lagrangian data and are therefore well suited to experimentally obtain data sets.

We define coherent pairs as the solution to an optimization problem that can be solved numerically using the Singular Value Decomposition (SVD). The result is a pair of functions whose signs can be used to partition the data into two sets that minimize the number of elements that “leak out” in a given time interval. In the limit of infinite data, this problem can be succinctly expressed as an inner product involving the Koopman operator^{26–29} and is conceptually similar to the analytical definition presented by Froyland.²⁰ The method here could be thought of as a different finite-dimensional approximation of this overarching problem and similar to the algorithms implemented in GAIO that approximate the Perron-Frobenius operator using indicator functions²³ or the spectral-collocation method presented in Refs. 30 and 31.

Our approach allows for more freedom in the choice of basis functions and is compatible with basis sets comprised of indicator functions, (piecewise) polynomials,³² or “mesh-free” radial basis functions (RBFs). Although any of these choices could produce useful results, there are some advantages to choose basis sets other than indicator functions. Intuitively, this results in the same choice that arises when deciding between a spectral method,^{33,34} which typically approximates a linear operator using a set of globally supported set of basis functions, and a finite-volume method,³⁵ which uses compactly supported functions instead. Although either choice can produce accurate results, spectral methods typically converge more rapidly than finite-volume methods provided that the functions of interest are smooth,^{33–36} and from “data rich” examples and pre-existing efforts using transfer operator methods,^{18,21,22,37} this appears to be the case for the functions that define coherent pairs in the applications of interest to us. The practical benefit of a higher convergence rate is that effective approximations of coherent pairs can be obtained with fewer basis functions, and hence, fewer data points, which implies our approach is well suited for the “data poor” regime that often occurs experimentally.

The remainder of the article is outlined as follows: In Sec. II, we give a definition of a coherent set in terms of a solution to a data-driven optimization problem. In Sec. III, we consider the infinite data limit, where this problem can be recast into one involving the Koopman operator. As a result, methods such as Generalized Laplace Analysis^{28,29,38} or Extended Dynamic Mode Decomposition (Extended DMD)³² could be re-tasked to compute coherent sets. Furthermore, this limit makes the connection between our approach and the analytical definition presented by Froyland²⁰ more clear. In

Sec. IV, we apply our method to three examples: the double gyre, the Bickley Jet, and numerically simulated drifters in the Sulu Sea, in order to demonstrate that the approach is effective in practice. Finally in Sec. V, we present some brief concluding remarks.

II. A DEFINITION OF A COHERENT PAIR

In this section, we construct the optimization problem whose solution defines our pairs of coherent sets. As we will demonstrate in Sec. III, this approach is conceptually equivalent to the one presented analytically by Froyland²⁰ and implemented using GAIO.²³ As a result, there will be many similarities between what was done in Refs. 17–19 and what we do here; indeed, the manipulations that follow are motivated by Refs. 17–20. The key difference is that those approaches are tailored to use pre-specified functions such as indicator functions, while our approach is compatible with any reasonable basis set whose span contains the constant function.

A. The intuitive problem

We assume we are given a collection of M drifters, whose evolution is completely determined by the velocity field of some underlying flow, that are initially contained in some domain, $\Omega_X \subseteq \Omega \subset \mathbb{R}^N$, at time n , but migrate to another domain, $\Omega_Y \subseteq \Omega \subset \mathbb{R}^N$, at time $n + 1$. These pairs of positions are collected into the set $\{(\mathbf{x}_m, \mathbf{y}_m)\}_{m=1}^M$, where \mathbf{x}_m is the position of the m -th drifter at time n and \mathbf{y}_m is the position of that drifter at time $n + 1$. Our objective is to partition these drifters into two sets— X_1 and X_2 at time n and Y_1 and Y_2 at time $n + 1$ —based on their *physical positions* at times n and $n + 1$, respectively. To do this, we define a pair of functions, $f_X : \Omega_X \rightarrow \pm 1$ and $f_Y : \Omega_Y \rightarrow \pm 1$; at time n , the sign of f_X determines whether a point is in X_1 or X_2 , and at time $n + 1$, the sign of f_Y is used to assign the data points to either Y_1 or Y_2 .

As shown in Ref. 20, the functions f_X and f_Y will only identify useful coherent sets if the flow that maps \mathbf{x}_m to \mathbf{y}_m is stochastic. When this mapping is deterministic, one can find a pair of functions such that $g(f_X, f_Y) = 1$ for any admissible partition of Ω_X and Ω_Y simply by choosing Y_1 to be the image of the set X_1 and defining f_X and f_Y appropriately. To produce a pair of *distinguished* coherent sets, some stochasticity is required. Because we will approximate f_X and f_Y using relatively small numbers of basis functions, this required “noise” is often created implicitly via our choice of basis functions. However, to ensure that the underlying system appears to be stochastic, we also add explicit but small perturbations to both \mathbf{x}_m and \mathbf{y}_m . For the problems we will discuss, the choice of the functions used to approximate f_X and f_Y appears to have a larger impact on the resulting coherent sets than the externally added noise, but this may not always be the case if a sufficiently large number of basis functions are used.

Intuitively, one wants to choose a pair of coherent sets in a way that minimizes the “leakage” that occurs over a finite interval in time, or equivalently, maximizes the number of points that remain within one of the two sets. Because the

signs of the functions f_X and f_Y determine set membership, this intuitive goal can be achieved by choosing f_X and f_Y to maximize

$$g(f_X, f_Y) = \frac{1}{M} \sum_{m=1}^M f_X(\mathbf{x}_m) f_Y(\mathbf{y}_m), \tag{1}$$

where $g(f_X, f_Y) = 1$ if no drifters switch sets.

Without additional constraints, a global maximum can be obtained trivially by assigning *all* the data points to one set or the other. To force the algorithm to partition the data into two *nonempty sets*, we include another pair of constraints that specify the relative sizes of X_1 and X_2 and Y_1 and Y_2 . In particular, we require that

$$\frac{1}{M} \sum_{m=1}^M f_X(\mathbf{x}_m) = \varepsilon_X, \quad \frac{1}{M} \sum_{m=1}^M f_Y(\mathbf{y}_m) = \varepsilon_Y, \tag{2}$$

where ε_X and ε_Y are two constants that determine the difference in the number of elements in X_1 and X_2 and Y_1 and Y_2 , respectively. To obtain two sets of equal size, we set $\varepsilon_X = \varepsilon_Y = 0$, but it is often advantageous to allow ε_X and ε_Y to vary as not all systems can be (or should be) decomposed into two sets of equal size.

However, even (1) and (2) together is not sufficient to uniquely define f_X and f_Y . Indeed, there are either no feasible solutions (e.g., M is odd and $\varepsilon_X = \varepsilon_Y = 0$) or many optimal solutions (e.g., choose $f_X(\mathbf{x}_m) = f_Y(\mathbf{y}_m)$ for all m) when M is finite. As a result, further alterations to this intuitive problem are required if a pair of distinguished coherent sets are to be identified.

B. A finite dimensional approximation

In particular, we will modify the set of admissible f_X and f_Y . First, we relax the constraint that $f_X : \Omega_X \rightarrow \pm 1$ and $f_Y : \Omega_Y \rightarrow \pm 1$ and allow $f_X : \Omega_X \rightarrow \mathbb{R}$ and $f_Y : \Omega_Y \rightarrow \mathbb{R}$. Next, we approximate f_X and f_Y with functions that lie in the subspace spanned by two sets of basis functions that we denote as $\{\psi_k\}_{k=1}^{K_X}$ and $\{\tilde{\psi}_k\}_{k=1}^{K_Y}$ for f_X and f_Y , respectively.

In the discussion that follows, we assume that the first elements in each set are the relevant constant functions, $\psi_1(\mathbf{x}) = 1$ and $\tilde{\psi}_1(\mathbf{y}) = 1$. This ordering is helpful because it will create a block structure in our finite dimensional approximation that makes it easy to show that the constant function would be a solution to our relaxed optimization problem if some constraints were relaxed. Furthermore, provided the constant function lies in the span of the basis sets provided, one can always create such a set by “rearranging” the basis functions. Next, we define the vector-valued functions

$$\psi_X(\mathbf{x}) = \begin{bmatrix} \psi_1(\mathbf{x}) = 1 \\ \psi_2(\mathbf{x}) \\ \vdots \\ \psi_{K_X}(\mathbf{x}) \end{bmatrix}, \quad \psi_Y(\mathbf{y}) = \begin{bmatrix} \tilde{\psi}_1(\mathbf{y}) = 1 \\ \tilde{\psi}_2(\mathbf{y}) \\ \vdots \\ \tilde{\psi}_{K_Y}(\mathbf{y}) \end{bmatrix}, \tag{3}$$

which allows our finite-dimensional approximations of f_X and f_Y to be written as

$$f_X = \sum_{k=1}^{K_X} a_k \psi_k = \boldsymbol{\psi}_X^T \mathbf{a}, \quad f_Y = \sum_{k=1}^{K_Y} \tilde{a}_k \tilde{\psi}_k = \tilde{\boldsymbol{\psi}}_Y^T \tilde{\mathbf{a}}, \tag{4}$$

given two vectors of coefficients \mathbf{a} and $\tilde{\mathbf{a}}$. Because the maximum values of $|f_X|$ and $|f_Y|$ are no longer bounded, we include two additional constraints

$$\frac{1}{M} \sum_{m=1}^M |f_X(\mathbf{x}_m)|^2 = \frac{1}{M} \sum_{m=1}^M |f_Y(\mathbf{y}_m)|^2 = 1 \tag{5}$$

to impose an overall scaling on both functions. In all that follows, we will assume that f_X and f_Y are smooth functions, and therefore, can be accurately approximated even if K_X and K_Y are relatively small.

With this approximation, the objective function, (1), is

$$g(f_X, f_Y) = \mathbf{a}^T \left(\frac{1}{M} \sum_{m=1}^M \psi_X(\mathbf{x}_m) \tilde{\psi}_Y^T(\mathbf{y}_m) \right) \tilde{\mathbf{a}} = \mathbf{a}^T \mathbf{A} \tilde{\mathbf{a}}. \tag{6}$$

Similarly, the constraints, (2) and (5), are

$$\frac{1}{M} \sum_{m=1}^M \psi_1(\mathbf{x}_m) f_X(\mathbf{x}_m) = \mathbf{e}_1^T \mathbf{G}_X \mathbf{a} = \varepsilon_X, \tag{7a}$$

$$\frac{1}{M} \sum_{m=1}^M \tilde{\psi}_1(\mathbf{y}_m) f_Y(\mathbf{y}_m) = \tilde{\mathbf{e}}_1^T \mathbf{G}_Y \tilde{\mathbf{a}} = \varepsilon_Y, \tag{7b}$$

$$\frac{1}{M} \sum_{m=1}^M f_X(\mathbf{x}_m) f_X(\mathbf{x}_m) = \mathbf{a}^T \mathbf{G}_X \mathbf{a} = 1, \tag{7c}$$

$$\frac{1}{M} \sum_{m=1}^M f_Y(\mathbf{y}_m) f_Y(\mathbf{y}_m) = \tilde{\mathbf{a}}^T \mathbf{G}_Y \tilde{\mathbf{a}} = 1, \tag{7d}$$

where \mathbf{e}_1 and $\tilde{\mathbf{e}}_1$ are the first unit vectors in \mathbb{R}^{K_X} and \mathbb{R}^{K_Y} , respectively, and

$$\mathbf{G}_X \triangleq \frac{1}{M} \sum_{m=1}^M \boldsymbol{\psi}(\mathbf{x}_m) \boldsymbol{\psi}(\mathbf{x}_m)^T, \tag{8a}$$

$$\mathbf{G}_Y \triangleq \frac{1}{M} \sum_{m=1}^M \tilde{\boldsymbol{\psi}}(\mathbf{y}_m) \tilde{\boldsymbol{\psi}}(\mathbf{y}_m)^T, \tag{8b}$$

$$\mathbf{A} \triangleq \frac{1}{M} \sum_{m=1}^M \boldsymbol{\psi}(\mathbf{x}_m) \tilde{\boldsymbol{\psi}}(\mathbf{y}_m)^T. \tag{8c}$$

Note that our choice of $\psi_1 = 1$ and $\tilde{\psi}_1 = 1$ was used in (7a) and (7b). With this notation, the relaxed, finite-dimensional optimization problem is

$$\max_{\mathbf{a}, \tilde{\mathbf{a}}} \mathbf{a}^T \mathbf{A} \tilde{\mathbf{a}} \tag{9a}$$

$$\text{subject to } \mathbf{e}_1^T \mathbf{G}_X \mathbf{a} = \varepsilon_X, \tag{9b}$$

$$\tilde{\mathbf{e}}_1^T \mathbf{G}_Y \tilde{\mathbf{a}} = \varepsilon_Y, \tag{9c}$$

$$\mathbf{a}^T \mathbf{G}_X \mathbf{a} = \tilde{\mathbf{a}}^T \mathbf{G}_Y \tilde{\mathbf{a}} = 1. \tag{9d}$$

A schematic of (9) is given in Fig. 1. In short, the objective is to choose f_X and f_Y to maximize the number of data points where $f_X(\mathbf{x}_m)$ and $f_Y(\mathbf{y}_m)$ have the same sign, which is

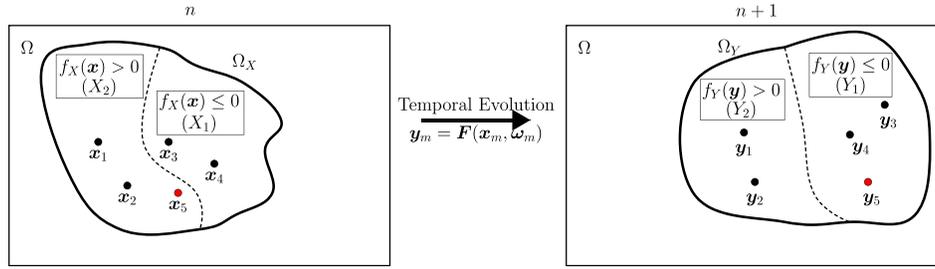


FIG. 1. A “cartoon” of the coherent set definition in this article. Given data from a discrete time dynamical system whose evolution operator at time n is F , where ω_m represents the “noise” added to the system, our objective is to identify two functions, $f_X : \Omega_X \rightarrow \mathbb{R}$ and $f_Y : \Omega_Y \rightarrow \mathbb{R}$, whose signs will be used to partition Ω_X and Ω_Y into X_1 and X_2 or Y_1 and Y_2 , respectively. These functions are determined by solving the optimization problem in (9). Intuitively, the computed functions minimize the number of “mis-classified” points, such as x_5 , that are assigned to the two different sets, X_1 or X_2 and Y_1 or Y_2 , at times n and $n + 1$. In general, $\Omega_Y \neq \Omega_X$, so the basis functions used to approximate f_X and f_Y could (and typically should) differ.

equivalent to minimizing the number of points that switch sets. The constraints are required to ensure that two *non-empty sets* are identified and impose an overall scaling on the functions. As written, (9) is a quadratically constrained quadratic program, which can be solved using specialized numerical routines.³⁹ However, we will show that this particular problem can also be solved using the SVD.

C. Computing coherent sets in practice

In this section, we will show that (9) can be solved using the SVD. The motivation for what follows is more mathematical than physical and is inspired by the results of Froyland.²⁰ Specifically, we note that

- (1) The pair of unit vectors, u and v , that maximize the quantity $u^T A v$ are the left and right singular vectors of A with the largest singular value, which we refer to as u_1 and v_1 .
- (2) With the addition of the constraints $u_1^T u = \varepsilon_X$ and $v_1^T v = \varepsilon_Y$, the optimal solution becomes $u = \varepsilon_X u_1 + \sqrt{1 - \varepsilon_X^2} u_2$ and $v = \varepsilon_Y v_1 + \sqrt{1 - \varepsilon_Y^2} v_2$, where u_2 and v_2 are the singular vectors associated with the *second* largest singular value.

The main difference between this problem, which can be solved using the SVD, and (9) is that the constraints could be written in terms of the standard Euclidean inner product, while (9) has constraints that are written in terms of weighted inner products.

Therefore, the first step is to transform our coordinates such that the constraints in (9) can be expressed in terms of “unweighted” inner products like in our model problem. To do this, we use the Cholesky Decomposition and let

$$G_X = L_X L_X^T, \quad G_Y = L_Y L_Y^T. \tag{10}$$

For this decomposition to exist, G_X and G_Y must be full rank, which will be the case if the sets of functions used to represent f_X and f_Y form a basis for a subspace of $L^2(\Omega_X, \rho)$ and $L^2(\Omega_Y, \nu)$, where ρ and ν are the spatial distribution of the x_m and y_m , respectively. If we define $b = L_X^T a$ and $\tilde{b} = L_Y^T \tilde{a}$, then the constraints simplify to

$$e_1^T b = \hat{\varepsilon}_X, \tag{11a}$$

$$\tilde{e}_1^T \tilde{b} = \hat{\varepsilon}_Y, \tag{11b}$$

$$b^T b = \tilde{b}^T \tilde{b} = 1, \tag{11c}$$

where $\hat{\varepsilon}_X = \varepsilon_X / L_X^{(11)}$, $\hat{\varepsilon}_Y = \varepsilon_Y / L_Y^{(11)}$, $L_X^{(11)} \in \mathbb{R}$ denote the element in the first row and column of L_X . These terms appear because L_X (or L_Y) is lower-triangular, and therefore $e_1^T L_X = L_X^{(11)} e_1^T$. We also rewrite the objective function and set $a^T A \tilde{a} = b^T \hat{A} \tilde{b}$, where

$$\hat{A} \triangleq L_X^{-1} A L_Y^{-T}. \tag{12}$$

This results in a transformed system of equations

$$\max_{b, \tilde{b}} \quad b^T \hat{A} \tilde{b}, \tag{13a}$$

$$\text{subject to:} \quad e_1^T b = \hat{\varepsilon}_X, \tag{13b}$$

$$\tilde{e}_1^T \tilde{b} = \hat{\varepsilon}_Y, \tag{13c}$$

$$b^T b = \tilde{b}^T \tilde{b} = 1, \tag{13d}$$

which is formally equivalent to our model problem because e_1 and \tilde{e}_1 are the left and right singular vectors of \hat{A} with $\sigma_1 = 1$. This is a result of our choice of $\psi_1 = \tilde{\psi}_1 = 1$ and is simple but tedious to show (see Appendix).

Therefore, the solution to (13) is of the form

$$b = \hat{\varepsilon}_X e_1 + \sqrt{1 - \hat{\varepsilon}_X^2} u_2, \tag{14a}$$

$$\tilde{b} = \hat{\varepsilon}_Y \tilde{e}_1 + \sqrt{1 - \hat{\varepsilon}_Y^2} v_2, \tag{14b}$$

where u_2 and v_2 are the left and right singular vectors associated with σ_2 , the largest singular value not equal to $\sigma_1 = 1$. In what follows, we will assume that $1 = \sigma_1 > \sigma_2$. In the unrelaxed problem, the maximum value of the objective function is 1, but in the relaxed problem, it is possible to find solutions associated with larger values. These solutions are associated with values of $\sigma_2 > 1$ (and do not appear to produce useful pairs of sets). Instead, we treat the magnitude of the largest singular value as a “sanity check” on the procedure. If $\sigma_2 > 1$, then our relaxed procedure is identifying solutions that exceed the theoretical maximum of the original problem and therefore is not a reliable surrogate for the

original problem. In practice, we have found that reducing the number of basis functions used to approximate both f_X and f_Y alleviates this issue.

Once \mathbf{b} and $\tilde{\mathbf{b}}$ have been computed, we let $\mathbf{a} = \mathbf{L}_X^{-T}\mathbf{b}$ and $\tilde{\mathbf{a}} = \mathbf{L}_Y^{-T}\tilde{\mathbf{b}}$ and approximate f_X and f_Y at any desired points using (4). The final step in the procedure is the partition of Ω_X and Ω_Y using the numerically computed f_X and f_Y . We define

$$X_1 = \{\mathbf{x} \in \Omega_X : f_X(\mathbf{x}) \leq 0\}, \tag{15a}$$

$$Y_1 = \{\mathbf{y} \in \Omega_Y : f_Y(\mathbf{y}) \leq 0\}, \tag{15b}$$

and let X_2 and Y_2 be their complements (or, equivalently, the subset where $f_X, f_Y > 0$). The values of ε_X and ε_Y effectively add a constant offset to both f_X and f_Y . As a result, ε_X and ε_Y can be determined after the fact, and following Ref. 19, we will choose them so that the computed f_X and f_Y maximize the fraction of consistently classified points (i.e., if $\mathbf{x}_m \in X_1$ then $\mathbf{y}_m \in Y_1$) in the pairs of sets (X_1, Y_1) and (X_2, Y_2) . Although we do not place any explicit constraints on the values of ε_X or ε_Y , we typically require that neither ε_X nor ε_Y can be so large or small that either X_1 or X_2 (or Y_1 and Y_2) contains a negligible number of data points.

D. Algorithm summary

In practice, this algorithm requires the user to provide three quantities: (i) a data set of snapshot pairs, $\{(\mathbf{x}_m, \mathbf{y}_m)\}_{m=1}^M$, (ii) two sets of basis functions that comprise the vector-valued functions ψ_X and ψ_Y , and (iii) the “noise” that will be added to the data. The first two quantities are important if this method is to perform well, but because they are highly problem dependent, we will defer the discussion of these choices until Sec. IV, where we apply the method to our example problems. The addition of noise will, in principle, affect the resulting sets, but in practice, appears to have a smaller impact than the data and basis functions provided the noise chosen is not too large. Given these quantities, the coherent sets are computed as follows:

- (1) (Optional) Augment the existing data set with noise by looping through the data multiple times and randomly perturbing \mathbf{x}_m and \mathbf{y}_m . These new data pairs are then added to the existing set of data and will be used in the steps that follow. In practice, this step is often unnecessary; we effectively inject noise into the problem by using a limited number of basis functions in our approximation of f_X and f_Y .
- (2) Compute the matrices in (8), their Cholesky decompositions in (10), and the matrix $\hat{\mathbf{A}}$ in (12).
- (3) Using the SVD, let $\hat{\mathbf{A}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$.
- (4) As a sanity check, examine σ_2 , which is the largest singular value that is not unity. In practice, we iterate through steps 1–3 and select the largest basis where $\sigma_2 < 1$.
- (5) Choose values of ε_X and ε_Y and compute \mathbf{b} and $\tilde{\mathbf{b}}$ using (14). To obtain sets with less “leakage,” we choose these values to minimize the fraction of misclassified points, which is similar to the concept of *coherence* in Refs. 18

and 19. In practice, simply letting $\varepsilon_X = \varepsilon_Y = 0$ is sufficient in many applications.

- (6) Compute $\mathbf{a} = \mathbf{L}_X^{-T}\mathbf{b}$ and $\tilde{\mathbf{a}} = \mathbf{L}_Y^{-T}\tilde{\mathbf{b}}$, which are the solutions of the original relaxed, finite-dimensional optimization problem.
- (7) Finally, compute the value of f_X or f_Y at any desired points using (4) and partition the domain based on the sign of f_X .

If more than two coherent sets are desired, we repeat the procedure outlined above in a recursive fashion using the data in X_1 and Y_1 and the data in X_2 and Y_2 separately. Similar to the work of Ma and Boltz,³⁷ this results in a larger number of coherent sets that can capture finer spatial features. In practice, we terminate this iteration procedure if more than 5% of the data in any pair of sets leaks out during the interval of interest so that all of the resulting sets will, visually, appear to be coherent.

The algorithm presented here runs in $\mathcal{O}(K^2 \max(K, M))$ time, where $K = \max(K_X, K_Y)$. This cost is either determined by the need to assemble $\mathbf{G}_X, \mathbf{G}_Y$, and \mathbf{A} , which is an $\mathcal{O}(K^2 M)$ computation, or to decomposing $\hat{\mathbf{A}}$, which is an $\mathcal{O}(K^3)$ operation. Assuming $K_X \sim K_Y$, this is the same asymptotic complexity as the algorithms used in GAIO if they are naively implemented. However, given equal numbers of basis functions, our approach will be *slower* than GAIO because it uses tree-like data structures to efficiently construct the needed matrices, which our approach is unable to do and avoids the additional Cholesky factorizations that our procedure requires. As we shall demonstrate shortly, our approach can often identify useful coherent pairs using far fewer basis functions, which in practice, helps to offset the larger cost-per-basis-function associated with this method.

III. CONNECTIONS TO THE KOOPMAN OPERATOR

The algorithm presented in Sec. II is both conceptually and mathematically related to the approach presented in Refs. 18–20; indeed, the primary difference between the approaches is that we can use a “richer” set of basis functions to represent f_X and f_Y . In this section, we examine the “infinite data” limit, which is the limit where our approach can be compared to these transfer operator-based methods.

In all that follows, we assume that our data set $\{(\mathbf{x}_m, \mathbf{y}_m)\}_{m=1}^M$ is constructed by randomly choosing initial conditions, \mathbf{x}_m , from the distribution, ρ . As before, \mathbf{y}_m is the location of the m -th drifter at time $n + 1$, and ν is the new distribution at that time. If the evolution operator from time n to $n + 1$ is \mathbf{F} , then $\mathbf{y}_m = \mathbf{F}(\mathbf{x}_m, \omega_m)$, where ω_m accounts for both the noise that is artificially added to the flow map and any stochasticity that naturally exists in the flow.

In the limit as $M \rightarrow \infty$, the ij -th element of \mathbf{G}_X is almost surely

$$\lim_{M \rightarrow \infty} \mathbf{G}_X^{(ij)} = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \psi_i(\mathbf{x}_m) \psi_j(\mathbf{x}_m) = \langle \psi_i, \psi_j \rangle_\rho, \tag{16}$$

where $\langle f, g \rangle_\rho = \int_{\Omega_X} f(\mathbf{x})g(\mathbf{x})\rho(\mathbf{x})dx$. This argument follows directly from the law of large numbers. Similarly,

$$\lim_{M \rightarrow \infty} \mathbf{G}_Y^{(ij)} = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \tilde{\psi}_i(\mathbf{y}_m) \tilde{\psi}_j(\mathbf{y}_m) = \langle \tilde{\psi}_i, \tilde{\psi}_j \rangle_\nu, \quad (17)$$

where $\langle f, g \rangle_\nu = \int_{\Omega_Y} f(\mathbf{y})g(\mathbf{y})\nu(\mathbf{y}) d\mathbf{y}$. In this limit, both \mathbf{G}_X and \mathbf{G}_Y are Gram matrices, where each element is an inner product of basis functions weighted by the density of the data.

Unlike \mathbf{G}_X and \mathbf{G}_Y , the ij -th element of \mathbf{A} depends upon both the randomly selected initial conditions, \mathbf{x}_m , and their images, \mathbf{y}_m , which are affected by the stochasticity in the dynamics. Assuming that the \mathbf{x}_m and ω_m are chosen independently

$$\begin{aligned} \lim_{M \rightarrow \infty} \mathbf{A}^{(ij)} &= \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \psi_i(\mathbf{x}_m) \tilde{\psi}_j(\mathbf{y}_m) \\ &= \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \psi_i(\mathbf{x}_m) \tilde{\psi}_j(\mathbf{F}(\mathbf{x}_m, \omega_m)) \\ &= \int_{\Omega_X} \mathbb{E} \left[\psi_i(\mathbf{x}) \tilde{\psi}_j(\mathbf{F}(\mathbf{x}_m)) \rho(\mathbf{x}) \right] d\mathbf{x} \\ &= \left\langle \psi_i, \mathbb{E} \left[\tilde{\psi}_j \circ \mathbf{F} \right] \right\rangle_\rho, \end{aligned} \quad (18)$$

where \mathbb{E} denotes the expected value over the stochasticity in the dynamics and represents the integral taken over the probability space.

In this formulation, the connection to the Koopman operator appears in (18). The Koopman operator was originally defined for Hamiltonian systems^{26,27} but in recent years has also been applied to dissipative systems^{28,32,38,40} and those with stochastic dynamics;^{28,32} this latter formulation is most relevant here. In this application, the Koopman operator, which we denote as \mathcal{K} , is defined for a discrete-time *Markov process* with the evolution operator \mathbf{F} . The appeal of studying the Koopman operator instead of \mathbf{F} is that \mathcal{K} is linear even when \mathbf{F} is nonlinear. However, the Koopman operator acts on scalar observables, such as the ψ_k or $\tilde{\psi}_k$, which map state space to scalars, and is infinite dimensional even when \mathbf{F} is finite dimensional.

For the observable $\tilde{\psi} : \Omega_Y \rightarrow \mathbb{R}$, the action of the Koopman operator is

$$\mathcal{K}\tilde{\psi} = \mathbb{E}[\tilde{\psi} \circ \mathbf{F}], \quad (19)$$

where \mathbb{E} is the expectation over the stochastic dynamics in \mathbf{F} and $\mathcal{K}\tilde{\psi} : \Omega_X \rightarrow \mathbb{R}$ is another function defined on a different domain. With the Koopman operator, the ij -th element of \mathbf{A} can be written succinctly as

$$\mathbf{A}^{(ij)} = \langle \psi_i, \mathcal{K}\tilde{\psi}_j \rangle_\rho, \quad (20)$$

and due to the linearity of the Koopman operator, the objective function can be written as

$$g(f_X, f_Y) = \langle f_X, \mathcal{K}f_Y \rangle_\rho, \quad (21)$$

which leads to the optimization problem

$$\max_{f_X, f_Y} \langle f_X, \mathcal{K}f_Y \rangle_\rho \quad (22a)$$

$$\text{subject to : } \langle 1, f_X \rangle_\rho = \varepsilon_X, \quad (22b)$$

$$\langle 1, f_Y \rangle_\nu = \varepsilon_Y, \quad (22c)$$

$$\langle f_X, f_X \rangle_\rho = \langle f_Y, f_Y \rangle_\nu = 1, \quad (22d)$$

that we would solve given an infinite amount of data and a complete set of basis functions.

The benefit of this formulation is that it makes the similarities between our method and the one presented in Ref. 20 clear. Because the Koopman operator, \mathcal{K} , is the adjoint of the (modified) Perron-Frobenius operator, \mathcal{L} , used there, the objective function can either be written as $\langle f_X, \mathcal{K}f_Y \rangle_\rho$ or $\langle \mathcal{L}f_X, f_Y \rangle_\nu$. The latter expression is of the same form as the objective function used by Froyland²⁰ and could be equivalent provided the noise added to the system is chosen appropriately.

Furthermore, the problem in (9) can also be derived by approximating the Koopman operator using extended dynamic mode decomposition³² with the set of snapshot pairs $\{(\mathbf{x}_m, \mathbf{y}_m)\}_{m=1}^M$, the basis functions ψ_k and $\tilde{\psi}_k$, and using Monte-Carlo integration to approximate any needed inner products. As a result and assuming f_X and f_Y are smooth, our procedure will converge to the true solution at a rate of either $\mathcal{O}(\sqrt{M})$, if the error is dominated by errors in the integrals, or at a rate determined by set of basis functions used to approximate f_X and f_Y . In principle, however, any method that can approximate the action of the Koopman operator directly from data, such as Generalized Laplace Analysis,²⁹ could also be used to compute coherent sets. This would lead to a different optimization problem and a different rate of convergence.

IV. EXAMPLE APPLICATIONS

In this section, we consider three examples that demonstrate the efficacy of our method. The first is the double gyre, which is defined on a fixed domain, and frequently used as a test problem for coherent set identification. The purpose of this example is to demonstrate that the approach described here produces coherent sets that are similar to the sets produced by GAIO using an approximation of the Perron-Frobenius operator. Next, we consider the Bickley Jet, which is an idealized but more realistic problem where the data are not initially confined to some trapping region, and therefore, mesh-free approximations of f_X and f_Y become necessary. Next, we consider the example of numerically simulated “drifters” in the Sulu Sea, which is a realistic example of how we envision this technique being used in practice. In this example, our objective is to identify an eddy that is already known to exist in the time frame of the simulation.

A. Choosing the basis functions

As mentioned previously, one important facet of this procedure is the choice of the basis functions that are the building blocks for ψ_X and ψ_Y . In each of these problems, we use a basis set of thin-plate splines, which are functions of the form

$$\psi_k(\mathbf{x}) = r^2 \log(r), \quad \text{where } r = \|\mathbf{x} - \xi_k\|, \quad (23a)$$

$$\tilde{\psi}_k(\mathbf{y}) = \tilde{r}^2 \log(\tilde{r}), \quad \text{where } \tilde{r} = \| \mathbf{y} - \tilde{\xi}_k \|, \quad (23b)$$

where ξ_k is the k -th RBF center, which is a vector in \mathbb{R}^2 that defines the center-position of the thin plate spline. We also define the basis functions used to construct f_Y in a similar manner, but call the associated centers $\tilde{\xi}_k$.

Thin plate splines are a special case of polyharmonic splines that are tailored for problems in \mathbb{R}^2 and commonly used for the interpolation of scattered data.⁴¹ Although they are not compactly supported, these functions have two useful properties: (i) They do not require the scaling parameter that many other radial basis functions do and (ii) they do not require a computational mesh to be defined.^{36,42,43} The ξ_k and $\tilde{\xi}_k$ are chosen by applying k -means clustering⁴⁴ to the collection of x_m and y_m snapshots, respectively. k -means clustering partitions a set of data into k -sets, which are chosen to minimize the total distance between the points and the mean of the set they are assigned to. We use the set of means that result from this procedure as the ξ_k and $\tilde{\xi}_k$, respectively.

To determine the number of basis functions, K_X and K_Y , we first choose a ‘‘conservative’’ pair of values, say, $K = K_X = K_Y = 5$. Next, we compute the leading singular value of the \hat{A} associated with the basis sets generated by this value of K . If the leading singular value is one, then we increment K and repeat the process until this constraint fails to hold. The results in this section are from the largest values of K_X and K_Y that did not violate our sanity check, which as a rule of thumb, corresponds to between 5 and 20 data points per basis function. This procedure is *ad hoc* but appears to produce a useful set of basis functions for the examples presented in this article.

This is, of course, neither the only possible choice of basis functions nor do we claim it is in any way optimal. However, the benefit of using the basis elements associated with mesh-free methods, such as the thin plate splines, is that they can be applied to problems on domains that are not simple rectangles, which makes them suitable for a wide range of applications. Similar to GAIO, this allows us to apply the same procedure to all of the examples that follow despite the fact only one of them is defined on a fixed domain.

B. The double gyre

Our first example is the double gyre, whose governing equations are

$$\dot{x} = -\pi A \sin(\pi h(x, t)) \cos(\pi y), \quad (24a)$$

$$\dot{y} = \pi A \cos(\pi h(x, t)) \sin(\pi y) \frac{\partial h}{\partial x}, \quad (24b)$$

where $h(x, t) = \epsilon \sin(\omega t)x^2 + (1 - 2\epsilon \sin(\omega t))x$ with $\epsilon = 0.25$, $\omega = 2\pi$, and $A = 0.25$. In these equations, $x \in [0, 2]$ and $y \in [0, 1]$. The double gyre with these parameters is a frequently used test case for coherent structure computations. See, for example, Refs. 6, 21, and 37, which compute coherent sets (albeit with slightly different definitions) for this problem and parameters.

The purpose of this example is to demonstrate that the computational procedure outlined in Sec. II produces

coherent sets that are similar to those produced using the definition in Refs. 17,18, 20, and 21, from a more limited amount of data. For the purposes of comparison, Fig. 2 shows the equivalent of f_X identified by GAIO, which uses 262 144 indicator functions and a total of 104 857 600 data pairs (i.e., 400 points per function). Because it uses indicator functions, a large basis set and, hence, a large amount of data are required if the resulting functions are to look smooth. By using tree-like data structures, this computation can be performed quickly even with hundreds of millions of data points.²³ However, in applications where experimental rather than numerical data are being used, obtaining such a large set may not be possible.

To highlight the performance of the method, we apply it with 51, 251, and 501 basis functions (e.g., 50 thin plate splines and the constant function) using 1000, 5000, and 10 000 data pairs, respectively. The data at the initial time (i.e., the x_m) are chosen by randomly selecting initial conditions from a uniform distribution on state space. Because the governing equations are discrete, we make the system stochastic by adding noise all the x_m and y_m . In this example, we make 20 copies of each of our data pairs and perturb the data by adding a random vector chosen from a normal distribution with a standard deviation of 10^{-3} ; with the basis sets we will use, neither the number of copies nor the precise nature of the noise will have a qualitative impact on the resulting functions. As a result, the values of M in our computation are 20 000, 100 000, and 200 000, which accounts for these additional copies, but similar results could be obtained in the noiseless case with $M = 1000, 5000,$ and 10 000. Finally, we impose that $\epsilon_X = \epsilon_Y = 0$ in order to facilitate comparison with GAIO.

Figure 3 shows the function, f_X , obtained using the three sets of data listed above. The black line denotes the $f_X = 0$ level set, which is used to partition state space into the pair of coherent sets. As shown above, our approximations of f_X appear to be converging to a particular function as the number of basis functions and data points increases; when run with 1001 basis functions and 20 000 sets of data ($M = 4 \times 10^5$), the resulting f_X is qualitatively similar to the function obtained with 501 basis functions.

As a benchmark for our approach, Fig. 4 shows the function equivalent to f_X computed using GAIO with 64, 256, and 512 basis functions with 25 data points per basis function initialized on a uniform grid. This figure should be compared to Fig. 3, which used slightly fewer basis functions

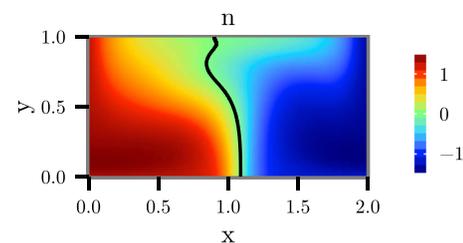


FIG. 2. The function equivalent to f_X computed using GAIO with 262 144 indicator functions (and the equivalent of 104 857 600 data pairs). The black line denotes the zero level set of both functions, which would be used to partition the domain into two coherent sets.

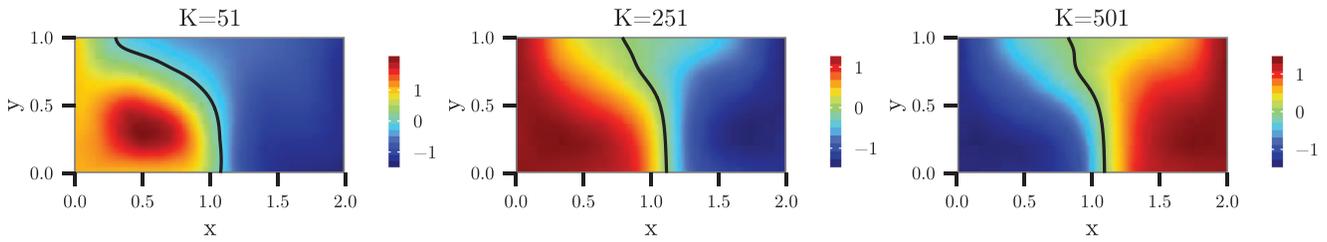


FIG. 3. The function f_X for the double gyre computed using 51, 251, and 501 basis functions (e.g., 50 thin plate splines and the constant function) with $\varepsilon_X = \varepsilon_Y = 0$. The black line indicates the zero level set, which partitions the domain into the two sets X_1 and X_2 . These results should be compared with the “true” solution in Fig. 2.

(i.e., 51, 251, and 501) with 20 randomly distributed initial conditions per basis function. As a result, each column in Figs. 3 and 4 is comparable. The most apparent difference between these two sets of results is the smoothness of f_X , which are both clearly discontinuous in Fig. 4 due to the basis set that is implicitly chosen by GAIO. More importantly for experimental applications, we obtain a solution that is *qualitatively* similar to the “true” solution with only 251 basis functions, while GAIO requires at least twice that amount.

We should note that there are quantitative differences between the coherent sets identified using our method and the ones identified by GAIO. In particular, there are small quantitative differences in the zero level sets near the point $x = 1$ and $y = 1$ for f_X and $x = 1$ and $y = 0$ for f_Y . Part of this difference is due to the noise added to the data; our approach explicitly adds normally distributed perturbations, and GAIO *implicitly* adds noise that is related to the width of each subdomain.²⁰ As a result, the part of the error due to differences in the added noise would not vanish even if the amount of data was effectively infinite. However, there is also a difference in the spaces spanned by the thin plate splines used here and the indicator functions used by GAIO; in practice, this difference has a larger impact on the resulting sets. In this example, a large number of thin plate splines would be required to capture the sudden “bulge” that occurs in f_X near the edge of the domain.

As shown here, our approach compares favorably to transfer operator based methods for coherent set identification and produces coherent pairs that are qualitatively similar to the ones identified by those methods. Our approach uses a smaller number of globally supported basis functions and appears to converge more rapidly because f_X and f_Y are smooth in this problem. As a result, useful and accurate approximations of these functions can be obtained with fewer data points than other methods may require.

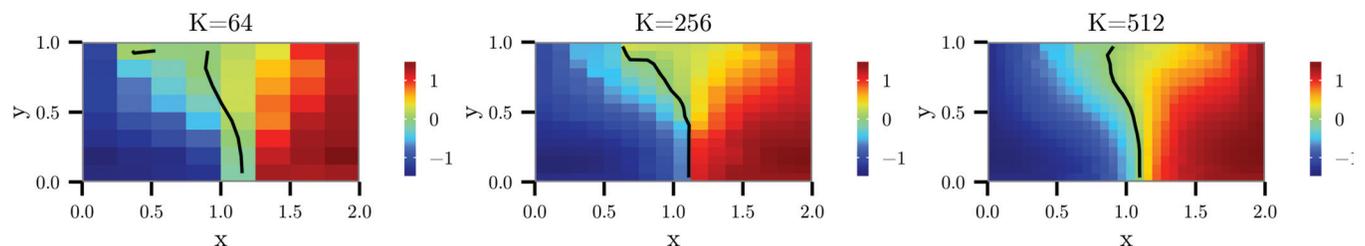


FIG. 4. The function equivalent to f_X obtained using GAIO with 64, 256, and 512 basis functions and 25 uniformly distributed data points per basis function; the images here are a benchmark for our results, which are shown in Fig. 3.

C. The Bickley Jet

In this example, we demonstrate the effectiveness of this method by computing a pair of coherent sets in the Bickley Jet flow, which is a dynamically consistent approximation of an idealized stratospheric flow.⁴⁵ We are concerned with sets that are optimal for the interval $t \in [10, 20]$ days, which was chosen so that these results may be compared with pre-existing results.²¹ This idealized system is Hamiltonian

$$\frac{\partial x}{\partial t} = -\frac{\partial \Phi}{\partial y}, \quad (25a)$$

$$\frac{\partial y}{\partial t} = \frac{\partial \Phi}{\partial x}, \quad (25b)$$

where

$$\begin{aligned} \Phi(x, y, t) = & c_3 y + U_0 L \tanh(y/L) \\ & + A_3 U_0 L \operatorname{sech}^2(y/L) \cos(k_3 x) \\ & + A_2 U_0 L \operatorname{sech}^2(y/L) \cos(k_2 x - \sigma_2 t) \\ & + A_1 U_0 L \operatorname{sech}^2(y/L) \cos(k_1 x - \sigma_1 t), \end{aligned} \quad (25c)$$

with $U_0 = 62.66$ m/s, $L = 1770$ km, $c_2 = 0.205 U_0$, $c_3 = 0.7 U_0$, $A_1 = 0.075$, $A_2 = 0.4$, $A_3 = 0.2$, $k_1 = 2/r_c$, $k_2 = 4/r_c$, $k_3 = 6/r_c$, $r_c = 6.371$, $\sigma_2 = k_2(c_2 - c_3)$, and $\sigma_1 = \frac{1+\sqrt{5}}{2} \sigma_2$. See Rypina *et al.*⁴⁵ for an explanation of these parameter values. Our initial data are 10^4 uniformly distributed on $x \in [0, 20]$ Mm and $y \in [-2.5, 2.5]$ Mm at $t = 10$ days, which we augment by making 20 copies where both x_m and y_m are perturbed randomly using numbers drawn from a normal distribution with a standard deviation of 10^{-3} . Even without noise, many initial conditions will leave this window, so $\Omega_X \neq \Omega_Y$, and choosing a different set of basis functions to

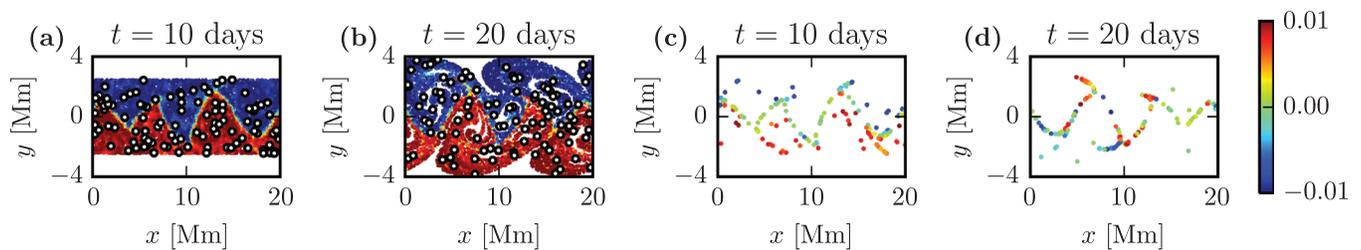


FIG. 5. (a) The 10^4 data points that comprise the data set colored by the numerically computed approximation of f_x . (b) The set of points in (a) at their new positions at $t = 20$ days; note that these points are still colored by f_x rather than f_y . (c) The set of 158 mis-classified points (i.e., those that “leaked” out over the 10 day window) at $t = 10$ days. (d) The same set of points at $t = 20$ days. The white circles in the left two images represent 100 of the 1000 ξ_k and $\tilde{\xi}_k$ that were used as the centers of the radial basis functions. Despite the complexity of the “true” coherent sets and the relatively small number of basis functions, there is little mixing between the numerically computed sets, and most of the mixing occurs on the boundary between the sets or in thin filaments extending into either side.

represent f_x and f_y is critical. We use 10^3 thin plate splines whose centers are chosen using the k -means procedure outlined at the start of this section; a subset of these locations are indicated by the white dots in Fig. 5, which makes it clear that the resulting distributions are qualitatively different at the two times.

Figure 5 shows the results obtained with these basis functions and data. For this problem, the geometry of the coherent pair is more complex, and both sets have a “sawtooth” pattern. Like before, the function in Fig. 5(a) changes rapidly in value from approximately -0.015 to 0.015 at the boundary between the coherent pair. Note that the color-scales in those images are restricted to -0.01 to 0.01 rather than the full range of values. The sets identified by partitioning the data based on the sign of f_x with 10^4 data points compare favorably with those in Ref. 21, which uses over a million data points. In particular, we partition the computational domain into the two subdomains that are located above and below the oscillating jet near $y=0$ that separates them. Furthermore, as shown in the figure, this approximation is “good enough” that only 158 of the 10^4 numerically classified points “leak” out of the set they were assigned to; similar to Refs. 19 and 21, this leakage occurs either on the boundary between sets or on thin filaments that penetrate into either side.

The difference between the Bickley Jet and the double gyre example in Sec. IV B is where the initial data were located. For the double gyre, Ω_x was a trapping region, and $\Omega_x = \Omega_y$, but that is not the case in this problem. While coherent sets were defined in Sec. II, one must also choose basis functions associated with “mesh-free” numerical methods if this computational procedure is to be viable. In this example, we used thin-plate splines and recovered a good approximation of the sets obtained via GAIO-like methods with far fewer data points.

D. Numerical drifters in the Sulu Sea

In this section, we consider a more realistic example generated by a numerical model ROMS⁴⁶ for the Philippine Archipelago.¹⁰ Similar to the work of Rypina *et al.*,¹⁰ the objective here is to use our coherent set definition to identify a mesoscale anticyclonic eddy that was present in the Sulu Sea. Our data come in the form of numerically simulated drifters that are sampled once every week. These drifters are

randomly and uniformly distributed over the computational domain. In an experiment rather than a simulation, initializing thousands of drifters is infeasible, and so the data available are truly limited in quantity. As such, we consider two cases: The first consists of 25 146 tracers randomly but uniformly distributed over the computational domain and constitutes a “data rich” example, which we will use to determine the “true” coherent sets. Then, we will reduce the amount of data to 400 uniformly but randomly distributed initial conditions, which is a more realistic amount of data, and compare the results obtained from this “data poor” set with the full data. As before, we make 20 copies of the data and add normally distributed noise with a standard deviation of 100 m to both the x_m and the y_m .

In this example, we are interested in identifying a meso-scale eddy with a 100-km radius within a much larger, 500 km by 1500 km, domain. While both the “rich” and “poor” data sets can be partitioned into a pair of coherent sets, due to the implicit constraint on the size of these sets, neither will immediately identify the eddy of interest. Therefore, it becomes necessary to iterate the procedure and to further subdivide space until the size of the coherent sets is on the same order as the eddy. As a result, we will iterate up to four times using the procedure described in Sec. II and Refs. 21 and 37. For each of the iterates, we choose ϵ_x and ϵ_y to maximize the fraction of consistently classified data pairs. We limit the range of values that the ϵ_x and ϵ_y can take on so that the resulting sets contain (roughly) the same number of points (i.e., the smaller set must contain at least 25% of the total data). This additional restriction is *ad hoc* and meant to prevent the algorithm from selecting “trivial” sets with only a handful of isolated data points. Initially, we use 250 basis functions for the full data and 120 for the reduced data set of 400 points. After every subsequent iteration, we divide the number of basis functions used in the computation by two and require that the number of basis functions is no more than 30% of the data points; as a result, later iterations are performed on smaller domains and with fewer basis functions.

In Fig. 6, we show the hierarchy of coherent sets that is optimal for the $t = 0$ to $t = 7$ days time window using the full set of the 25 146 data points available, where each of the points is “colored” by which of the 27 sets they were assigned to. Note that the number of sets is a result of the recursion procedure outlined above; more or fewer sets can

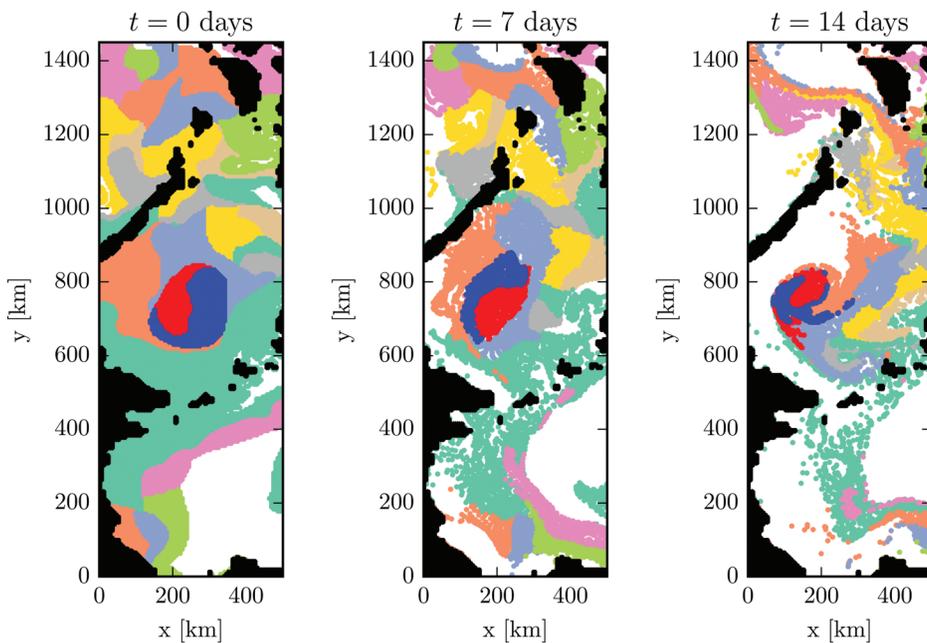


FIG. 6. The three images above show the evolution of the “complete” set of 25 146 data points at $t=0, 7,$ and 14 days. Two of the 27 coherent sets identified by the method, which correspond to the eddy of interest, are shown in red and blue; the other 25 sets, which may or may not have a physical interpretation, are indicated by the other colors. In all three images, the black points denote land. It should be noted that these coherent sets are only optimal from $t=0$ to 7 days. At $t=14$ days, this lack of optimality can be seen in the long filaments that have formed in the blue and red sets.

be generated by changing the amount of allowable “leakage,” the maximum number of recursive iterations, and the cutoff points ϵ_x and ϵ_y . There is an additional plot of the data at $t=14$ days that demonstrates that the coherent sets identified by our method remain coherent even at longer times. We should reiterate, however, that these sets are, by construction, only optimal from $t=0$ to 7 days; the results at $t=14$ days are extrapolation and not guaranteed to still be coherent at that time. In this example, the persistence of the identified coherent sets until day 14 is consistent with Rypina *et al.*¹⁰ who also found the eddy to be present over a 2-week period. Note that even at $t=0$, the data have “holes,” which are either due to the presence of land, which is indicated by the black regions, or because initial conditions at those regions leave the computational window in Fig. 6 before a week has elapsed.

The identified coherent sets are “optimal” sets but that does not necessarily mean that all of them necessarily have a simple and straightforward physical interpretation. However, one physically meaningful pair of coherent sets is indicated by the red and blue regions near $x=200$ km and $y=700$ km in the figure and corresponds to the eddy identified by Rypina *et al.*¹⁰ From $t=0$ to $t=7$, the red and blue sets move counterclockwise around each other without much stretching and folding, which is consistent with what is expected for a cyclonic motion associated with an eddy. Note however, that because the sets are not optimal at $t=14$ days, they begin to leak out of the eddy at that time forming long filaments. To summarize, with a large amount of drifter data, such as the data set generated numerically using the velocity field, this approach can produce a partition of state space that contains physically meaningful time-varying sets.

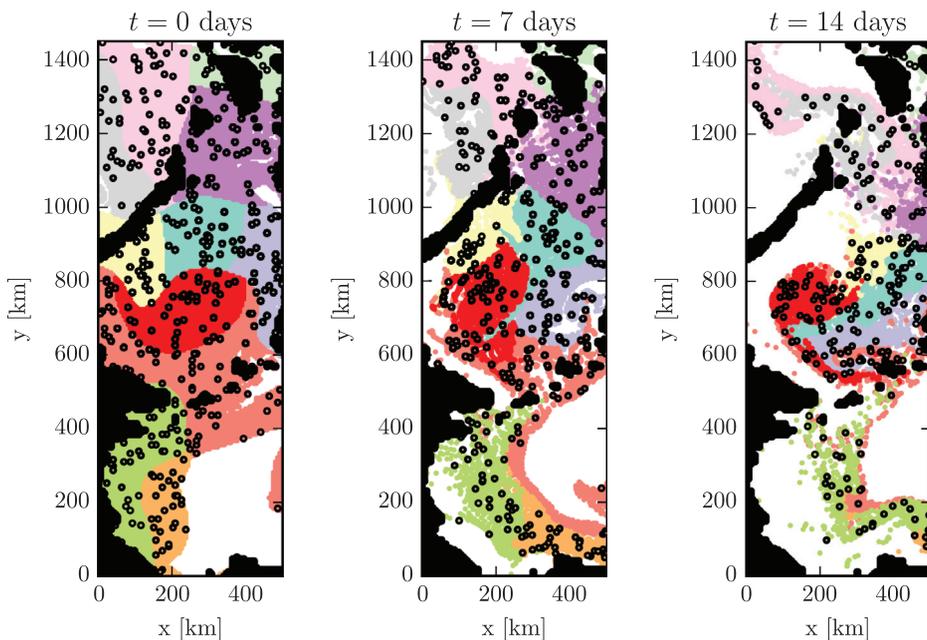


FIG. 7. This is a reproduction of Fig. 6 using only the 400 data points indicated by the black circles. As in that figure, the other 24 746 points are colored based on which coherent set they are assigned to by the approximation of f_X obtained from the indicated set of 400 points. Note that the eddy is once again identified but is now contained within a single coherent set that is shown in red.

However, the purpose of this article is to demonstrate that these results can be obtained with limited quantities of Lagrangian data, so we repeat the computation above with 400 data points instead of the full set of 25 146. We follow the same iteration procedure as before and compute additional sets by recursing up to four times provided the identified coherent sets “mis-classified” at most 5% of the data points available to each stage of the recursion procedure (e.g., the first level of the full 25 146 point data set was allowed to mis-classify up to 1257 points, but the first level of the 400 point data set is only allowed to mis-classify 20). The results of this computation are shown in Fig. 7. To aid the eye, the colored points are, once again, the complete data set that is shown in Fig. 6, where the colors denote the various coherent sets. The 400 points used in the computation are indicated by the large black circles in the figure.

As shown, the relatively small amount of data and the concomitant reduction in the number of basis functions has had an impact on the resolution and accuracy of the resulting method. Visually, the coherent sets we identify are larger in area than those that we obtain with full data because fewer iterations of the coherent set algorithm can be performed before our sets allow more than 5% of their points to escape. Nonetheless, we once again identify the eddy, which is now indicated solely by the red set. Once again, this set is only optimal from $t = 0$ to 7 days, and long filaments are again visible at $t = 14$ days.

It should be noted that whether the eddy is contained within a single coherent set or a pair is determined by the basis functions and data provided to the method. In the “data rich” example above, this subdivision occurs at the final step of the procedure, so the red and blue sets in Fig. 6 can be merged by terminating the procedure one iteration sooner. However, changes in computational parameters such as the location or number of the thin plate splines can result in this subdivision occurring before the last iterate. Due to the recursive nature of the procedure, the data are assigned to coherent sets in a “greedy” manner, and once the eddy has been subdivided, it will remain so in all future iterates. Other procedures for identifying multiple sets that do not have this limitation have been developed,^{47,48} but their integration within the framework presented here will be the focus of future work.

In this section, we considered a more realistic example: numerically simulated drifters in the Sulu sea. First, we applied our procedure to a relatively large set of data and demonstrated that, under ideal conditions, it was able to identify an eddy that is known to exist in this flow. Next, we limited the amount of data to 400 randomly chosen drifters. The cost of working in this “data poor” regime is a loss of resolution; fewer coherent sets could be identified before the amount of “leakage” grew past our threshold. In the end, however, our procedure once again identified the eddy of interest.

V. CONCLUSIONS

In this article, we presented a method for computing coherent sets that are optimal over a finite interval in time, which is conceptually related to that of Froyland.²⁰

However, our interest is in the “data poor” regime, which is common in problems involving experimental, rather than computational experiments. In the double gyre example, we demonstrated that the coherent sets identified using a limited number of thin plate splines agreed well with the sets obtained using a larger number of indicator functions. The benefit of using thin plate splines or other mesh-free basis functions is that they can also be used in problems where a computational grid is not easily defined. This is useful in the second example involving the Bickley Jet, where the initial domain is not a trapping region, and the domain at the final time resembles a “sawtooth.” By using radial basis functions, the same procedure used for the double gyre can also be used here without alteration. Our final example is identifying an eddy in the Sulu Sea, which possessed a changing computational domain in combination with a relatively small coherent set of physical interest. In that example, we also demonstrated that our approach can identify the eddy of interest even with relatively small amounts of data that approach the number of drifters used in recent massive drifter deployment experiments.^{49,50}

In all three examples, the noise added to the x_m and y_m was normally distributed with a standard deviation that was small compared to the spatial scales on which the problem was defined. However, it appears that the “noise” in the dynamics introduced implicitly by our basis functions typically has a far larger impact on the resulting coherent sets. If this is not the case, one improvement to the procedure would be to use observation-based spatially dependent anisotropic diffusivities (see, for example, Rypina *et al.*⁵¹) to represent the stochastic portion of the flow rather than arbitrarily choosing a distribution as we do here.

Because of the crucial role transport barriers play in understanding systems with chaotic mixing, algorithmic methods for identifying these barriers are useful tools for researchers in application areas such as geophysical fluid dynamics, combustion, and even those focused on ecological problems. In some situations, one either knows or can approximate the velocity field of the flow, which enables standard techniques and software packages such as FTLE fields or GAIO to be used. However, in other applications, the velocity field cannot be obtained analytically or numerically, and Lagrangian data from drifting buoys are all that is available to us. Ultimately, algorithms such as the one presented here are the first steps towards adapting the techniques we would use in a data rich environment for use in practical problems where the needed Lagrangian data are sparse and difficult/expensive to obtain.

ACKNOWLEDGMENTS

The authors acknowledge I. G. Kevrekidis for helpful discussions on set-oriented methods and suggestions for this article. M.O.W. gratefully acknowledges support from the NSF (DMS-1204783). I.R. was supported by ONR (MURI Award N000141110087) and by the NSF (Grant No. 85464100). C.W.R. was supported by AFOSR (Grant No. FA9550-14-1-0289).

APPENDIX: THE LEADING SINGULAR VECTORS OF THE APPROXIMATION

To prove that the first unit vectors are also singular vectors of \hat{A} , we must first prove that \hat{A} is block diagonal. To show this structure arises, it is convenient to define the data matrices

$$\Psi_X = \begin{bmatrix} \psi_X(x_1)^T \\ \psi_X(x_2)^T \\ \vdots \\ \psi_X(x_M)^T \end{bmatrix}, \quad \Psi_Y = \begin{bmatrix} \psi_Y(y_1)^T \\ \psi_Y(y_2)^T \\ \vdots \\ \psi_Y(y_M)^T \end{bmatrix}, \quad (\text{A1})$$

where $\Psi_X \in \mathbb{R}^{M \times K_X}$ and $\Psi_Y \in \mathbb{R}^{M \times K_Y}$. Note that $A = \frac{1}{M} \Psi_X^T \Psi_Y$, $G_X = \frac{1}{M} \Psi_X^T \Psi_X$, and $G_Y = \frac{1}{M} \Psi_Y^T \Psi_Y$.

Because the Cholesky decomposition is unique, we can also write

$$\Psi_X = \sqrt{M} Q_X L_X^T, \quad \text{and} \quad \Psi_Y = \sqrt{M} Q_Y L_Y^T, \quad (\text{A2})$$

where Q_X and Q_Y are orthonormal matrices. While Q_X and Q_Y , in general, differ, this is a QR decomposition, so the first columns of Q_X and Q_Y are normalized version of the first columns of Ψ_X and Ψ_Y . However, because of our choice of basis functions (in particular, $\psi_1 = \tilde{\psi}_1 = 1$), the first columns of Ψ_X and Ψ_Y are identical. Therefore, the first columns of Q_X and Q_Y , which we refer to as $q_X^{(1)}$ and $q_Y^{(1)}$, are also identical. Furthermore, $Q_Y^T q_X^{(1)} = \tilde{e}_1$ and $Q_X^T q_Y^{(1)} = e_1$ by the orthonormality of Q_X and Q_Y . Finally, $\hat{A} = L_X^{-1} A L_Y^{-T} = L_X^{-1} L_X (Q_X^T Q_Y) L_Y^T L_Y^{-T} = Q_X^T Q_Y$. Because the first columns of Q_X and Q_Y are identical

$$\hat{A} = Q_X^T Q_Y = \begin{bmatrix} 1 & 0 \\ 0 & \hat{A}_{22} \end{bmatrix}, \quad (\text{A3})$$

which shows the matrix \hat{A} possess the desired block structure. Because of this structure, it is now clear that e_1 and \tilde{e}_1 are singular vectors associated with the singular value $\sigma = 1$.

¹J. Ottino, "Mixing, chaotic advection, and turbulence," *Annu. Rev. Fluid Mech.* **22**, 207–254 (1990).

²H. Aref, "The development of chaotic advection," *Phys. Fluids* **14**, 1315–1325 (2002).

³S. Wiggins, "The dynamical systems approach to Lagrangian transport in oceanic flows," *Annu. Rev. Fluid Mech.* **37**, 295–328 (2005).

⁴I. Mezić, S. Loire, V. A. Fonoberov, and P. Hogan, "A new mixing diagnostic and Gulf oil spill movement," *Science* **330**, 486–489 (2010).

⁵M. J. Olascoaga and G. Haller, "Forecasting sudden changes in environmental pollution patterns," *Proc. Natl. Acad. Sci. U. S. A.* **109**, 4738–4743 (2012).

⁶T. Ma and E. M. Bollt, "Differential geometry perspective of shape coherence and curvature evolution by finite-time nonhyperbolic splitting," *SIAM J. Appl. Dyn. Syst.* **13**, 1106–1136 (2014).

⁷G. Haller, "Finding finite-time invariant manifolds in two-dimensional velocity fields," *Chaos* **10**, 99–108 (2000).

⁸S. L. Brunton and C. W. Rowley, "Fast computation of finite-time Lyapunov exponent fields for unsteady flows," *Chaos* **20**, 017503 (2010).

⁹S. C. Shadden, F. Lekien, and J. E. Marsden, "Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows," *Physica D* **212**, 271–304 (2005).

¹⁰I. Rypina, L. J. Pratt, J. Pullen, J. Levin, and A. L. Gordon, "Chaotic advection in an archipelago," *J. Phys. Oceanogr.* **40**, 1988–2006 (2010).

¹¹G. Haller and G. Yuan, "Lagrangian coherent structures and mixing in two-dimensional turbulence," *Physica D* **147**, 352–370 (2000).

¹²I. Rypina, S. Scott, L. J. Pratt, and M. G. Brown, "Investigating the connection between complexity of isolated trajectories and Lagrangian coherent structures," *Nonlinear Process. Geophys.* **18**, 977–987 (2011).

¹³G. Haller and T. Sapsis, "Lagrangian coherent structures and the smallest finite-time Lyapunov exponent," *Chaos* **21**, 023115 (2011).

¹⁴F. J. Beron-Vera, Y. Wang, M. J. Olascoaga, G. J. Goni, and G. Haller, "Objective detection of oceanic eddies and the Agulhas leakage," *J. Phys. Oceanogr.* **43**, 1426–1438 (2013).

¹⁵M. Budišić and I. Mezić, "Geometry of the ergodic quotient reveals coherent structures in flows," *Physica D* **241**, 1255–1269 (2012).

¹⁶C. Mendoza, A. Mancho, and S. Wiggins, "Lagrangian descriptors and the assessment of the predictive capacity of oceanic data sets," *Nonlinear Process. Geophys.* **21**, 677–689 (2014).

¹⁷G. Froyland, K. Padberg, M. H. England, and A. M. Treguier, "Detection of coherent oceanic structures via transfer operators," *Phys. Rev. Lett.* **98**, 224503 (2007).

¹⁸G. Froyland, S. Lloyd, and N. Santitissadeekorn, "Coherent sets for nonautonomous dynamical systems," *Physica D* **239**, 1527–1541 (2010).

¹⁹G. Froyland, N. Santitissadeekorn, and A. Monahan, "Transport in time-dependent dynamical systems: Finite-time coherent sets," *Chaos* **20**, 043116 (2010).

²⁰G. Froyland, "An analytic framework for identifying finite-time coherent sets in time-dependent dynamical systems," *Physica D* **250**, 1–19 (2013).

²¹E. M. Bollt and N. Santitissadeekorn, *Applied and Computational Measurable Dynamics* (SIAM, 2013), Vol. 18.

²²G. Froyland and K. Padberg, "Almost-invariant sets and invariant manifolds—Connecting probabilistic and geometric descriptions of coherent structures in flows," *Physica D* **238**, 1507–1523 (2009).

²³M. Dellnitz, G. Froyland, and O. Junge, "The algorithms behind GAIO-set oriented numerical methods for dynamical systems," in *Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems* (Springer, 2001), pp. 145–174.

²⁴B. K. Horn and B. G. Schunck, "Determining optical flow," in *1981 Technical Symposium East* (International Society for Optics and Photonics, 1981), pp. 319–331.

²⁵F. A. Mussa-Ivaldi, "From basis functions to basis fields: Vector field approximation from sparse data," *Biol. Cybern.* **67**, 479–489 (1992).

²⁶B. O. Koopman, "Hamiltonian systems and transformation in Hilbert space," *Proc. Natl. Acad. Sci. U. S. A.* **17**, 315 (1931).

²⁷B. Koopman and J. v. Neumann, "Dynamical systems of continuous spectra," *Proc. Natl. Acad. Sci. U. S. A.* **18**, 255 (1932).

²⁸I. Mezić, "Spectral properties of dynamical systems, model reduction and decompositions," *Nonlinear Dyn.* **41**, 309–325 (2005).

²⁹M. Budišić, R. Mohr, and I. Mezić, "Applied Koopmanism," *Chaos* **22**, 047510 (2012).

³⁰G. Froyland, O. Junge, and P. Koltai, "Estimating long-term behavior of flows without trajectory integration: The infinitesimal generator approach," *SIAM J. Numer. Anal.* **51**, 223–247 (2013).

³¹P. Koltai, *Efficient Approximation Methods for the Global Long-Term Behavior of Dynamical Systems: Theory, Algorithms and Examples* (Logos Verlag Berlin GmbH, 2011).

³²M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, "A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition," *J. Nonlinear Sci.* preprint arXiv:14084408 (in press).

³³L. N. Trefethen, *Spectral Methods in MATLAB* (SIAM, 2000), Vol. 10.

³⁴J. P. Boyd, *Chebyshev and Fourier Spectral Methods* (Courier Dover Publications, 2013).

³⁵R. J. LeVeque, *Finite Volume Methods for Hyperbolic Problems* (Cambridge University Press, 2002), Vol. 31.

³⁶H. Wendland, "Meshless Galerkin methods using radial basis functions," *Math. Comput. Am. Math. Soc.* **68**, 1521–1531 (1999).

³⁷T. Ma and E. M. Bollt, "Relatively coherent sets as a hierarchical partition method," *Int. J. Bifurcation Chaos* **23**, 1330026 (2013).

³⁸I. Mezić, "Analysis of fluid flows via spectral properties of the Koopman operator," *Annu. Rev. Fluid Mech.* **45**, 357–378 (2013).

³⁹C. J. Albers, F. Critchley, and J. C. Gower, "Quadratic minimisation problems in statistics," *J. Multivar. Anal.* **102**, 698–713 (2011).

- ⁴⁰C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, "Spectral analysis of nonlinear flows," *J. Fluid Mech.* **641**, 115–127 (2009).
- ⁴¹A. Iske, *Multiresolution Methods in Scattered Data Modelling* (Springer Science & Business Media, 2004), Vol. 37.
- ⁴²G.-R. Liu, *Meshfree Methods: Moving Beyond the Finite Element Method* (CRC Press, 2010).
- ⁴³G. E. Fasshauer, "Solving partial differential equations by collocation with radial basis functions," in *Proceedings of Chamonix* (Citeseer, 1996), Vol. 1997, pp. 1–8.
- ⁴⁴C. M. Bishop *et al.*, *Pattern Recognition and Machine Learning* (Springer, New York, 2006), Vol. 1.
- ⁴⁵I. I. Rypina, M. G. Brown, F. J. Beron-Vera, H. Kocak, M. J. Olascoaga, and I. A. Udovydchenkov, "On the Lagrangian dynamics of atmospheric zonal jets and the permeability of the stratospheric polar vortex," *J. Atmos. Sci.* **64**, 3595–3610 (2007).
- ⁴⁶A. F. Shchepetkin and J. C. McWilliams, "The regional oceanic modeling system (ROMS): A split-explicit, free-surface, topography-following-coordinate oceanic model," *Ocean Modell.* **9**, 347–404 (2005).
- ⁴⁷P. Deuffhard, W. Huisinga, A. Fischer, and C. Schütte, "Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains," *Linear Algebra Appl.* **315**, 39–59 (2000).
- ⁴⁸P. Deuffhard and M. Weber, "Robust Perron cluster analysis in conformation dynamics," *Linear Algebra Appl.* **398**, 161–184 (2005).
- ⁴⁹I. I. Rypina, A. R. Kirincich, R. Limeburner, and I. A. Udovydchenkov, "Eulerian and Lagrangian correspondence of high-frequency radar and surface drifter data: Effects of radar resolution and flow components," *J. Atmos. Oceanic Technol.* **31**, 945–966 (2014).
- ⁵⁰A. C. Poje, T. M. Özgökmen, B. L. Lipphardt, B. K. Haus, E. H. Ryan, A. C. Haza, G. A. Jacobs, A. J. H. M. Reniers, M. J. Olascoaga, G. Novelli, A. Griffa, F. J. Beron-Vera, S. S. Chen, E. Coelho, P. J. Hogan, A. D. Kirwan, H. S. Huntley, and A. J. Mariano, "Submesoscale dispersion in the vicinity of the deepwater horizon spill," *Proc. Natl. Acad. Sci. U. S. A.* **111**, 12693–12698 (2014).
- ⁵¹I. I. Rypina, I. Kamenkovich, P. Berloff, and L. J. Pratt, "Eddy-induced particle dispersion in the near-surface North Atlantic," *J. Phys. Oceanogr.* **42**, 2206–2228 (2012).