

RESEARCH ARTICLE

Simultaneous coherent structure coloring facilitates interpretable clustering of scientific data by amplifying dissimilarity

Brooke E. Husic^{1*}, Kristy L. Schlueter-Kuck², John O. Dabiri^{2,3*}

1 Department of Chemistry, Stanford University, Stanford, California, United States of America, **2** Department of Mechanical Engineering, Stanford University, Stanford, California, United States of America, **3** Department of Civil and Environmental Engineering, Stanford University, Stanford, California, United States of America

* bhusic@stanford.edu (BEH); jodabiri@stanford.edu (JOD)



OPEN ACCESS

Citation: Husic BE, Schlueter-Kuck KL, Dabiri JO (2019) Simultaneous coherent structure coloring facilitates interpretable clustering of scientific data by amplifying dissimilarity. PLoS ONE 14(3): e0212442. <https://doi.org/10.1371/journal.pone.0212442>

Editor: David L. Hu, Georgia Institute of Technology, UNITED STATES

Received: September 3, 2018

Accepted: February 1, 2019

Published: March 13, 2019

Copyright: © 2019 Husic et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The three fluid mechanics datasets and all adjacency matrices used to create the models in this work are available on github at <https://github.com/brookehus/sCSC>. This repository also contains example MATLAB and Python codes, including Jupyter notebook tutorials. The all-atom molecular dynamics simulations of Protein G were previously published in Ref. [51], and the trajectories are available at no cost for non-commercial use through contacting trajectories@deshawresearch.com.

Abstract

The clustering of data into physically meaningful subsets often requires assumptions regarding the number, size, or shape of the subgroups. Here, we present a new method, simultaneous coherent structure coloring (sCSC), which accomplishes the task of unsupervised clustering without *a priori* guidance regarding the underlying structure of the data. sCSC performs a sequence of binary splittings on the dataset such that the most dissimilar data points are required to be in separate clusters. To achieve this, we obtain a set of orthogonal coordinates along which dissimilarity in the dataset is maximized from a generalized eigenvalue problem based on the pairwise dissimilarity between the data points to be clustered. This sequence of bifurcations produces a binary tree representation of the system, from which the number of clusters in the data and their interrelationships naturally emerge. To illustrate the effectiveness of the method in the absence of *a priori* assumptions, we apply it to three exemplary problems in fluid dynamics. Then, we illustrate its capacity for interpretability using a high-dimensional protein folding simulation dataset. While we restrict our examples to dynamical physical systems in this work, we anticipate straightforward translation to other fields where existing analysis tools require *ad hoc* assumptions on the data structure, lack the interpretability of the present method, or in which the underlying processes are less accessible, such as genomics and neuroscience.

Introduction

Modern science increasingly leverages machine learning on large datasets in the sciences, from electronic structure [1] to whole genome sequences [2] to distributed ocean sensor measurements [3]. Many of these datasets capture the dynamics of a system evolving in time, encoding trends with predictive power. Analyzing these datasets using a statistically robust and interpretable framework is a longstanding challenge that often involves clustering, or the unsupervised learning of coherent groups within the dataset.

Funding: BEH and JOD received no specific funding for this work. KLS-K was supported by the U.S. National Science Foundation and by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

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

Clustering is a notoriously challenging problem which, unlike supervised learning, features no direct measure of model success or validity and often requires heuristic assessments of effectiveness [4]. Thus, many classes of clustering algorithms have been developed for different problems. Some commonly used techniques include partition-based methods such as *k*-means [5], or their fuzzy counterparts [6]; density-based methods such as DBSCAN [7]; and connectivity-based methods such as divisive and agglomerative hierarchical clustering [8, 9].

Each of the aforementioned methods exhibits drawbacks with respect to *a priori* assumptions and algorithmic limitations. For example, partition-based clustering such as *k*-means requires the modeler to prescribe the number of partitions in a dataset before constructing the model. If multiple results are obtained from different values of *k*, these results are not interrelated; similarly, the model cannot be used to determine relationships between the *k* clusters of a single model. While connectivity-based methods feature interrelated clusters, these also require the determination of where to cut the corresponding dendrogram to obtain the clustering result. Although density-based methods do not require *a priori* or *a posteriori* determination of the number of clusters to use, these methods are generally not robust to datasets containing a range of cluster densities [10].

Here, we present a new method, simultaneous coherent structure coloring (sCSC), which minimizes the assumptions required in an unsupervised clustering task. sCSC focuses solely on the efficient separation of the most dissimilar states in the system, resulting in a quantitative structure that automatically captures the clusters in the dataset and their interrelationships without *a priori* knowledge of the system. The method is demonstrated for simulated and empirical systems of fluid and molecular dynamics, and its straightforward extension to other types of data is discussed.

Background

The use of clustering for data analysis is ubiquitous. However, our motivation emerged from research on the identification of coherent structures from fluid dynamics. A variety of mathematical frameworks have been developed to identify coherent structures. The broad class of Lagrangian methods has been developed to describe flows that are unsteady (i.e., not well-summarized by instantaneous snapshots) in a way that is not dependent on their frame of reference (i.e., may not contain velocity or acceleration data). Two recent reviews of Lagrangian methods for the detection of coherent structures in fluids can be found in Refs. [11] and [12].

Existing algorithms for coherent structure analysis that involve clustering exhibit various limitations. The fuzzy *c*-means approach presented in Ref. [3], for example, introduces a dynamic distance between particle trajectories, but ultimately requires the choice of *c*, i.e. how many clusters to use. To avoid explicitly choosing the number of coherent structures, a spectral clustering method was introduced in Ref. [13], which utilizes the spectral gap in the graph Laplacian to determine the number of coherent structures. However, it was subsequently shown in Ref. [14] that such a gap is only robust when the number of trajectories used exceeds 10^3 .

The method of coherent structure coloring (CSC), introduced in Refs. [14] and [15], was designed to address these and other limitations of clustering algorithms for coherent structure determination based on trajectories of particle spatial coordinates. In this work, we extend CSC in the context of its own limitations, as described below.

While we restrict the focus of the rest of the paper to clustering, this is not the only way to identify coherent structures from frame-independent particle trajectories. Over the past two decades, both the fluid and molecular dynamics communities have developed methods to identify “almost-invariant” sets through data-driven approximations to the Perron-Frobenius

operator and its adjoint, the Koopman operator. The former, also referred to as the transfer or transition operator, propagates probability densities forward in time, whereas the latter propagates observables [16].

For fluid systems, Dellnitz and Junge [17], Froyland and Dellnitz [18], and Mezić [19] used finite approximations to the Perron-Frobenius eigenfunctions to divide the space occupied by a dynamical system into almost-invariant sets and almost-cycles. At the same time, Schütte *et al.* [20] and Deuffhard *et al.* [21] introduced the use of approximations to the same operator, under a reversibility constraint, to determine the “metastable” states of molecular systems simulated on the atomic level.

A decade later, researchers in their respective fields independently determined equivalent algorithms for optimizing the estimation of the approximated eigenfunctions of the Perron-Frobenius and Koopman operators [22–24]. In both cases, linear models are generated using a data-driven, objective protocol to model highly nonlinear dynamics, where the eigenvectors and eigenvalues can be used to identify coherent sets.

In fact, in our final example in the current study, we utilize both types of methods on a molecular dynamics simulation dataset. We first create an optimized model that approximates the Perron-Frobenius operator [22], which is difficult to visualize due to its high dimensionality. Then, to reduce our model to a visualizable and interpretable coarse-grained model, we use the clustering method presented in this work to identify the major coherent structures.

Simultaneous coherent structure coloring (sCSC)

Coherent structure coloring theory

Many datasets we wish to explore in the physical sciences are generated by complex dynamical systems that exhibit instabilities and chaos. A key consequence of these processes is that states of the system (e.g. fluid particle trajectories or protein conformations) that are proximal but belonging to different coherent sets will separate exponentially faster as the system evolves than states belonging to the same cluster [25, 26].

On this basis, we previously hypothesized that these complex datasets can be clustered more robustly and effectively by amplifying state differences rather than state similarity [14, 15]. The rationale for this approach is that the exponential separation of dissimilar states can provide more sensitive detection of clusters than a focus on state similarity, the latter requiring longer observation to become apparent [25, 26]. In other words, we aim to identify coherent clusters indirectly, by prioritizing the separation of states with greatest dissimilarity and confidently ruling out the possibility of their membership in the same cluster. Those states that remain together after the separation process will subsequently emerge as belonging to the same cluster.

To amplify the dissimilarity between states, we solve an optimization problem to maximize a figure of merit z that quantifies total state dissimilarity in the dataset. Specifically, this figure of merit depends on a scalar value x_i assigned to each state i in the system, where the squared difference in the scalar value assigned to each of pair states (e.g. $(x_1 - x_2)^2$ for states 1 and 2) is weighted by a measure of their dissimilarity. Formally, the clustering parameter z is given by

$$z \equiv \frac{1}{2} \sum_i^n \sum_j^n (x_i - x_j)^2 a_{ij}, \quad (1)$$

where the summations of i and j are each taken over the full set of n states to be clustered, and a_{ij} is an element of the adjacency matrix A containing the pairwise dissimilarity between states i and j . The construction of this matrix requires the calculation of $\binom{n}{2} = (n - 1)n/2$ adjacency

values. Example definitions of the (symmetric) pairwise dissimilarity can include the standard deviation for comparison of time-dependent signals, or the Jensen-Shannon divergence for comparison of probability distributions [27, 28]. Both definitions represent measures of dissimilarity, where identical data points receive $a_{ij} = 0$. Thus, assuming all data points are unique, the matrix A will be dense.

Given the adjacency matrix A , we seek to find state assignments x_i that will maximize z , subject to the constraint that the magnitude of the $n \times 1$ vector X containing the n scalar values x_i must remain finite (e.g. to avoid the trivial case that maximizes z for $x_1 = \infty$ and $x_2 = -\infty$). It is straightforward to show that the constrained optimization of Eq 1 with finite X can be written as the generalized eigenvalue problem [29]:

$$LX = \lambda DX, \quad (2)$$

where D is a diagonal matrix with entries equal to the row-sums of the adjacency matrix, i.e. $\sum_j a_{ij}$ for each row i , and $L = D - A$ is the graph Laplacian. This maximization is expressed using the Lagrangian form; see [14] for more details.

Each of the n eigenvectors X_n of Eq 2 represents a solution that assigns to each state a scalar value x_i based on its dissimilarity to the other states in the system. Those states with scalar assignments in each X that are most dissimilar can be presumed to belong to different clusters of the data when the data is partitioned according to that particular solution of Eq 1. The eigenvector X_1 associated with the maximum eigenvalue λ_1 contains the scalar assignments x_i that maximize the figure of merit z . This can be considered the single most effective partitioning of the dataset.

Given the analogy between this approach and the problem of fuzzy graph coloring [30], wherein the connected nodes of a graph with large weights are assigned the most disparate values, we call this method coherent structure coloring (CSC) [14]. The technique has recently been demonstrated to successfully identify coherent eddies and jets associated with individual fluid particle trajectories in model geophysical flows [15].

Simultaneous inclusion of multiple CSC solutions

A key limitation of the original CSC method [14] is that it relies on only a single eigenvector associated with the largest eigenvalue of Eq 2. Hence, although multiple dimensions of dissimilarity are almost always present in real data, the method cannot simultaneously distinguish between multiple types of dissimilarity in a dataset. Moreover, the method applied to individual fluid particle trajectories in a subsequent study required a subjectively defined threshold to calculate eigenspace distances [15], and it was shown to produce degenerate results for fluid particles in chaotic regions of the flow (cf. Fig 7 in Ref. [15]).

Importantly, because the adjacency matrix A introduced in the previous section is real and symmetric, the remaining eigenvectors associated with lesser eigenvalues provide additional, linearly independent (i.e. orthogonal) solutions for partitioning the data, albeit less effectively [31]. The key innovation of the present work is to use *all* of the eigenvectors in a top-down fashion to simultaneously cluster the system states.

To perform sCSC, we begin with the most effective partition given by the eigenvector associated with the maximum eigenvalue, and proceed through the set of orthogonal eigenvectors in order of decreasing eigenvalue. This approach simultaneously reveals the coherent sets of the system, and eliminates the subjective user intervention required in the previous method [15].

Given a dissimilarity measure and resulting eigenvector solutions, the simultaneous coherent structure coloring (sCSC) algorithm begins by assigning to each state in the system a

binary membership based on its corresponding scalar value along each orthogonal coordinate direction. A bifurcation is appropriate given that each one-dimensional coordinate has two extreme ends toward which the optimization of Eq 1 pulls dissimilar states.

The states are bifurcated along each coordinate dimension by using agglomerative clustering with average linkage (although other linkages or splitting methods could be used for this step; see e.g. Ref. [32], Table 1.) and assigning to each state a value of 0 or 1 based on its membership within either of the two largest clusters of the resulting dendrogram. Each eigenvector contributes a separate bit to the binary code associated with each of the states in the system, with the leading bit corresponding to the largest eigenvalue and the remaining bits concatenated in decreasing order of their corresponding eigenvalues. Though we suggest using a bifurcation in general, the method does not prohibit the division of each eigenvector coordinate into three or more discrete bins, thus creating a *k*-way splitting and associated base-*k* codes.

For each subsequent eigenvector, the bifurcation is performed for all data points (i.e. states), and each is assigned a 0 or a 1. For the *k*th eigenvector bifurcation, this enables 2^k numerically possible clusters (Fig 1). For example, the first splitting produces branches 0 and 1, and the second splitting enables the population of 2^2 unique clusters by appending 0 or 1 to each branch of the existing binary code ({00, 01, 10, 11}). However, it may be the case that the numerically possible branch 01 is not occupied because there is no data point that receives both a label of 0 in the first bifurcation and a label of 1 in the second bifurcation. Thus, we hypothesize that branch 0 (and its only occupied split, branch 00) evidences a coherent region of the data. In this way, a natural stopping criterion emerges from unoccupied bit codes during the binary splitting.

The binary codes generated by the aforementioned process can be visualized in a dendrogram, with each branch pair connecting those states that differ only at the least significant bit of their binary code. The length of each branch pair is a measure of the dissimilarity between the groups connected by the branches, and it corresponds to the value of the summation in Eq (1) computed only over those states connected by the branches. Bits for progressively smaller eigenvalues are included at progressively lower levels of the dendrogram. The dissimilarity between the groups connected at lower levels therefore generally becomes smaller as well.

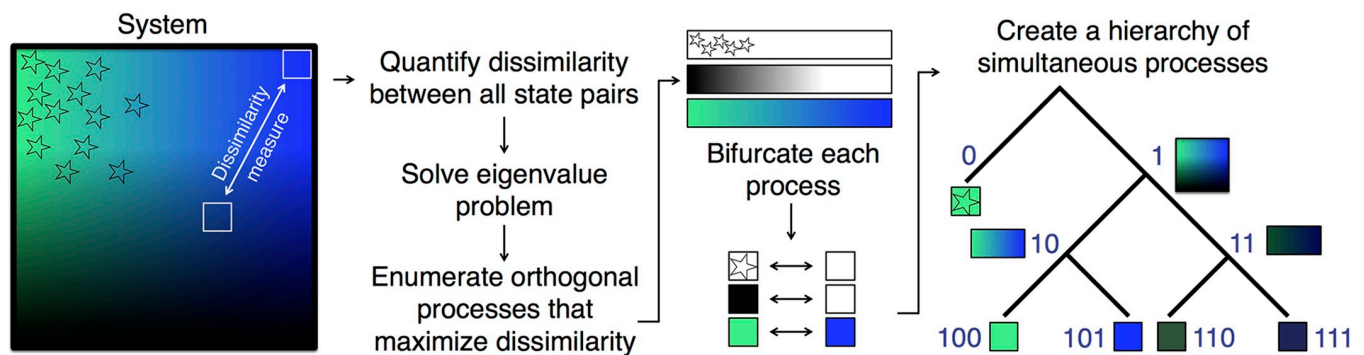


Fig 1. Conceptual scheme illustrating the sCSC algorithm. First, the dissimilarity between all pairs of states are tabulated in an adjacency matrix. For this example system, states are represented by a uniform grid of squares, two of which are illustrated in the left panel. The adjacency matrix is then used to solve an eigenvalue problem (Eq (2)) that maximizes the dissimilarity measure. The solutions to the eigenvalue problem identify orthogonal processes in order of their ability to separate the system; in this case, we have stars ↔ no stars, bright ↔ dark, and green ↔ blue, which we have asserted are decreasingly effective in explaining dissimilarity in this notional system. These three processes are bifurcated into two extremes (middle panel). Then, each state is encoded according to each bifurcation. For the first orthogonal process (stars ↔ no stars), we bifurcate the entire system. For the next orthogonal process (bright ↔ dark), we bifurcate the system separately and illustrate only states which become bifurcated along this division. For example, there is no state that contains stars and is dark, so branch 0 of the corresponding dendrogram is not further bifurcated. Finally, we bifurcate both branches 10 and 11 according to green or blue.

<https://doi.org/10.1371/journal.pone.0212442.g001>

While in principle the sCSC dendrogram should naturally truncate when no further splits occur, large amounts of data points or statistical noise may lead to insignificant (i.e., low- z) clusters or explore a combinatorially unfavorable number of splits. In that case, one may choose to truncate the dendrogram after a certain number of eigenvectors according to visual inspection, or determine a cutoff based on the magnitude of z or the eigenvalue.

As in standard divisive and hierarchical clustering methods, the clustering models produced with sCSC are dependent on the adjacency definition supplied by the user. Because the adjacency matrix summarizes pairwise dissimilarities only, this has the benefit of not requiring adherence to the triangle inequality [33]—in fact, the data points need not exist in a well-defined space at all. However, with this flexibility comes the drawback that a poor dissimilarity metric may obscure patterns in the data. The dissimilarity measures used in this study have been shown to be effective in previous studies [14, 28], and in general may require domain-specific knowledge to determine for a given dataset.

In the next two sections, we apply sCSC to benchmark problems in fluid dynamics in order to demonstrate its effectiveness in identifying coherent structures in the absence of *a priori* assumptions. Then, we demonstrate the use of sCSC to determine the number and shape of flow structures involved in vortex ring entrainment using data obtained from empirical measurements the laboratory. Finally, to highlight the interpretability of the sCSC dendrogram for high-dimensional datasets, we use sCSC to visualize an interpretable representation of an atomistic protein folding simulation. Finally, we discuss the relationship of this method to other unsupervised clustering methods, and the possibility of extending sCSC beyond physical dynamical systems.

Coherent structure identification from analytical geophysical flow simulations and empirical measurements

Quadruple-eddy simulated ocean flow

A key challenge in geophysical fluid dynamics is to extract and characterize coherent fluid motions from sparsely sampled turbulent flows of air or water. The coherent structures, often manifested as eddies and jets, can dominate the transport of heat, salt, nutrients, and pollutants [34, 35]. Therefore, they can serve as the basis for low-order models that capture the salient physics [36], or as a template for data assimilation into large-scale weather forecasting models [37]. Turbulent flow structures in the ocean also impact the behavior and ecology of marine life [38].

Distributed sensor networks such as the Argo collection of 3800 ocean drifters [39] sample the flow field in a Lagrangian sense, recording the properties of the water as each drifter is carried by the prevailing currents. Here we demonstrate the capability of the sCSC method to extract coherent fluid structures from such collections of Lagrangian measurements.

To do so, we first apply sCSC to a common model of Lagrangian ocean drifters in a simplified flow field comprising only four eddies, the unsteady quadruple-eddy flow [3, 11]. While this model represents a simplification of the full physics, it is valuable due to its common use for the evaluation and comparison with existing methods to identify coherent structures [3, 11, 14, 15].

As shown in Fig 2A, drifter trajectories within the two eddies at the upper-left and lower-right rotate clockwise, whereas trajectories within the other two eddies rotate counter-clockwise. Simultaneous with this rotation, an east-west oscillation of the eddy field occurs, which causes exchange of drifters between the eastern and western eddies. This exchange, which depends on the location and timing of the drifter release relative to the east-west oscillation

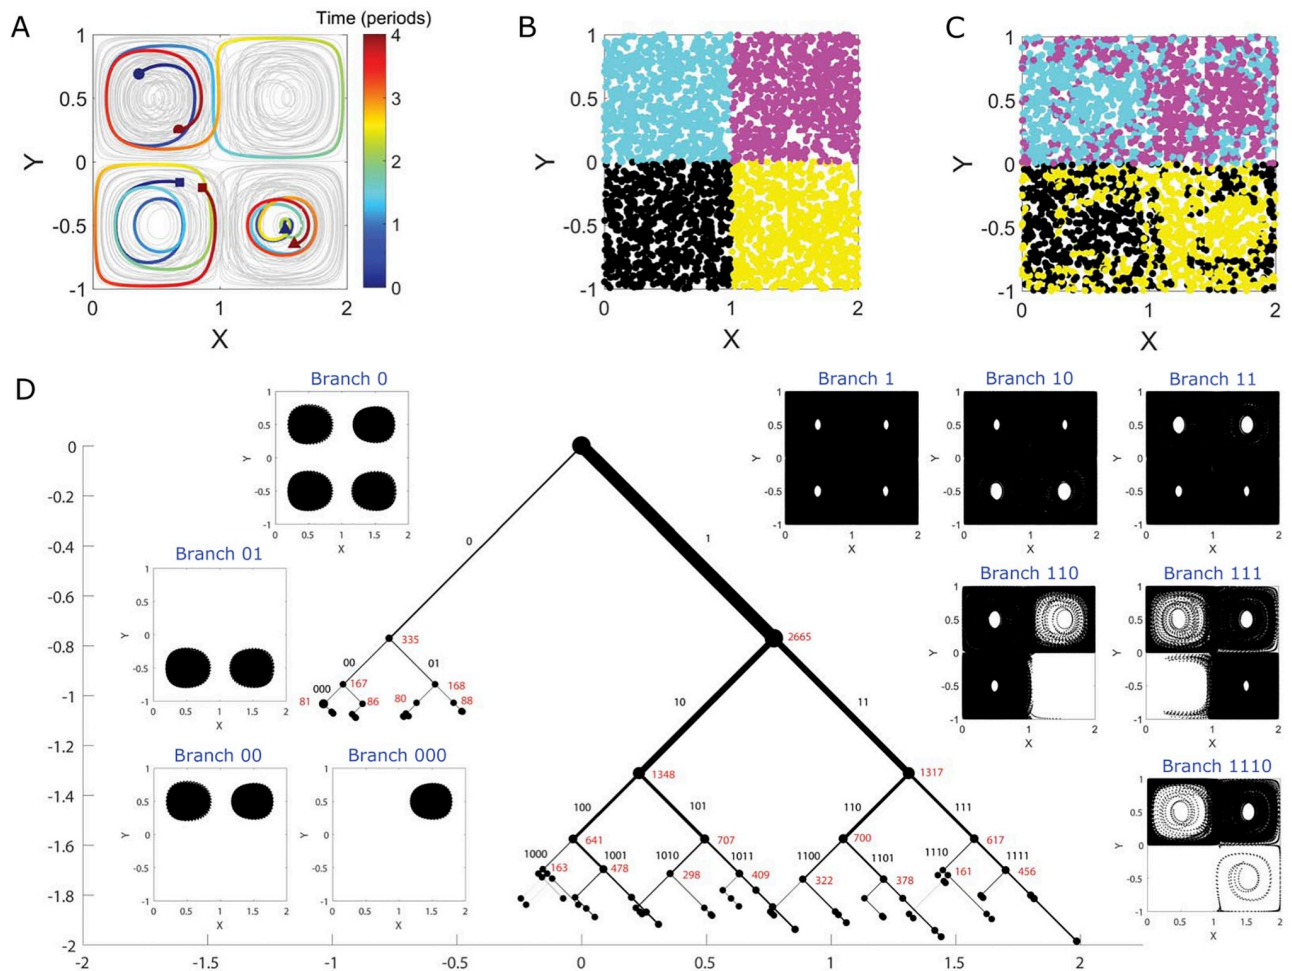


Fig 2. Quadruple-eddy ocean flow model. (A) Trajectories of 50 selected drifters randomly initialized in the flow (gray). The trajectories of 3 drifters are highlighted for 4 periods of horizontal oscillation, from their initial positions (blue) to their final positions (red). These drifters illustrate qualitatively different trajectories in the flow, including those that switch quadrants (dots), those that remain in a single eddy core (triangles), and those that spiral radially between the center and the boundary of a quadrant (squares). In panel (B), the initial positions of 3000 randomly initialized drifters are colored according to their initial quadrant in the flow. The drifters maintain their color assignment in panel (C), showing how the unsteady eddy motion leads to mixing of the drifters after the 4 periods of horizontal oscillation. The east-west oscillation of the eddy field leads horizontal mixing of the flow. The resulting sCSC dendrogram is shown in panel (D), with every position occupied by all 3000 drifters plotted in black dots in the corresponding inset branch plot (note that drifter positions often appear as continuous black patches due to the high density of overlapping positions occupied by the drifters.) The width of each branch is proportional to the fraction of the states that it contains. The corresponding binary code of each branch is labeled in black text, and the number of trajectories associated with each node is labeled in red text. The dendrogram is plotted to the seventh eigenvector, although labels below the fourth eigenvector are omitted for clarity. The horizontal and vertical axes are measured in units of the parameter z , and the branches are plotted at 45-degree angles. We have visualized the first 7 eigenvectors for brevity of presentation.

<https://doi.org/10.1371/journal.pone.0212442.g002>

cycle, is illustrated in the transition from initial drifter positions in Fig 2B to their final positions in Fig 2C.

Each drifter trajectory represents a state of this fluid dynamic system, and the pairwise dissimilarity between each of the states is given by the standard deviation of the instantaneous distance between drifter positions at time t_k , $r_{ij}(t_k)$, divided by the average distance between each pair of drifters, \bar{r}_{ij} , for T total time points [14]:

$$a_{ij} = \frac{1}{\bar{r}_{ij}} \left[\sum_{k=1}^T (\bar{r}_{ij} - r_{ij}(t_k))^2 \right]^{\frac{1}{2}}. \tag{3}$$

This measure anticipates that coherent structures will comprise drifters whose relative positions do not vary as the flow evolves, leading to a small values of the pairwise dissimilarity measure (i.e. a small standard deviation) within each cluster. By contrast, pairs of drifters that straddle the boundary between coherent structures can experience exponential separation over time and a correspondingly large standard deviation of their instantaneous separation [26].

Without requiring the specification of the number of eddies, the sCSC method reveals a clear, physically interpretable structure for this complex flow (Fig 2D). The primary bifurcation of the flow is between trajectories that remain in the eddy cores of their original quadrant (branch 0) and trajectories that do not (branch 1). The trajectories of branch 0 are then further subdivided into trajectories that remain within eddy cores in the northern half of the flow (branch 00) and those that remain within eddy cores in the southern half of the flow (branch 01), reflecting the absence of north-south drifter exchange. Finally, the trajectories associated with the individual quadrants are identified at the level of the third bifurcation (e.g. branch 000 shown in Fig 2D inset, as well as branches 001, 010, and 011 for the other three individual quadrants, not shown in inset). An additional visualization of the major coherent structures identified—namely, branches 00, 01, 10, and 11 in Fig 2D—is presented in Fig 3A.

Whereas the application of *k*-means clustering or other conventional tools would require *a priori* guidance to determine that four independent structures exist in branch 0 (i.e. one eddy per quadrant) [12], this result is revealed naturally by the sCSC dendrogram, as

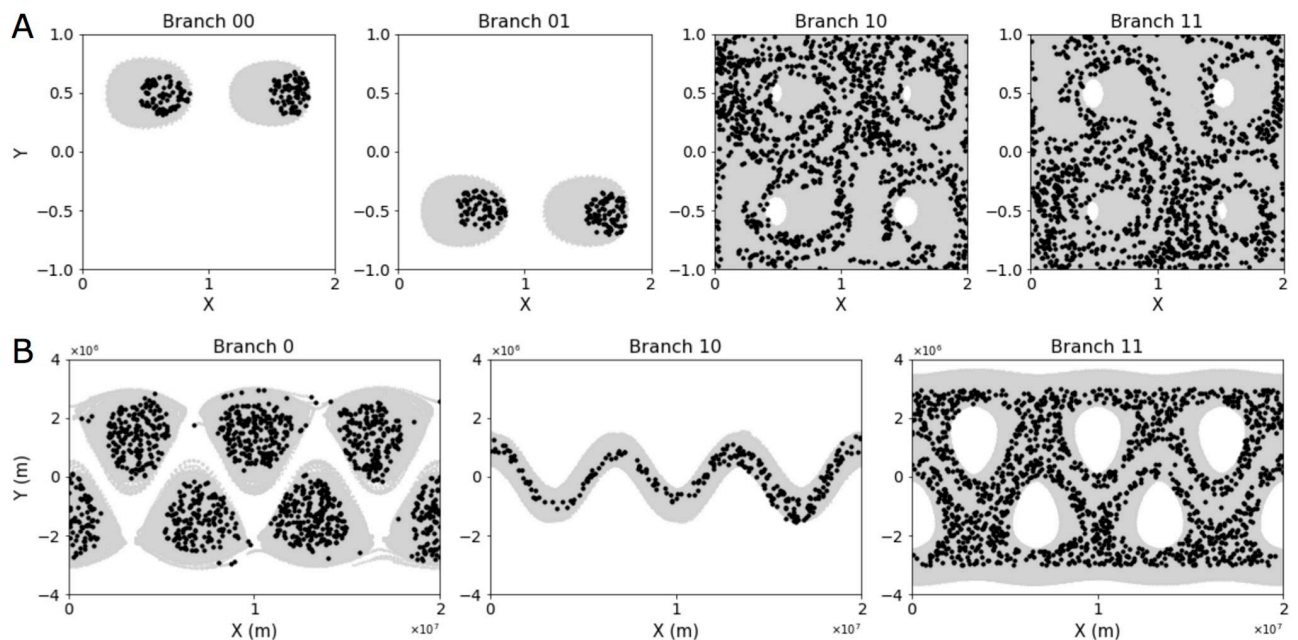


Fig 3. Selected coherent structures from the quadruple-eddy and Bickley jet. (A) Four selected branches from the quadruple-eddy analysis are visualized by plotting the particle trajectories belonging to a designated coherent structure (gray) and their starting positions (black). Branches 00 and 01 show the top and bottom eddy cores, respectively, whereas branches 10 and 11 show the incoherent region. The latter two branches also display an imbalance between the top and bottom regions of the flow. (B) Three selected branches from the Bickley jet analysis are visualized by plotting the particle trajectories belonging to a designated coherent structure (gray) and their starting positions (black). Branch 0 shows the eddy cores, which do not mix. Branch 10 contains the meandering jet, and branch 11 accounts for the incoherent surroundings. For both sets of plots, the starting positions show that the particles belonging to coherent structures at a given time point (in this case, the starting point) are more compactly located than the total space explored by their trajectories over time. A particle found at a given instant in the space that is common among different coherent structures can therefore not be attributed to a coherent structure based on that time point alone. Visually equivalent results can be produced from other time points.

<https://doi.org/10.1371/journal.pone.0212442.g003>

further bifurcations after branch 000 do not produce additional coherent states; all of the trajectories that remain together after the third bifurcation remain together after subsequent bifurcation.

To be sure, the presence of the four eddy cores can also be revealed by a contour map of the largest finite-time Lyapunov exponent (FTLE) corresponding to the quadruple-eddy velocity field (see Figs 2 and 4 in Ref. [14]). The key advantage of the sCSC approach is that a similar result can be achieved with two orders-of-magnitude less data: Schlueter-Kuck and Dabiri showed in Ref. [14] that the FTLE gradient calculation is well-posed when the number of drifters is on the order of 10^5 . By contrast, the same cores can be identified by as few as 300 drifters using the present method, and the cores can be identified as long as drifters are present in the cores over timescales longer than the eddy turnover time.

The structure of branch 1 is less well organized and reflects the chaotic advection of trajectories that spiral radially within a quadrant and/or switch quadrants in the unsteady flow. Nonetheless, the dendrogram structure does indicate geometric symmetries within the chaotic motions, such as a preference for three quadrants among the trajectories in branches 110 and 111; and a more constrained preference for two quadrants exists at branch 1110. A general observation is that geometric symmetries appear as balanced dendrogram bifurcations. This is in contrast to the structure of random noise, which is characterized by a trivial sCSC dendrogram with a single branch that contains nearly all of the states and a splintering of a small number of fully-converged states at each level of the dendrogram (see S1 Fig).

Bickley jet simulated atmospheric flow

A more complex geophysical flow model is the Bickley jet, which serves as a common model for zonal jets in the atmosphere [40]. This flow is composed of a central meandering jet as well as flanking eddies that are periodic along the east-west axis (Fig 4A). The sCSC dendrogram corresponding to this flow (for the same dissimilarity measure as the quadruple-eddy flow, Eq (3)) is similarly effective in extracting the salient coherent features (Fig 4D). The flanking eddies are identified in branch 0.

However, a key difference from the previous quadruple-eddy example is that the individual eddies are largely indistinguishable from one another. This result reflects the homogeneity of fluid dynamics within the flanking eddies, which was not present among the trajectories in the quadruple-eddy flow (contrast e.g. Figs 2C and 4C). A notable exception is the eddy located at the meridional axis of symmetry, i.e. branch 011. An additional analysis tracking the distance of particles from the eddy cores showed that fewer particles belonging to this center eddy travel past a given contour threshold during the simulated time-series than particles from other eddies. Branch 1 of the Bickley jet dendrogram collects those trajectories that are not associated with the flanking eddies. A subset of those trajectories, namely branch 10, is the meandering zonal jet. The remaining trajectories (branch 11) form a chaotic background flow that is robust to further bifurcation. These three coherent structures are further visualized in Fig 3B.

The sCSC structure of both of these simulated geophysical flows can be exploited to create low-order models of the governing fluid transport processes, without the need for *ad hoc* assumptions regarding the number of coherent structures present. Because similar results can be achieved despite significant missing or noisy data (see Ref. [14]), the inherently limited data collection that can be achieved in the ocean and atmosphere can be more effectively leveraged to potentially improve the accuracy of weather forecasting, for example [37]. Hence, the sCSC method can be a powerful tool for both very large and very sparse datasets.

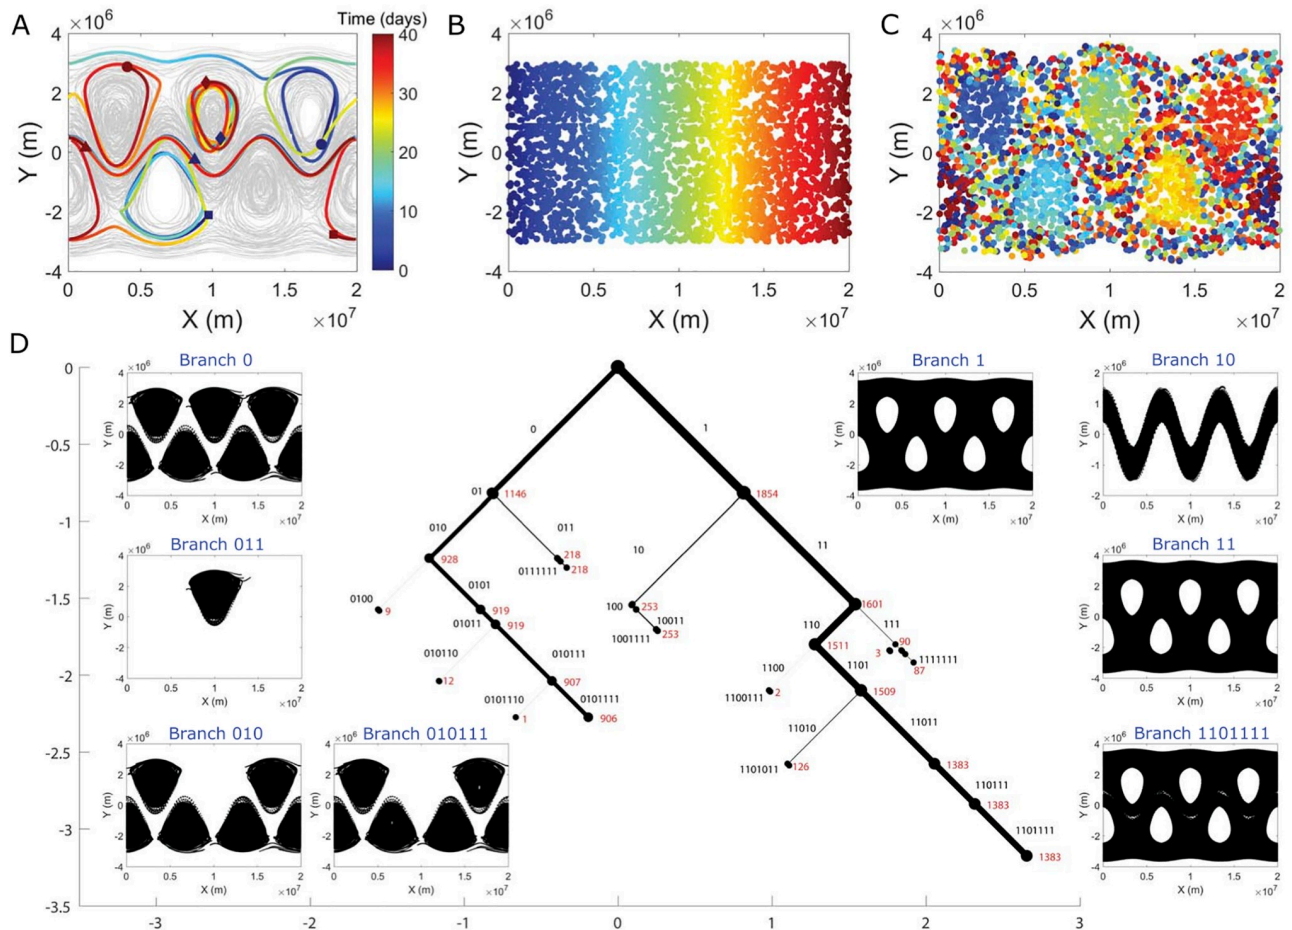


Fig 4. Bickley jet atmospheric flow model. (A) Trajectories of 75 selected Lagrangian particles randomly initialized in the flow (gray). The trajectories of 4 particles are highlighted for a 40-day integration period, from their initial positions (blue) to their final positions (red). These particles illustrate qualitatively different trajectories in the flow, including those that remain in a single flanking eddy (diamonds), those that pass between multiple eddies (dots and squares), and those in the meandering jet (triangles). In panel (B), the initial positions of 3000 particles are colored according to their position along the east-west axis of the flow. The particles maintain their color assignment in panel (C), showing how the unsteady jet and eddy motions lead to mixing of the particles after 40 days. A periodic boundary condition is applied in the east-west direction. The resulting sCSC dendrogram is shown in panel (D), with every position occupied by all 3000 particles plotted in black dots in the corresponding inset branch plot (note that particle positions often appear as continuous black patches due to the high density of overlapping positions occupied by the particles). The width of each branch is proportional to the fraction of the states that it contains. The corresponding binary code of each branch is labeled in black text, and the number of trajectories associated with each node is labeled in red text. The dendrogram is plotted to the seventh eigenvector, although many of the labels are omitted for clarity. The horizontal and vertical axes are measured in units of the parameter z , and the branches are plotted at 45-degree angles. We have visualized the first 7 eigenvectors for brevity of presentation.

<https://doi.org/10.1371/journal.pone.0212442.g004>

Empirical measurement of vortex ring formation and entrainment

Vortex ring formation is a prominent phenomenon in engineered and biological systems as diverse as aerodynamic flow control, animal swimming, and the human cardiovascular system [41–43]. The growth and dynamics of vortex rings are dictated by the extent to which they entrain surrounding fluid [44]. Moreover, knowledge of the precise region of the flow that is ultimately entrained by a forming vortex ring can be used to predict how a vortex delivers mass, momentum, and energy to the surrounding flow. For example, pathological vortex ring formation in the human left ventricle has been shown to provide an effective diagnostic of heart failure [43]. Despite the importance of vortex ring entrainment, methods to quantify the region of the flow impacted by vortex rings have shown limited success, particularly in cases

for which the FTLE field cannot be calculated due to the sparsity of measurements. Here, we demonstrate the ability of the sCSC technique to precisely identify the region of a flow that is entrained by a forming vortex ring—knowledge that has been previously inaccessible in cases where measurement data is sparse, such as when the flow is interrogated using non-invasive clinical methods such as ultrasound or magnetic resonance imaging.

Vortex rings were formed in the laboratory using a piston-cylinder apparatus described in previous work [45]. A motor-driven piston pushes water through a vertical hollow cylinder of diameter $D = 2.49$ cm that is submerged in a tank with cross-sectional area of 61 cm by 61 cm and height of 91 cm. As the flow exits the cylinder at a nominal speed of 7 cm s^{-1} , the fluid boundary layer at the inner surface of the cylinder rolls up into a toroidal vortex ring, which propagates away from the cylinder via self-induction.

A set of 1174 fluid particle trajectories in the domain encountered by the vortex ring were analyzed using the present sCSC method and the dissimilarity measure in Eq (3) to identify regions of the ambient flow that were entrained by the vortex ring. As illustrated in Fig 5A, it is impossible to determine which fluid particles have been entrained by the vortex ring based on visual inspection of the trajectories alone. A comparison FTLE analysis performed by Schluter-Kuck and Dabiri on 30,500 advected particles (see Ref. [14], Fig 9) showed that 1174 trajectories are not sufficiently close to one another to compute the FTLE field, because the required gradient calculations are not well-posed for sparse trajectories. The alternative use of existing techniques based on heuristics, such as k -means or the spectral eigengap, rely on knowledge of the number of eddies to guide clustering; in the present case, it is not known *a priori* how many structures comprise the flow.

The sCSC dendrogram (Fig 5D) avoids the need for explicit determination of the number of eddies, as it unambiguously identifies the fluid particles entrained by the vortex ring as those belonging to branch 0. Branch 1 identifies all other particles and further bifurcations of that branch reveal underlying geometric symmetries, as in Branch 1 of the quadruple-eddy flow in Fig 2D.

Plots of the initial and final positions of the fluid particles in Fig 5B and 5C show that the fluid entrained by the vortex ring occupies a well-defined region in the immediate path of the vortex ring, a result that is consistent with intuition but that can now be characterized quantitatively for the first time. The void created by the evolution of the blue particles from Fig 5B to 5C is filled by the fluid ejected from the cylinder. The entrained blue particles ultimately occupy positions around the vortex ring that are consistent with the FTLE analysis in Ref. [14]. This provides another demonstration of the interpretability of the sCSC results: notably, these result have been achieved without any of the ad hoc assumptions required by existing methods of entrainment quantification [44, 46].

Visualizing macrostate modeling of molecular dynamics

In this section we highlight the interpretability of the sCSC dendrogram for a high-dimensional dynamical dataset. Specifically, we focus on an atomistic simulation of protein folding. Whereas fluid dynamics datasets typical represent only a few spatiotemporal coordinates, atomistic molecular dynamics (MD) datasets can contain thousands of degrees of freedom with complex interrelationships.

While MD is resource-intensive, advances in simulation parameters [47], bespoke hardware [48], and distributed computing frameworks [49], have enabled MD analyses to yield insight into complex biophysical systems at biologically meaningful timescales [50]. Thus, these simulations have the potential to uncover biophysical phenomena such as the misfolding

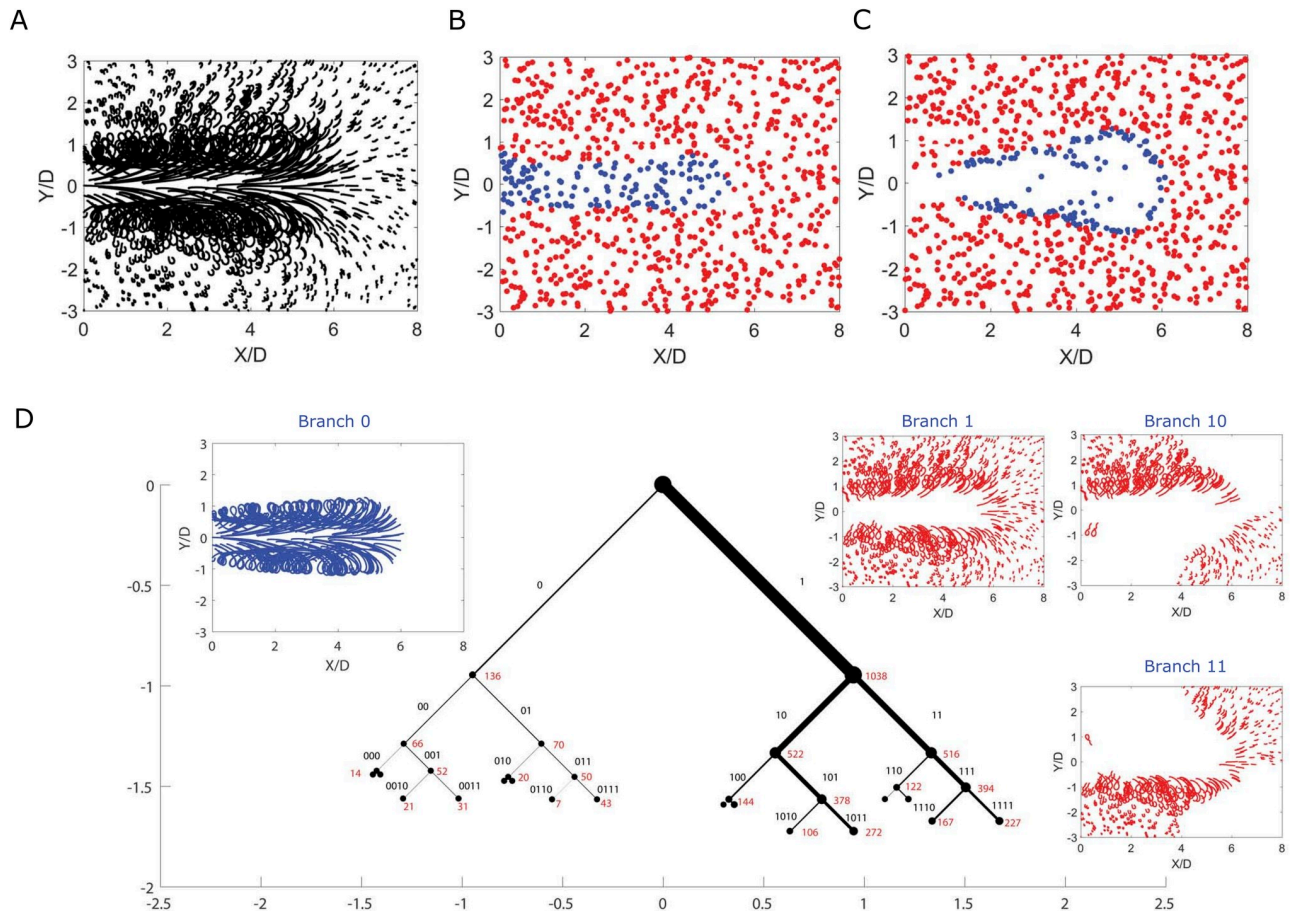


Fig 5. sCSC analysis of vortex ring formation. (A) Trajectories of 1174 Lagrangian particles initialized in the flow. (B) Initial positions of the 1174 particles. Blue particles are those revealed by the sCSC analysis to be entrained by the vortex ring; red particles are not entrained. (C) Final positions of the same particles tracked in panel B. The sCSC dendrogram on which this analysis is based is shown in panel (D), with every position occupied by all 1174 particles plotted in the corresponding inset branch plot. The width of each branch is proportional to the fraction of the states that it contains. The corresponding binary code of each branch is labeled in black text, and the number of trajectories associated with each node is labeled in red text. The dendrogram is plotted to the fourth eigenvector, although many of the labels are omitted for clarity. The horizontal and vertical axes are measured in units of the parameter z , and the branches are plotted at 45-degree angles. We have visualized the first 4 eigenvectors for brevity of presentation, but this does not affect the model results.

<https://doi.org/10.1371/journal.pone.0212442.g005>

mechanisms involved in a variety of diseases, stable configurations yet undiscovered by crystallography, and small molecule binding sites and kinetics for drug discovery.

However, without complementary analysis methods designed to communicate statistically rigorous and understandable conclusions resulting from such computational experiments, the benefits of advances in MD cannot be fully realized. While many methods have been developed to perform these analyses [50], it remains a challenge to display their results in a meaningful way. sCSC can be used to augment already-existing methods for analyzing MD simulations such that the results can be visualized and interpreted.

To demonstrate the use of sCSC to visualize an MD analysis, we use an ultralong MD simulation performed by Lindorff-Larsen *et al.* [51] of the folding and unfolding of Protein G, a 56-residue protein expressed in streptococcal bacteria. The simulation details are described in the Supporting Materials of Ref. [51]. We use a Markov state model (MSM) analysis to define the states of the system, which is a discrete approximation to the Perron-Frobenius operator [20]. This established mathematical framework codifies the system using a kinetic master

equation [50]. The master equation takes the form of a stochastic transition probability matrix, in which each state of the system is identified by a probability distribution of transitioning to every other state.

After constructing a quantitatively accurate and optimized MSM [22], we are interested in clustering these state into a smaller number of interpretable “macrostates”, since it is conceptually difficult to describe hundreds unique states in a physically interpretable way [52].

For our MSM, we found that 175 states optimally describes the system according to a variational evaluation (the MSM construction protocol is consistent with current best practices and is described in detail in the Methods). Minimum variance clustering analysis (MVCA), an effective coarse-graining method for MSMs, has recently been developed by one of the authors and uses a pairwise information theoretic dissimilarity metric in order to group states into a smaller number of macrostates, namely, the Jensen-Shannon divergence between the probability distributions characterizing the rows of the MSM transition probability matrix [28] (see also [Methods Eq \(4\)](#)).

By using the same pairwise dissimilarity metric as MVCA, the $\binom{175}{2}$ state adjacencies can be input into the sCSC algorithm to produce a visualization of a set of macrostates in the protein folding dataset, which is displayed in [Fig 6](#). Nine branches of the sCSC dendrogram are depicted by sampling one structure from each original MSM state contained in that branch. Since the nine depicted branches contain all 175 original MSM states, these branches can be interpreted as a possible set of system macrostates. By superimposing a representative conformation from each MSM state and coloring the protein according to its secondary structure, we can visualize the MD trajectory by interpreting the sCSC groupings.

First, we note that the folded structure (branch 0) is identified in the first sCSC solution and is separated from the denatured, unfolded ensemble, which comprises the rest of the dendrogram (branch 1). We see that the folded branch isolates a well-defined conformation with low variance across sampled conformations. The incorporation of subsequent sCSC eigenvectors identifies groups of structures unified by their protein secondary structure features. Various branches contain similar secondary structure elements (similar colors in the structure visualization in [Fig 6](#)), elucidating substructures exhibited during the folding of Protein G. For example, branch 1110 contains β -sheet secondary structure (yellow), whereas branch 11110 contains noticeable α -helical secondary structure (pink). Branch 1111111 is the least coherent, containing the most unstructured states. Summary statistics for each macrostate can be found in [S1 Table](#).

In this example, we have chosen to highlight secondary structure changes so we can understand which secondary structure elements characterize different subprocesses within folding. We see that the yellow β -sheet secondary structure appears in several macrostates—often along with the blue 3_{10} -helix, thought to be an intermediate structure during α -helix formation [53]—which might indicate that the formation of the pink α -helix is a rate-limiting step in the folding process. However, we could also choose to quantify and visualize macrostate contact maps, radii of gyration, or distance to various structures in order to gain complementary insight into the folding system.

For other dynamical processes characteristic of proteins, such as conformational change, allostery, and drug binding, we might choose to visualize parameters related to specific sites of interest or observables that can be probed experimentally. The choice of how to describe the macrostates is independent of the clustering process, but the depictions or statistics that enable the best interpretation of the system will depend on the dynamics of interest.

The sCSC dendrogram analysis offers advantages in the interpretation of high-dimensional MD datasets after an adequate kinetic model has been constructed. Utilizing the pairwise state

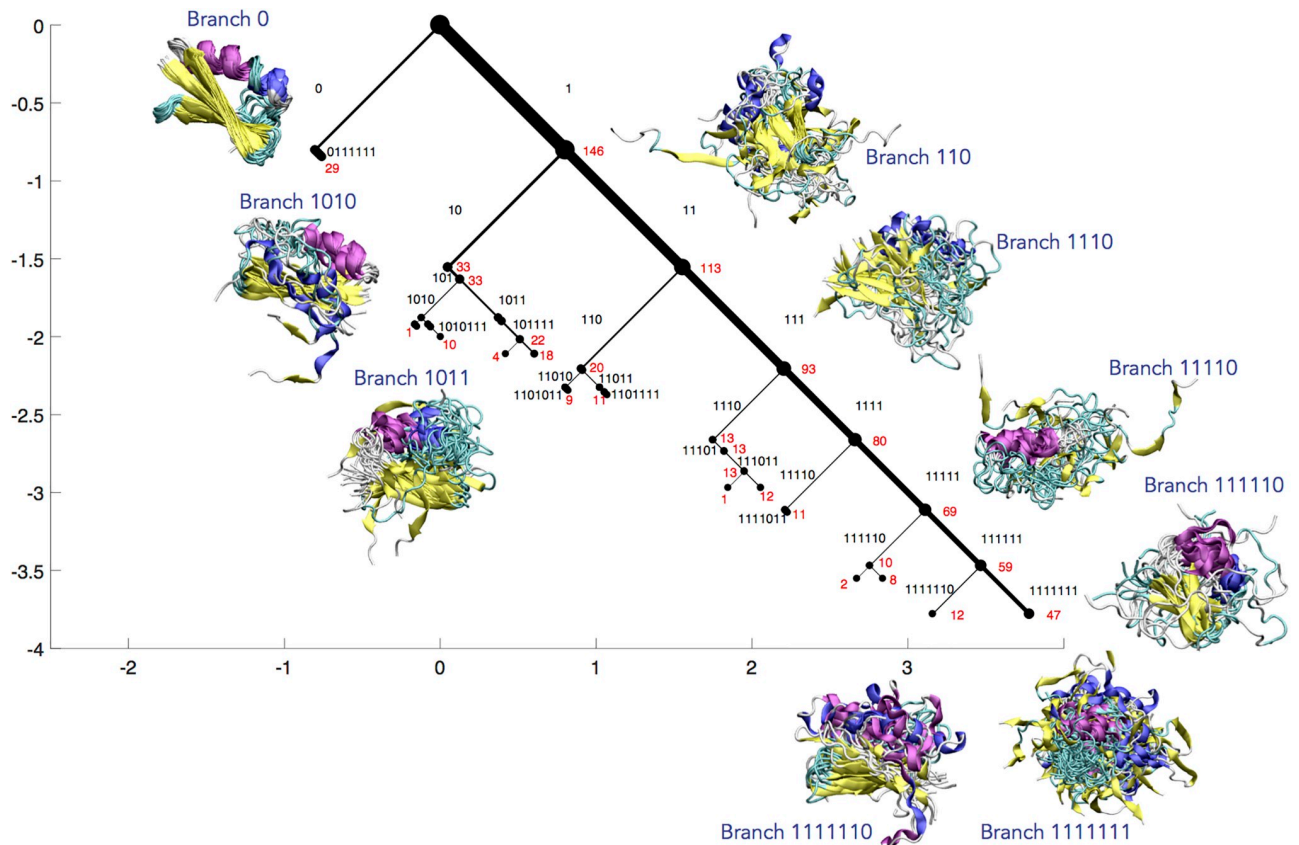


Fig 6. Atomistic protein folding simulation. The sCSC dendrogram for protein folding simulations of Protein G isolates the folded state and eight other macrostates that characterize the folding process. Each collective structure represents one sample drawn from each MSM state contained in the macrostate. The 9 macrostates contain all 175 original states, and are distinguished by consistency in secondary structure: pink regions represent α -helices, blue represents 3_{10} helices, and yellow represents β -sheet regions. Branch 0, which contains the folded macrostate, is the most ordered and is separated first. Branch 1111111 represents the most disordered macrostate. The other branches represent macrostates with varying degrees of order and secondary structure elements. The width of each branch is proportional to the fraction of the states that it contains. The corresponding binary code of each branch is labeled in black text, and the number of MSM states associated with each node is labeled in red text. The horizontal and vertical axes are measured in units of the parameter z . We have visualized the first 7 eigenvectors for brevity of presentation.

<https://doi.org/10.1371/journal.pone.0212442.g006>

adjacency from this kinetic model for an sCSC analysis produces a hierarchical representation of structural motifs according to the extent of their dissimilarity, which provides insight into the protein conformations that characterize subprocesses within folding. As in the fluid dynamics examples in the previous section, truncating the dendrogram when bifurcations are unoccupied produces an objective way to visualize the converged clusters. Finally, the generation of orthogonal sCSC solutions enables orthogonal dynamical processes in the protein folding simulation to be incorporated in analogy to the simple model in Fig 1. We anticipate that this type of interpretable visualization will be useful for communicating the results obtained from high-dimensional datasets.

Discussion

The present approach addresses the previously stated challenges with common clustering algorithms: it does not require a choice of cluster number or dendrogram cutting, it leverages the concept of dissimilarity in a computationally tractable way, and it maintains an interpretable hierarchical relationship among splittings.

Perhaps the most important advantage of this approach relative to commonly used tools is that the number, shape, and size of clusters in the data emerges naturally from the sCSC dendrogram rather than being specified *a priori*. As the set of eigenvectors that is included in the analysis is increased to include those associated with lesser eigenvalues, the number of unoccupied binary codes generally increases. This is because progressively fewer groups of states that survived the preceding orthogonal partitions will be subsequently separated at lower levels of the dendrogram. In this way, the set of clusters in the dataset is revealed to be those branches of the dendrogram structure whose shape converges as the number of eigenvectors included in the analysis increases. The sCSC dendrogram indicates not only the number of these converged clusters but also the relative strength of the partitions between clusters, via the length of the connecting branches in z -space.

While sCSC conceptually resembles divisive hierarchical clustering, the number of possible divisions in the latter scales as $2^{c-1} - 1$ with the number of clusters c , which is generally not feasible for large c unless the initial dataset is sparse [54]. sCSC scales in the same way as agglomerative hierarchical clustering, requiring a computationally nontrivial but tractable calculation of $\binom{n}{2}$ dissimilarity values for n initial data points. However, unlike agglomerative clustering—which also requires the calculation of $\binom{n}{2}$ dissimilarities—small differences between states at lower levels of the sCSC dendrogram have no impact on the clusters that form at higher levels, as the top-down approach begins by using the most significant partitions indicated by the eigenvectors associated with the largest eigenvalues.

When applying sCSC, domain knowledge should inform selection of an appropriate dissimilarity measure, but *ad hoc* and *a priori* assumptions about the structure of the data itself are not needed. While we have demonstrated sCSC only for simulated physical systems, we anticipate that these features will make sCSC a powerful tool for interrogating both new and longstanding research problems, including those in fields where the underlying processes are less accessible, such as genomics and neuroscience. For example, genetic ancestries can potentially be clustered on the basis of the sCSC structure that emerges from the dissimilarity of single-nucleotide polymorphisms (SNPs) among individuals within a population. In the latter case, differences in neuronal activation can be amplified using sCSC to identify emergent functions that involve coordination of spatially distant neurons. These and other applications can be pursued immediately given the tools developed here.

Methods

Quadruple-eddy ocean flow model

The velocity field of the quadruple-eddy ocean flow model is given by,

$$\begin{aligned} \frac{dx}{dt} &= -\pi A \sin(\pi f) \cos(\pi y) \\ \frac{dy}{dt} &= -\pi A \cos(\pi f) \cos(\pi y)(2ax + b), \end{aligned}$$

where $x = [0, 2]$ and $y = [-1, 1]$ are the dimensionless east-west and north-south spatial coordinates (i.e. normalized by the quadrant side length), t is time in dimensionless units, and

$$\begin{aligned} a &= \epsilon \sin(\omega t) \\ b &= 1 - 2\epsilon \sin(\omega t) \\ f &= ax^2 + bx. \end{aligned}$$

In the present unsteady implementation of the model, $A = 0.1$, $\epsilon = 0.1$, and $\omega = 2\pi/10$. 3000 Lagrangian drifters were randomly initialized in the domain and advected in the flow using a fifth-order Runge-Kutta integration scheme. The duration of advection was 40 dimensionless time units, corresponding to 4 periods of horizontal oscillation of the flow.

Bickley jet atmospheric model

The Bickley jet flow is given by the streamfunction $\psi = \psi_0 + \psi_1$, where,

$$\begin{aligned}\psi_0 &= c_3 y - UL \tanh(y/L) \\ \psi_1 &= UL \operatorname{sech}^2(y/L) \sum_{n=1}^3 \epsilon_n \cos(k_n(x - \sigma_n t)).\end{aligned}$$

In the present study, we use similar parameter values as in Ref. [13]: $U = 62.66 \text{ ms}^{-1}$, $L = 1770 \text{ km}$, $k_n = 2n/r_0$, $c = [0.1446U, 0.205U, 0.461U]$, $\sigma = c - c(3)$, and $\epsilon = [0.0075, 0.15, 0.3]$, and the flow is computed on the interval $x = [0, 20 \times 10^6] \text{ m}$, $y = [-3 \times 10^6, 3 \times 10^6] \text{ m}$, over the time interval $t = [0, 40]$ days, divided into 601 discrete time steps. The flow was treated as periodic in x . 3000 Lagrangian fluid particles were randomly initialized in the domain and advected in the flow using a fifth-order Runge-Kutta integration scheme.

Markov state models

Markov state models (MSMs) are a kinetic master equation framework for describing and analyzing time-series data such as molecular dynamics (MD) simulations by approximating the continuous Perron-Frobenius operator using a discrete transition probability matrix [20]. A MSM requires partitioning the phase space explored by a system into discrete states (henceforth “microstates”), and is represented by a transition probability matrix defined for a Markovian lag time τ at which transitions between the microstates are independent of the history of the system. For protein folding analyses, phase space (positions and velocities) is conventionally approximated by conformation space (positions), and states are chosen according to an objective optimization protocol, in this case a variational principle [22]. The Markovian lag time chosen to analyze a system must be long enough for memoryless inter-state transitions, but short enough to resolve dynamics; for protein folding dynamics, lag times on the order of tens of nanoseconds are typical.

The MSM transition probability matrix is constrained to be stochastic, symmetric with respect to a stationary distribution, ergodic, and aperiodic. It can thus be decomposed into eigenvalues and eigenvectors, $T(\tau)\lambda = \psi\lambda$, where the eigenvalues are on the unit interval $|\lambda_i| \leq 1$ and the highest eigenvalue $\lambda_1 = 1$ is unique. The variational principle states that the sum of estimated eigenvalues is bounded from above by the sum of true eigenvalues; thus many state decompositions can be tested according to the sum of a set number of eigenvalues for a set Markovian lag time and the state decomposition resulting in the highest sum of approximated eigenvalues can be chosen for further analysis.

The MSM for the simulation analyzed in this work was constructed according to the protocol used in Ref. [55] for a set lag time of 50 ns according to a previous analysis for the same system performed in Ref. [56]. First, the Cartesian coordinates from the raw simulation data are transformed into the sines and cosines of the ϕ and ψ side chain dihedral angles for each amino acid of the protein. Next, the vector of dihedrals is again transformed using time structure-based independent component analysis (tICA) [57] with a tICA lag time of 50 ns wherein each tICA solution vector was weighted according to its associated eigenvalue [58]. Then, mini-batch k -means was used to cluster the simulation frames according to their weighted

tICA representations for 265 different numbers of cluster centers randomly chosen between 10 and 5000. Finally, a MSM was constructed on each k -means state decomposition in which the transition probability matrix is obtained using a maximum likelihood estimator of the data such that detailed balance is achieved. For each model, five MSMs were fit to a randomly chosen half of the data and then applied to the other half of the data, and the latter was used to sum the first 50 MSM eigenvalues as that model's score. The winning model was chosen to be the one that achieved the single maximum score from parameter sets with mean scores within one standard deviation of the maximum mean score. For our analysis of 265 different microstate numbers, the best model according to this variational analysis had 175 microstates and was used for analysis in the main text.

Coarse-graining MSMs with MVCA

Minimum variance clustering analysis (MVCA) was recently published by one of the authors as an algorithm for coarse-graining an MSM transition probability matrix into a smaller number of macrostates by grouping the original microstates [28]. MVCA achieves a coarse-grained model by using agglomerative hierarchical clustering with Ward's minimum variance method [59] to cluster the microstates, where the pairwise dissimilarity between microstates is quantified using an information theoretic measure between the probability distribution characterized by the corresponding row of the MSM transition matrix.

If two microstates are defined by transition probability distributions P and Q , their pairwise dissimilarity can be written using the Jensen-Shannon divergence [27],

$$\text{div}_{\text{JS}}(P||Q) = \frac{1}{2} \sum_i P_i \log \frac{P_i}{M_i} + \frac{1}{2} \sum_i Q_i \log \frac{Q_i}{M_i} \quad (4)$$

where M is the elementwise mean of P and Q , and each term is the Kullback-Leibler divergence to the mean. We quantify the dissimilarity between microstates using the square root of Eq 4 [28, 60].

From this set of pairwise similarities, MVCA goes on to cluster the microstates using agglomerative hierarchical clustering with Ward's method. In the analysis presented in this work, the set of pairwise dissimilarities is instead used to construct the adjacency matrix for sCSC.

Supporting information

S1 Fig. sCSC dendrogram applied to random data. To evaluate the sCSC dendrogram structure resulting from random noise, an adjacency matrix was constructed based on 3000 two-dimensional trajectories whose instantaneous positions over 2000 time steps were selected randomly from uniform distributions over the spatial coordinate intervals $x = (0, 1)$, $y = (0, 1)$. These states were analyzed using pairwise dissimilarity based on the normalized standard deviation. The result is a single main branch with a small splintering of trajectories at each eigenvector level. The splintering at each level converges throughout the seven eigenvectors included in the analysis. The width of each branch is proportional to the fraction of the states that it contains. The corresponding binary code of each branch is labeled in black text, and the number of trajectories associated with each node is labeled in red text. (TIFF)

S1 Table. Macrostate statistics for Protein G. Summary statistics for the nine macrostates identified from the sCSC model of Protein G. The first column identifies the branch label in Fig 6. The second column shows the number of the original 175 MSM states contained in the

macrostate. For the next three columns, statistics are gathered using the one sampled state from each macrostate that is used for visualization in Fig 6. The third column reports the average \pm one standard deviation pairwise RMSD over all atoms for all pairs of original MSM states within the same macrostate (e.g. the average RMSD for Branch 0, which contains 29 microstates, each of which is represented by one sampled structure from each original MSM state, is the average of the pairwise RMSD for $\binom{29}{2} = 406$ possible pairs). The fourth and fifth columns report the average percentage \pm one standard deviation of the α -helical and β -sheet secondary structure for each original MSM structure sample in the macrostate according to the simplified dictionary of protein secondary structure (DSSP) protocol [61] implemented in the MDTraj [62] software package. The last column indicates which Protein G residues were aligned to create the superpositions illustrated in Fig 6 for each branch. (PDF)

Acknowledgments

The authors are grateful to Muneeb Sultan, Jared Dunnmon, Nicole Xu, and the referees for insightful manuscript feedback and to D. E. Shaw Research for providing the Protein G dataset.

Author Contributions

Conceptualization: Brooke E. Husic, John O. Dabiri.

Data curation: Kristy L. Schlueter-Kuck, John O. Dabiri.

Formal analysis: Brooke E. Husic, John O. Dabiri.

Investigation: Brooke E. Husic, John O. Dabiri.

Methodology: Brooke E. Husic, John O. Dabiri.

Project administration: John O. Dabiri.

Resources: John O. Dabiri.

Software: Brooke E. Husic, John O. Dabiri.

Supervision: John O. Dabiri.

Validation: Brooke E. Husic, John O. Dabiri.

Visualization: Brooke E. Husic, John O. Dabiri.

Writing – original draft: Brooke E. Husic, John O. Dabiri.

Writing – review & editing: Brooke E. Husic, Kristy L. Schlueter-Kuck, John O. Dabiri.

References

1. Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, et al. MoleculeNet: a benchmark for molecular machine learning. *Chem Sci*. 2018; 9(2):513–530. <https://doi.org/10.1039/c7sc02664a> PMID: 29629118
2. Kavvas ES, Catoi E, Mih N, Yurkovich JT, Seif Y, Dillon N, et al. Machine learning and structural analysis of Mycobacterium tuberculosis pan-genome identifies genetic signatures of antibiotic resistance. *Nat Commun*. 2018; 9(1):4306. <https://doi.org/10.1038/s41467-018-06634-y> PMID: 30333483
3. Froyland G, Padberg-Gehle K. A rough-and-ready cluster-based approach for extracting finite-time coherent sets from sparse and incomplete trajectory data. *Chaos*. 2015; 25:087406. <https://doi.org/10.1063/1.4926372> PMID: 26328577

4. Friedman J, Hastie T, Tibshirani R. The elements of statistical learning. vol. 1. Springer series in statistics New York, NY, USA.; 2001.
5. Macqueen J. Some methods for classification and analysis of multivariate observations. In: 5th Berkeley Symposium on Mathematical Statistics and Probability; 1967. p. 281–297.
6. Dunn JC. A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters. *J Cybern.* 1973; 3(3):32–57. <https://doi.org/10.1080/01969727308546046>
7. Ester M, Kriegel HP, Sander J, Xu X. A density-based algorithm for discovering clusters in large spatial databases with noise. In: Proceedings of the Second ACM SIGKDD International Conference on Knowledge, Discovery, and Data Mining. vol. 96; 1996. p. 226–231.
8. Newman MEJ, Girvan M. Finding and evaluating community structure in networks. *Phys Rev E.* 2004; 69:026113.
9. Kaufman L, Rousseeuw PJ. Finding groups in data: an introduction to cluster analysis. vol. 344. John Wiley & Sons; 2009.
10. Ali T, Asghar S, Sajid NA. Critical analysis of DBSCAN variations. In: 2010 International Conference on Information and Emerging Technologies; 2010. p. 1–6.
11. Allshouse MR, Peacock T. Lagrangian based methods for coherent structure detection. *Chaos.* 2015; 25:097617. <https://doi.org/10.1063/1.4922968> PMID: 26428570
12. Hadjighasem A, Farazmand M, Blazeovski D, Froyland G, Haller G. A critical comparison of Lagrangian methods for coherent structure detection. *Chaos.* 2017; 27:053104. <https://doi.org/10.1063/1.4982720> PMID: 28576102
13. Hadjighasem A, Karrasch D, Teramoto H, Haller G. Spectral-clustering approach to Lagrangian vortex detection. *Phys Rev E.* 2016; 93:063107. <https://doi.org/10.1103/PhysRevE.93.063107> PMID: 27415358
14. Schlueter-Kuck KL, Dabiri JO. Coherent structure colouring: identification of coherent structures from sparse data using graph theory. *J Fluid Mech.* 2017; 811:468–486. <https://doi.org/10.1017/jfm.2016.755>
15. Schlueter-Kuck KL, Dabiri JO. Identification of individual coherent sets associated with flow trajectories using coherent structure coloring. *Chaos.* 2017; 27(9):091101. <https://doi.org/10.1063/1.4993862> PMID: 28964141
16. Klus S, Koltai P, Schütte C. On the numerical approximation of the Perron-Frobenius and Koopman operator. *J Comput Dynam.* 2016; 3(1):51–79.
17. Dellnitz M, Junge O. On the approximation of complicated dynamical behavior. *SIAM J Numer Anal.* 1999; 36(2):491–515. <https://doi.org/10.1137/S0036142996313002>
18. Froyland G, Dellnitz M. Detecting and locating near-optimal almost-invariant sets and cycles. *SIAM J Sci Comput.* 2003; 24(6):1839–1863. <https://doi.org/10.1137/S106482750238911X>
19. Mezić I. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dyn.* 2005; 41(1-3):309–325. <https://doi.org/10.1007/s11071-005-2824-x>
20. Schütte C, Fischer A, Huisinga W, Deuffhard P. A direct approach to conformational dynamics based on hybrid Monte Carlo. *J Comput Phys.* 1999; 151(1):146–168. <https://doi.org/10.1006/jcph.1999.6231>
21. Deuffhard P, Huisinga W, Fischer A, Schütte C. Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains. *Linear Algebra Appl.* 2000; 315(1):39–59. [https://doi.org/10.1016/S0024-3795\(00\)00095-1](https://doi.org/10.1016/S0024-3795(00)00095-1)
22. Noé F, Nüske F. A Variational Approach to Modeling Slow Processes in Stochastic Dynamical Systems. *Multiscale Model Simul.* 2013; 11(2):635–655. <https://doi.org/10.1137/110858616>
23. Williams MO, Kevrekidis IG, Rowley CW. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *J Nonlinear Sci.* 2015; 25(6):1307–1346. <https://doi.org/10.1007/s00332-015-9258-5>
24. Wu H, Nüske F, Paul F, Klus S, Koltai P, Noé F. Variational Koopman models: slow collective variables and molecular kinetics from short off-equilibrium simulations. *J Chem Phys.* 2017; 146(15):154104. <https://doi.org/10.1063/1.4979344> PMID: 28433026
25. Guckenheimer J, Holmes P. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields.* vol. 42. Springer; 1983.
26. Haller G, Yuan G. Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D.* 2000; 147(3-4):352–370.
27. Lin J. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory.* 1991; 37(1):145–151. <https://doi.org/10.1109/18.61115>

28. Husic BE, McKiernan KA, Wayment-Steele HK, Sultan MM, Pande VS. A minimum variance clustering approach produces robust and interpretable coarse-grained models. *J Chem Theory Comput.* 2018; 14(2):1071–1082. <https://doi.org/10.1021/acs.jctc.7b01004> PMID: 29253336
29. Hall KM. An r-Dimensional Quadratic Placement Algorithm. *Management Sci.* 1970; 17(3):219–229. <https://doi.org/10.1287/mnsc.17.3.219>
30. Munoz S, Ortuno MT, Ramirez J, Yanez J. Coloring fuzzy graphs. *Omega.* 2005; 33:211–221. <https://doi.org/10.1016/j.omega.2004.04.006>
31. Press WH, Teukolsky SA, Vetterling WT, Flannery BP. *Numerical Recipes.* Cambridge University Press; 2007.
32. Müllner D. fastcluster: Fast hierarchical, agglomerative clustering routines for R and Python. *J Stat Softw.* 2013; 53(9):1–18.
33. Hansen P, Jaumard B. Cluster analysis and mathematical programming. *Math Program.* 1997; 79. [https://doi.org/10.1016/S0025-5610\(97\)00059-2](https://doi.org/10.1016/S0025-5610(97)00059-2)
34. Huhn F, von Kameke A, Perez-Munuzuri V, Olascoaga MJ, Beron-Vera FJ. The impact of advective transport by the South Indian Ocean Countercurrent on the Madagascar plankton bloom. *Geophys Res Lett.* 2012; 39(6):L06602. <https://doi.org/10.1029/2012GL051246>
35. Hughes CW, Miller PI. Rapid water transport by long-lasting modon eddy pairs in the southern midlatitude oceans. *Geophys Res Lett.* 2017; 44(24). <https://doi.org/10.1002/2017GL075198>
36. Kaiser E, Noack BR, Cordier L, Spohn A, Segond M, Abel M, et al. Cluster-based reduced-order modeling of a mixing layer. *J Fluid Mech.* 2014; 754:365–414. <https://doi.org/10.1017/jfm.2014.355>
37. Maclean J, Santitissadeekorn N, Jones CKRT. A coherent structure approach for parameter estimation in Lagrangian Data Assimilation. *Physica D.* 2017; 360:36–45. <https://doi.org/10.1016/j.physd.2017.08.007>
38. Sengupta A, Carrara F, Stocker R. Phytoplankton can actively diversify their migration strategy in response to turbulent cues. *Nature.* 2017; 543:555–558. <https://doi.org/10.1038/nature21415> PMID: 28297706
39. Argo. Argo float data and metadata from Global Data Assembly Centre (Argo GDAC). 2000;.
40. Rypina II, Brown MG, Beron-Vera FJ, Kocak H, Olascoaga MJ, Udovydchenkov IA. On the Lagrangian Dynamics of Atmospheric Zonal Jets and the Permeability of the Stratospheric Polar Vortex. *J Atmos Sci.* 2007; 64:3595–3610. <https://doi.org/10.1175/JAS4036.1>
41. Krueger PS, Gharib M. The significance of vortex ring formation to the impulse and thrust of a starting jet. *Phys Fluids.* 2003; 15:1271–1281. <https://doi.org/10.1063/1.1564600>
42. Dabiri JO. Optimal vortex formation as a unifying principle in biological propulsion. *Annu Rev Fluid Mech.* 2009; 41:17–33. <https://doi.org/10.1146/annurev.fluid.010908.165232>
43. Gharib M, Rambod E, Kheradvar A, Sahn DJ, Dabiri JO. Optimal vortex formation as an index of cardiac health. *Proc Natl Acad Sci USA.* 2006; 103:6305–6308. <https://doi.org/10.1073/pnas.0600520103> PMID: 16606852
44. Dabiri JO, Gharib M. Fluid entrainment by isolated vortex rings. *J Fluid Mech.* 2004; 511:311–331. <https://doi.org/10.1017/S0022112004009784>
45. Schlueter-Kuck KL, Dabiri JO. Pressure evolution in the shear layer of forming vortex rings. *Phys Rev Fluids.* 2016; 1:012501(R). <https://doi.org/10.1103/PhysRevFluids.1.012501>
46. Olcay AB, Krueger PS. Measurement of ambient fluid entrainment during vortex ring formation. *Exp Fluids.* 2008; 44:235–247. <https://doi.org/10.1007/s00348-007-0397-9>
47. Wang LP, McKiernan KA, Gomes J, Beauchamp KA, Head-Gordon T, Rice JE, et al. Building a More Predictive Protein Force Field: A Systematic and Reproducible Route to AMBER-FB15. *J Phys Chem B.* 2017; 121(16):4023–4039. <https://doi.org/10.1021/acs.jpcc.7b02320> PMID: 28306259
48. Shaw DE, Maragakis P, Lindorff-Larsen K, Piana S, Dror RO, Eastwood MP, et al. Atomic-Level Characterization of the Structural Dynamics of Proteins. *Science.* 2010; 330(6002):341–346. <https://doi.org/10.1126/science.1187409> PMID: 20947758
49. Shirts M, Pande VS. Screen Savers of the World Unite! *Science.* 2000; 290(5498):1903–1904. <https://doi.org/10.1126/science.290.5498.1903> PMID: 17742054
50. Husic BE, Pande VS. Markov State Models: From an Art to a Science. *J Am Chem Soc.* 2018; 140(7):2386–2396. <https://doi.org/10.1021/jacs.7b12191> PMID: 29323881
51. Lindorff-Larsen K, Piana S, Dror RO, Shaw DE. How Fast-Folding Proteins Fold. *Science.* 2011; 334(6055):517–520. <https://doi.org/10.1126/science.1208351> PMID: 22034434
52. Pande VS, Beauchamp K, Bowman GR. Everything you wanted to know about Markov State Models but were afraid to ask. *Methods.* 2010; 52(1):99–105. <https://doi.org/10.1016/j.ymeth.2010.06.002> PMID: 20570730

53. Sundaralingam M, Sekharudu Y. Water-inserted α -helical segments implicate reverse turns as folding intermediates. *Science*. 1989; 244:1333–1337. <https://doi.org/10.1126/science.2734612> PMID: 2734612
54. Boley D. Principal Direction Divisive Partitioning. *Data Min Knowl Discov*. 1998; 2(4):325–344. <https://doi.org/10.1023/A:1009740529316>
55. Husic BE, Pande VS. Ward Clustering Improves Cross-Validated Markov State Models of Protein Folding. *J Chem Theory Comput*. 2017; 13(3):963–967. <https://doi.org/10.1021/acs.jctc.6b01238> PMID: 28195713
56. Beauchamp KA, McGibbon R, Lin YS, Pande VS. Simple few-state models reveal hidden complexity in protein folding. *Proc Natl Acad Sci*. 2012; 109(44):17807–17813. <https://doi.org/10.1073/pnas.1201810109> PMID: 22778442
57. Schwantes CR, Pande VS. Improvements in Markov State Model Construction Reveal Many Non-Native Interactions in the Folding of NTL9. *J Chem Theory Comput*. 2013; 9(4):2000–2009. <https://doi.org/10.1021/ct300878a> PMID: 23750122
58. Noé F, Clementi C. Kinetic Distance and Kinetic Maps from Molecular Dynamics Simulation. *J Chem Theory Comput*. 2015; 11(10):5002–5011. <https://doi.org/10.1021/acs.jctc.5b00553> PMID: 26574285
59. Ward JH Jr. Hierarchical grouping to optimize an objective function. *J Amer Statist Assoc*. 1963; 58(301):236–244. <https://doi.org/10.1080/01621459.1963.10500845>
60. Endres DM, Schindelin JE. A new metric for probability distributions. *IEEE Transactions on Information Theory*. 2003; 49(7):1858–1860. <https://doi.org/10.1109/TIT.2003.813506>
61. Kabsch W, Sander C. Dictionary of protein secondary structure: Pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers*. 1983; 22(12):2577–2637. <https://doi.org/10.1002/bip.360221211> PMID: 6667333
62. McGibbon RT, Beauchamp KA, Harrigan MP, Klein C, Swails JM, Hernández CX, et al. MDTraj: A Modern Open Library for the Analysis of Molecular Dynamics Trajectories. *Biophys J*. 2015; 109(8):1528–1532. <https://doi.org/10.1016/j.bpj.2015.08.015> PMID: 26488642