


Article

Uncertainty Quantification of Trajectory Clustering Applied to Ocean Ensemble Forecasts

Guilherme S. Vieira ¹, Irina I. Rypina ² and Michael R. Allshouse ^{1,*}

¹ Department of Mechanical and Industrial Engineering, Northeastern University, Boston, MA 02115, USA; salvadorvieira.g@northeastern.edu

² Department of Physical Oceanography, Woods Hole Oceanographic Institution, Woods Hole, MA 02543, USA; irypina@whoi.edu

* Correspondence: m.allshouse@northeastern.edu; Tel.: +1-617-373-2138

Received: 21 August 2020; Accepted: 11 October 2020; Published: 17 October 2020



Abstract: Partitioning ocean flows into regions dynamically distinct from their surroundings based on material transport can assist search-and-rescue planning by reducing the search domain. The spectral clustering method partitions the domain by identifying fluid particle trajectories that are similar. The partitioning validity depends on the accuracy of the ocean forecasting, which is subject to several sources of uncertainty: model initialization, limited knowledge of the physical processes, boundary conditions, and forcing terms. Instead of a single model output, multiple realizations are produced spanning a range of potential outcomes, and trajectory clustering is used to identify robust features and quantify the uncertainty of the ensemble-averaged results. First, ensemble statistics are used to investigate the cluster sensitivity to the spectral clustering method free-parameters and the forecast parameters for the analytic Bickley jet, a geostrophic flow model. Then, we analyze an operational coastal ocean ensemble forecast and compare the clustering results to drifter trajectories south of Martha's Vineyard. This approach identifies regions of low uncertainty where drifters released within a cluster predominantly remain there throughout the window of analysis. Drifters released in regions of high uncertainty tend to either enter neighboring clusters or deviate from all predicted outcomes.

Keywords: Lagrangian transport; spectral clustering; uncertainty quantification; parameter sensitivity; ocean ensemble forecast; drifter data; search-and-rescue

1. Introduction

Fluid flows, even if unsteady and aperiodic, may admit persistent patterns generally referred to as coherent structures that reveal flow characteristics related to the transport of fluid particles [1–3]. Coherent structures of the elliptic type [4–6] are portions of fluid that do not significantly mix with the rest of the domain. From a Lagrangian perspective, the perimeter delimiting material within these structures remains nearly constant as they move [7,8], and fluid can be transported over long distances while surrounded by more vigorous mixing [2,9,10]. Eulerian methods compute coherent structures directly from instantaneous velocity fields and may not highlight features that have a persistent impact on transport. Additionally, these methods require full velocity fields, which may be unavailable or hard to reconstruct from sparse sets of observed trajectories. Lagrangian methods, in turn, can identify coherent structures based on the trajectories themselves and identify the dominant features over a given time interval [2,7,10–12].

One approach for identifying these coherent structures uses cluster analysis. Clustering algorithms reveal underlying structures in data sets by partitioning the data so that similar elements are assigned to the same cluster, while dissimilar elements are assigned to different ones. These algorithms have

been extensively studied and widely applied in image segmentation and anomaly detection, as well as biological and physical processes [13,14]. For fluid transport analysis, clustering techniques can efficiently identify elliptic structures when applied to fluid particle trajectories [5,15], which we refer to as trajectory clustering.

Here, we analyze particle trajectories using the spectral clustering algorithm [16–19]. This method has been used to identify coherent structures in analytic and simulated flows [15,20,21]. The clustering performs a systematic partitioning of the trajectories into coherent and incoherent sets, providing a conceptual simplification of the underlying dynamical system for a general flow. This method is also frame-invariant, and hence the identified structures are the same in all frames that translate or rotate relative to each other. One drawback of the spectral clustering method, however, is that there are a number of free-parameters that the user has to select, which could impact the results [20].

While ocean drifter data can be clustered *a posteriori* to identify coherent structures, an *a priori* analysis to predict coherent structures relies on trajectories obtained from a forecast model. The clustering accuracy is limited by that of the trajectories, which, in turn, is limited by the accuracy of the velocity model. In ocean modeling, however, several sources of uncertainties pose a challenge to the Lagrangian approach [22,23]. To simplify ocean models and reduce computational expenses, the governing equations are only resolved on a restricted range of spatial and temporal scales, and the influence of scales outside this window is either parameterized or neglected. Uncertainties also arise from the limited knowledge of processes within the scale window, which require approximate representations or parameterizations. Moreover, measurements used for model initialization and parameter estimation are limited in coverage and accuracy, leading to imprecise initial conditions and model parameters. Finally, models of interactions between the ocean and the atmosphere are approximate, and ocean boundary conditions are inexact. All of the above lead to differences between the actual values and the modeled values of physical fields and properties.

To account for model uncertainties intrinsic to the modeling process, an effective option is to perform ensemble statistics. Different sets of model parameter values generate an ensemble of possible outcomes, which are then processed to provide probabilistic information about the variability on the end results. Search-and-rescue planning, for instance, already considers ensemble statistics to produce probability-distribution maps for the target's location [24]. The vast uncertain parameter space together with the continuous motion of floating objects driven by unsteady flows, however, can lead to error accumulation in the predicted trajectories [25]. Coherent structures have been shown to depend less on individual trajectories and are more robust to model parameter variations and noise, highlighting the main structures in the flow even in the event of imperfect or scarce trajectories [23,26]. The robustness of the clustering algorithm to perturbations in the individual trajectories via an ensemble set of realizations has yet to be studied and could aid in quantifying uncertainty for the trajectory clusters.

The detection of robust clusters could then be used in the deployment of drifters to observe such features in the ocean. Other Lagrangian methods have been used in the past in experiments to interpret the observed behavior of drifters deployed in the ocean [27–33]. In these studies, drifters were released, tracked, and their trajectories were compared to results from Lagrangian analyses performed *a posteriori* using velocities from models, satellite altimetry, or radar measurements. Few studies, however, have used Lagrangian methods to plan and execute field experiments [25,34,35]. The drifter-release experiment presented in [36], whose drifter data are used in this paper, is one of the first experiments targeting coherent structures based on an *a priori* trajectory clustering obtained from forecast simulations.

Our approach accounts for and quantifies uncertainty when clustering trajectories. We apply the spectral clustering algorithm varying the method free-parameters to understand how the clustering results are sensitive to the implementation, analyze ensemble simulations to understand how the model parameters impact the resulting clusters, and finally apply the method to a forecast data set to compare the clustering prediction to experimental drifter data. The method by Hadjighasem et al. [15] is modified with a more broadly used similarity function and soft clustering membership probabilities,

allowing for a probabilistic view of the clusters for each individual model realization. Two different systems are analyzed with our approach: an analytic flow model and forecast simulations provided by a coastal ocean model. First, the Bickley jet analytic flow [37,38] with model parameter variations is used to mimic model uncertainty. Then, we analyze ensemble-forecasts generated by an ocean model [39,40] of the coastal region of the Martha's Vineyard island. Ensemble statistics of the resulting clusters provide a probabilistic view of the coherent structures identified by the method. The forecast clustering results are used to identify coherent structures in a drifter-release experiment, and the drifter trajectories are compared to the forecast cluster behavior and associated uncertainty.

The paper is organized as follows. Section 2 presents the spectral clustering method used in this work and how clustering results are processed to provide ensemble statistics. Section 3 introduces the quasi-periodic Bickley jet system, performs ensemble statistics for a single set of particle trajectories while varying the method free-parameters to quantify the sensitivity of the clustering results to the implementation. Section 4 varies Bickley jet system model parameters to assess the robustness of the identified coherent structures in a scenario of uncertainty in the velocity field. Section 5 presents the clustering analysis used to target coherent structures in a drifter-release experiment south of Martha's Vineyard, and compares the forecast results to the observed drifter trajectories. Finally, conclusions are presented in Section 6.

2. Method

This section presents the spectral clustering method and how uncertainty is measured. We adapt the method of Hadjighasem et al. [15] to use soft clustering (fuzzy c -means) to assign a cluster membership probability to each particle trajectory. Section 2.1 introduces the method along with specifics of the similarity measure selected and the soft clustering approach. Section 2.2 presents the general procedure we use to calculate statistics and quantify uncertainty when considering an ensemble of possible clustering results.

2.1. Spectral Clustering Method with Soft Memberships for Trajectory Clustering

To analyze flows from a Lagrangian perspective, we consider massless fluid particles that move with the velocity field $\mathbf{u}(t, \mathbf{x})$. The trajectory $\mathbf{x}_i(t)$ of fluid particle i departing from $\mathbf{x}_i(t_0)$ at time $t = t_0$ is

$$\mathbf{x}_i(t) = \mathbf{x}_i(t_0) + \int_{t_0}^t \mathbf{u}(\tau, \mathbf{x}_i(\tau)) d\tau. \quad (1)$$

A total of N particles are initialized in a grid that uniformly covers the targeted domain at t_0 , and their individual trajectories are numerically integrated and output at discrete time instances within the time interval $[t_0, t_f]$.

We use the set of trajectories $\{\mathbf{x}_i(t)\}_{1 \leq i \leq N}$ to partition the spatial domain into clusters. The spectral clustering algorithm performs an eigenanalysis to project the trajectory set onto a subspace that may yield clusters maximizing the within-cluster similarity and minimizing the between-clusters similarity [41]. Particles clustered together should move as a compact group, with limited mixing with particles outside of the cluster, while particles in different clusters should experience dissimilar trajectories [5,15].

The spectral analysis requires the construction of a positive, weighted, undirected graph known as the similarity graph. This graph quantifies the pairwise similarities between trajectories. Each graph node represents a particle trajectory, and the graph edges are weighted by the similarity between the nodes they connect [15,16,18]. To compute these weights, one must first define a measure for similarity between trajectories. One possible metric for the dissimilarity between trajectories \mathbf{x}_i and \mathbf{x}_j is their time-averaged distance

$$r_{ij} = \frac{1}{t_f - t_0} \int_{t_0}^{t_f} \text{dist}(\mathbf{x}_i(\tau), \mathbf{x}_j(\tau)) d\tau, \quad (2)$$

where $\text{dist}(\cdot, \cdot)$ is the Euclidean distance. Then, a similarity measure for the edge weights $w_{ij} = w(r_{ij})$ must be chosen as a function of the time-averaged distance. The only restriction on this functional dependence is that the weight should be a monotonically decreasing function of the distance. This choice of similarity function controls the graph edge weight distribution, and therefore can impact the clustering results [18].

Hadjighasem et al. [15] use a similarity function that is the inverse of the time-averaged distance, $w_{ij} = l_x/r_{ij}$ (a constant l_x is included here to make the weight dimensionless and does not impact the results). A more widely used function for graph partitioning is the Gaussian function [16–18]

$$w_{ij} = \exp\left(-r_{ij}^2/2\sigma^2\right), \quad (3)$$

which we will use. The similarity radius $\sigma > 0$ is a method free-parameter that controls the spatial width of the connected neighborhoods [16]. Note that $w_{ii} = 1$ and $w_{ij} \rightarrow 0$ for $r_{ij} \gg \sigma$. The choice of the Gaussian measure has a number of advantages over the l_x/r_{ij} choice: first, it is bounded, with $0 \leq w_{ij} \leq 1$; second, no offset value is set in the diagonal entries [15]; and third, the sparsification step can be skipped (or at least has a negligible impact on the results, see Supplementary Materials A), as (3) goes to zero faster than the measure in [15] for large r_{ij} . The similarity radius σ in (3) determines the rate of decay of w_{ij} to zero as r_{ij} increases, which is analogous to the graph sparsification in [15] where $w_{ij} = 0$ for r_{ij} above a defined threshold. A direct comparison between these two similarity functions and the impact of the σ -choice are presented in Sections 3.1 and 3.2, respectively.

The similarity matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ stores the w_{ij} values, and the diagonal degree matrix \mathbf{D} is computed such that $d_{ii} = \sum_{j=1}^N w_{ij}$. The generalized eigenvectors \mathbf{q} of the unnormalized graph Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W}$ are computed from the generalized eigenproblem

$$\mathbf{L}\mathbf{q} = \lambda\mathbf{D}\mathbf{q}. \quad (4)$$

The normalized vectors $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_N$, corresponding to the generalized eigenvalues $0 \leq \lambda_1 \leq \dots \leq \lambda_N$, differentiate properties in the graph and facilitate the clustering process [18]. While all of the eigenvectors provide information, the dominant eigenvectors, associated with the smallest eigenvalues, reveal the most dynamically relevant characteristics. Each trajectory is characterized by a value within each eigenvector, and these are the values ultimately used to cluster trajectories.

The eigenvector matrix $\mathbf{Q} \in \mathbb{R}^{N \times M}$ whose columns are the dominant eigenvectors $\mathbf{q}_1, \dots, \mathbf{q}_M$, with $M \ll N$, stores the characteristics used to cluster the trajectories. The number of retained eigenvectors M is specified depending on the system. Let $\mathbf{y}_i \in \mathbb{R}^M$ be the characterization vector corresponding to the i -th row of \mathbf{Q} , which contains the condensed differentiating information of trajectory i . The eigenanalysis provides a suitable low dimensional representation of the data set. The characterization vectors $\{\mathbf{y}_i\}_{1 \leq i \leq N}$ are then partitioned into K clusters. The relationship between the number of clusters K and the number of dominant eigenvectors M used for clustering depends on characteristics of the system, and is specified for each of the case studies in the following sections.

We cluster the characterization vectors using a fuzzy c -means algorithm [5,19], instead of the conventional k -means often performed at this step [15,18]. The c -means algorithm assigns to each trajectory i probabilities $p_{i,k} \in [0, 1]$ of being a member of cluster k , with $1 \leq k \leq K$, such that $\sum_{k=1}^K p_{i,k} = 1$. This step requires the prescription of a fuzziness parameter $m > 1$ that controls how “tight” the clustering membership probabilities are. As $m \rightarrow 1$, the clustering result approaches the k -means result where $p_{i,k} \in \{0, 1\}$. For greater m , the “looser” distribution in membership probabilities allows the identification of particles that present intermediate behaviors between clusters, as opposed to always assigning a trajectory as member of a single cluster. The cluster centers are initialized based on the dominant eigenvectors before starting the c -means iterative process detailed by Froyland and Padberg-Gehle [5]. Finally, Hadjighasem et al. [15] use an eigengap heuristic to determine M and K that relies on a large cluster differentiation to work properly, and in many practical cases an eigengap

may be less pronounced or nonexistent [18]. We leave the number of dominant eigenvectors and clusters as method free-parameters and study the related uncertainty associated to the clustering results for each K -choice.

2.2. Uncertainty Quantification for Multiple Realizations

In Sections 4 and 5, we will be interested not in applying the spectral clustering algorithm to a single realization of the flow, but in collecting clustering results from different realizations and combining them to get statistical information about the variability of the clusters with the method free-parameters σ , m and K , or with velocity model parameters. We therefore need a method to combine clustering results from different realizations. Regardless of the differences between realizations, we initialize the N particles on the same grid, so that trajectory i is uniquely labeled across realizations, based on $\mathbf{x}_i(t_0)$. Each realization $I \in \{1, \dots, R\}$ generates a full set of membership probabilities $p_{i,k}^{(I)}$, with $i \in \{1, \dots, N\}$ and $k \in \{1, \dots, K\}$. The cluster labels for different realizations are matched based on the similarity between cluster centers.

To quantify the uncertainty of a trajectory i being a member of cluster k , the probabilities $p_{i,k}^{(I)}$ are used to compute the mean membership probabilities over all R realizations

$$\bar{p}_{i,k} = \frac{1}{R} \left(\sum_{I=1}^R p_{i,k}^{(I)} \right), \quad (5)$$

and the corresponding sample standard deviations

$$S_{i,k} = \sqrt{\frac{1}{R-1} \sum_{I=1}^R \left(p_{i,k}^{(I)} - \bar{p}_{i,k} \right)^2}. \quad (6)$$

Both the mean and the standard deviation of the realizations are bounded values, with $0 \leq \bar{p}_{i,k} \leq 1$ and $0 \leq S_{i,k} \leq 0.5$. We perform this calculation for each trajectory, for each cluster.

3. The Bickley Jet System and Sensitivity to Method Free-Parameters

We analyze the analytic, quasi-periodic Bickley jet system to evaluate the spectral clustering method free-parameter and velocity field model sensitivity. This system is an idealized model of a meandering zonal jet under geostrophic balance and has been extensively used to illustrate coherent structures in fluid flows [20,37,38,42]. The model features a sheared zonal flow on which a superposition of Rossby-like waves propagate. The streamfunction ψ prescribing the two-dimensional velocity field $\mathbf{u} = -\partial_y \psi \mathbf{e}_x + \partial_x \psi \mathbf{e}_y$ with a superposition of three waves is

$$\psi(t, x, y) = -UL \tanh\left(\frac{y}{L}\right) + UL \operatorname{sech}^2\left(\frac{y}{L}\right) \sum_{n=1}^3 A_n \cos[k_n(x - c_n t + \phi_n)], \quad (7)$$

where U and L are a characteristic speed and length, respectively. The domain is periodic in the x -direction, with periodicity $l_x = 2\pi R_e \cos(60^\circ)$, where $R_e = 6371$ km is Earth's radius. The rectangular domain corresponds to $x/l_x \in [0, 1]$, and we limit our analysis to $y/l_x \in [-0.15, 0.15]$. The Rossby-like waves correspond to the three longest wave modes in the periodic domain, with amplitudes A_n , wave numbers $k_n = 2\pi n/l_x$, phase speeds c_n , and phases ϕ_n , for $n = 1, 2, 3$. To model the self-consistent state obtained by Rypina et al. [37] for modes 2 and 3 on the periodic domain, we fix $U = 62.74 \text{ ms}^{-1}$, $L = 1767$ km, $c_2/U = 0.2051$, and $c_3/U = 0.4615$. The mode-1 wave has speed $c_1/U = 0.1446$ chosen based on the golden ratio to break periodicity [15,37]. Provided that $A_1 \ll \min(A_2, A_3)$, the mode-1 wave acts as a small perturbation to the system's periodicity. In this section, we fix the mode amplitudes to $A_1 = 0.0075$, $A_2 = 0.15$, and $A_3 = 0.30$, values used by Hadjighasem et al. [15], and all three waves are in phase: $\phi_1 = \phi_2 = \phi_3 = 0$. Note, however,

that while the system dynamics depend on the values of A_n and ϕ_n [37], there is no physical basis for the stated choice of amplitudes and phases, and the impact of varying these parameters will be explored in Section 4.

The present section discusses the application of the spectral clustering algorithm to the Bickley jet system, and studies the sensitivity of the results to user-defined method free-parameters. We present the impact of the choice of similarity measure in Section 3.1. Then, study the sensitivity to the method free-parameters: similarity radius σ , in Section 3.2, tightness of the cluster memberships m , in Section 3.3, and number of clusters K , in Section 3.4.

3.1. Gaussian Similarity Measure

To cluster the Bickley jet system according to the method described in Section 2.1, particles are initialized in a uniform grid of 400 by 120 positions uniformly covering the domain, and advected from $t_0 = 0$ to $t_f = 40$ days, matching [15]. The distance function in (2) takes into consideration the x -periodicity of the domain, and the $M = 6$ dominant eigenvectors of (4) are used to partition the domain into $K = 7$ clusters, to account for 6 materially coherent vortices, which are the coherent clusters, and an incoherent cluster, the chaotic sea [15]. The method uses a fuzziness parameter $m = 2$.

The membership probabilities for the clusters identifying the 6 materially coherent vortices at time t_0 are presented in Figure 1. Figure 1a presents our results, obtained using a Gaussian similarity function (3), with similarity radius $\sigma/l_x = 0.020$, and no sparsification of the similarity matrix. The membership probabilities, plotted in different colors, highlight the 6 coherent vortices. We assign labels to the vortices, from left to right, and vortex 6 appears split at t_0 due to the x -periodicity. The membership probabilities of being part of the chaotic sea, the seventh cluster, are complementary to the ones plotted and are omitted throughout. Note that the use of a soft membership assigns to particles located at the periphery of the vortices lower probabilities of being a member, which relates to lower similarity in the dynamics (some of them may, for instance, be trapped inside the vortices for just a fraction of the time window of analysis, then merge with the chaotic sea, or vice-versa).

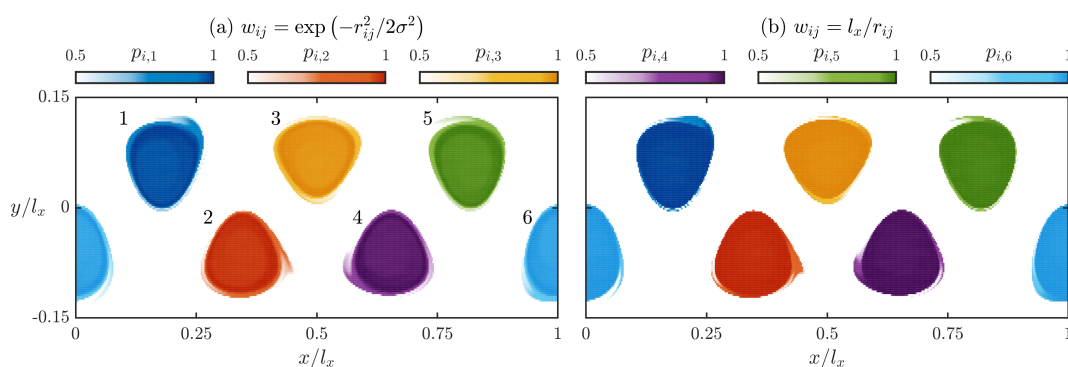


Figure 1. Spectral clustering membership probabilities for clusters $k \in \{1, \dots, 6\}$ identifying materially coherent vortices, with fuzziness $m = 2$ and number of clusters $K = 7$ (incoherent cluster probabilities are omitted). The similarity functions w_{ij} used are (a) the proposed Gaussian similarity measure (3) with $\sigma/l_x = 0.020$, and (b) the inverse measure l_x/r_{ij} from [15] sparsifying values less than $1/0.075$. The six color maps presented here are also used in Figures 2, 4a,b, and 7.

The results for the Gaussian similarity measure in Figure 1a are similar to those obtained with the l_x/r_{ij} similarity measure used by Hadjighasem et al. [15], presented in Figure 1b. For the latter, matrix entries $w_{ij} = 0$ for $r_{ij}/l_x > 0.075$, and the offset diagonal value is chosen as 100 times the largest matrix entry. The clustering results are particularly sensitive to the sparsification threshold, which relates to the choice of σ for the similarity measure (3). Both similarity functions yield similar results for the selected parameter values, with smoother transitions in membership probabilities (from 1 to 0) for the Gaussian measure.

When using the l_x/r_{ij} similarity measure, the degree of sparsification is an additional parameter for the clustering method that can impact the results. This parameter can be eliminated for the Gaussian similarity measure as no sparsification was necessary for the result in Figure 1a, but it is worthwhile to sparsify the matrix to reduce computational costs and storage. We demonstrate in the Supplementary Materials A that sparsifying entries of \mathbf{W} that satisfy $w_{ij} < \exp(-4^2/2) \approx 3 \cdot 10^{-4}$, corresponding to $r_{ij} > 4\sigma$, has negligible impact on the clustering membership probability results, and hence we sparsify according to this rule hereafter. Note that this result is valid for the Bickley jet, but the impact of sparsification may vary for an arbitrary flow.

3.2. Similarity Radius

We highlighted in the previous section how the similarity radius σ is closely related to the sparsification threshold used by Hadjighasem et al. [15]. While the sparsification threshold selection was mentioned in [15], the clustering results are highly sensitive to this parameter [36], and we now address this sensitivity.

While there may be some intuition on the size of the structures of interest, this is not always helpful in prescribing σ or in understanding how the σ -choice impacts the resulting clusters. Hadjighasem et al. [15] choose their threshold by defining which values of \mathbf{W} to keep based on a fixed percent sparsification of the matrix. However, for a fixed percent sparsification, the graph connections that are retained are ultimately a function of the number of particles and their distribution in the initial grid.

Based of this relationship between the sparsification level in [15] and the parameter σ in the Gaussian similarity function (3), we vary σ to demonstrating the clustering sensitivity to changes in sparsification. First, we define an interval bounding all relevant σ -values to be tested. For σ/l_x distributed on the interval $[0.005, 0.040]$ with steps of 0.001 (36 cases), the method is applied with $m = 2$ to identify $K = 7$ clusters. Figure 2 presents the membership probabilities for each of the six coherent clusters for $\sigma/l_x = 0.005, 0.015, 0.030$ and 0.040 . Compared to the results for $\sigma/l_x = 0.020$ presented in Figure 1a, smaller values of σ (Figure 2a,b) tend to shrink the coherent clusters to the vortex cores only, while larger choices of σ (Figure 2c,d) assign higher membership probabilities to filaments that correspond to particles that do not belong to the coherent vortex from the start, but have a long residence time on the perimeter of the vortex.




Figure 2. Membership probabilities for vortices $k \in \{1, \dots, 6\}$, for different choices of the parameter σ , plotted at $t = t_0$. Values correspond to (a) $\sigma/l_x = 0.005$, (b) 0.015 , (c) 0.030 , and (d) 0.040 . Corresponding colorbars presented in Figure 1. (See Supplementary Materials for a video of the time evolution of the clustered trajectories.)

To determine how the membership probabilities for each trajectory depend on σ , we use the information from the different realizations to compute the membership probability means $\bar{p}_{i,k}$ and sample standard deviations $S_{i,k}$ for each trajectory and cluster, as described in Section 2.2. These statistical measures are presented in Figure 3. In Figure 3a, we present $\bar{p}_{i,1}$ for vortex $k = 1$, and in Figure 3b the corresponding standard deviation $S_{i,1}$. Figure 3c,d condense the information for all vortices $k \in \{1, \dots, 6\}$. The superimposed mean $\bar{p}_{i,\hat{k}}$ in Figure 3c and the superimposed standard deviation $S_{i,\hat{k}}$ in Figure 3d are such that, for each trajectory i , \hat{k} is the cluster that maximizes the mean membership probability, hence $\hat{k} = \operatorname{argmax}_k \bar{p}_{i,k}$.

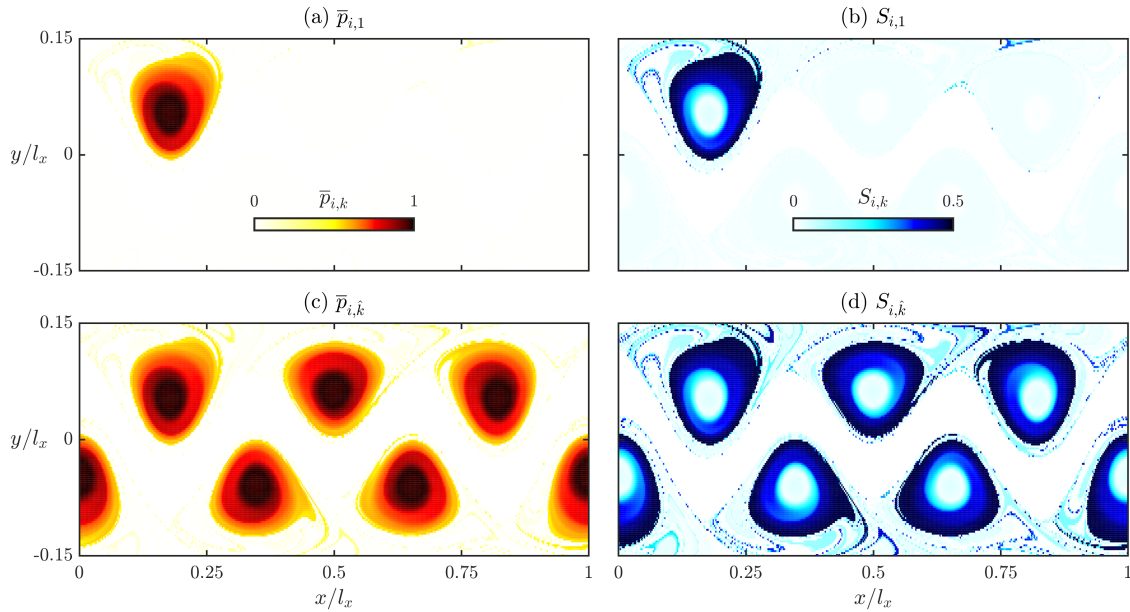


Figure 3. Clustering statistics varying σ , plotted at $t = t_0$, for σ uniformly distributed in $[0.005, 0.040]$. Membership probability (a) mean for $k = 1$, (b) standard deviation for $k = 1$, (c) superimposed mean, and (d) superimposed standard deviation for $k \in \{1, \dots, 6\}$. The two color maps presented here are also used in Figures 4c,d, 5, 6, 8, and 9a.

Figure 3a,c demonstrate that the vortex cores have the greatest mean membership, which relates to the fact that those particles are identified with high membership probabilities in all realizations. Particles further away from the cores have a lower mean, as a result of lower probabilities of being part of the respective vortices, in particular for low σ values. We also notice sharp drops in the probability after a certain vortex size. The uncertainty of particles being part of a vortex is highlighted by the standard deviations in Figure 3b,d, where we again see negligible standard deviation (low uncertainty) on the membership probabilities for the cores. As a result that there are no realizations that identify the central jet and some portions of the chaotic sea as part of the clusters, there is also a low standard deviation for those trajectories. The highest uncertainty is obtained for particles between the core and the chaotic sea, and some filaments are also highlighted with higher standard deviation. Those filaments correspond to particles consistently identified as part of a vortex, with high membership probabilities, for σ greater than a threshold, but not for lower values of σ .

3.3. Fuzziness Parameter

Next, we consider the sensitivity of the choice of the fuzziness parameter m , introduced by the c -means step that assigns cluster membership probabilities. In this analysis, $\sigma/l_x = 0.020$ while m is varied on the interval $[1.05, 3.00]$ with steps of 0.05. The results are presented in Figure 4. The membership probabilities for the extreme values ($m = 1.05$ and 3.00) are illustrated in Figure 4a,b, and the superimposed mean and standard deviation in Figure 4c,d.




Figure 4. Membership probabilities for $k \in \{1, \dots, 6\}$, with (a) $m = 1.05$, and (b) $m = 3.00$. Clustering statistics plotted at $t = t_0$ for m uniformly distributed in $[1.05, 3.00]$. Membership probability (c) superimposed mean and (d) superimposed standard deviation for $k \in \{1, \dots, 6\}$. Corresponding colorbars presented in Figures 1 and 3.

While $m = 1.05$ (in Figure 4a) yields a bimodal probability distribution, with 99% of the $p_{i,k} > 0.5$ cases also being greater than 0.95 (so tending to the corresponding k -means result), the use of $m = 3.00$ (Figure 4b) produces smoother probability transitions from the vortex core to the perimeter, as well as overall lower membership probabilities, with only 10% of the $p_{i,k} > 0.5$ cases greater than 0.95. The superimposed mean and standard deviations reveal a similar trend to the one observed for the dependence on σ , with two major differences: (i) while varying m introduces uncertainty in the membership probability of trajectories starting at the vortex perimeters, the m -choice does not identify filaments as part of the vortex for the current $\sigma/l_x = 0.020$ value, and (ii) the magnitude of the standard deviations related to m are about half of the ones associated to σ (see Figure 4d in comparison to Figure 3d).

A direct comparison between the σ - and m -sensitivity is presented in Figure 5, where the mean and standard deviation profiles are plotted along the cross-section $x/l_x = 0.5$, passing through the center of vortex 3 at $t = t_0$, for $y/l_x = [-0.02, 0.15]$. These profiles are interpolated from the membership probability statistics for the particle positions in Figures 3c,d and 4c,d. Varying σ generates higher mean membership probabilities in the vortex core and a more gradual drop in the mean probability from the core to the perimeters (Figure 5a), which is associated to wider peaks of high standard deviation (Figure 5b). While the core has slightly higher standard deviations for the m -variation study, the maximum standard deviations, which occur at the vortex perimeters, are only half of the ones associated to σ . There is, however, a remarkable consistency on the vertical coordinates where sharp drops in $\bar{p}_{i,3}$ and $S_{i,3}$ occur, for both σ and m variations: $y/l_x = 0$ (bottom) and 0.13 (top). These results reflect the fact that the method applied to the Bickley jet system is more sensitive to the choice of σ than m .

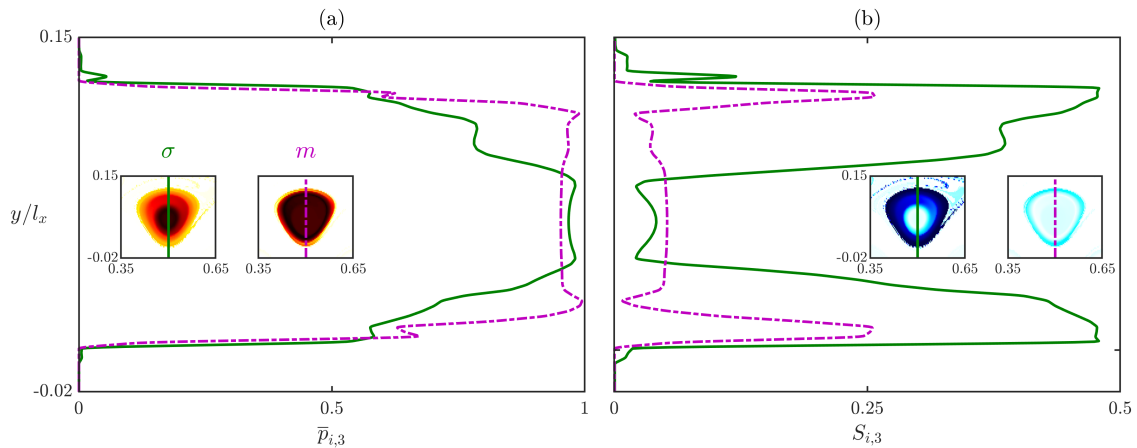


Figure 5. Membership probability profiles for (a) mean $\bar{p}_{i,3}$ and (b) standard deviation $S_{i,3}$ corresponding to vortex $k = 3$, for the variable- σ (solid green line) and variable- m (dash-dotted magenta line) studies. Insets present the clustering statistics for vortex 3, from Figures 3c,d and 4c,d. The profiles correspond to the cross-section $x/l_x = 0.5$ and are interpolated from the membership probabilities at $t = t_0$, presented in the insets. Corresponding colorbars presented in Figure 3.

3.4. Number of Clusters

Finally, the number of clusters in the system is not necessarily known beforehand, and here we address how the clustering results statistically vary for different choices of K . As presented in [15], for a suitable sparsification level of the similarity matrix, the eigengap heuristic can be used to infer the use of $M = 6$ dominant eigenvectors to choose $K = 7$ for the Bickley jet, and thus identify the 6 materially coherent vortices, plus the chaotic sea. It is known, however, that the existence of an eigengap is not guaranteed for any system, and depends on, among other properties, the connectivity of the similarity graph [18]. The number of clusters for systems not as distinctly separated as the Bickley jet will, therefore, be more challenging to determine. Without any prior knowledge about the number of clusters in the system, one could consider clustering based on different numbers of dominant eigenvectors of (4) into different K , which might result in merging clusters, splitting clusters, missing clusters, or even identifying new ones. As we vary K , we fix the relationship between M and K to $K = M + 1$ to always include the chaotic sea as an independent cluster [15].

As a result that the free-parameter K cannot be varied continuously, and because varying K while fixing σ and m would only mean changing the number of eigenvectors to use in the c -means step, we adopt a different strategy to quantify the K -uncertainty. For each choice of K , we perform statistics using realizations in which σ and m are varied. The $(\sigma/l_x, m)$ pairs are sampled from uniform distributions over $[0.005, 0.040] \times [1.05, 3.00]$, corresponding to the same intervals in Sections 3.2 and 3.3. We consider 100 samples from uniform distributions, and compute mean and sample standard deviations for $K = 6, 7$, and 8. One should now expect missing or merging vortices for $K < 7$, while splitting vortices or identifying new structures for $K > 7$. For the ensemble statistics, if vortex k is not identified in realization I , we set $p_{i,k}^{(I)} = 0$ for all i . For $K = 8$, a new coherent cluster corresponding to the jet is consistently identified, in all realizations. The jet is therefore considered a seventh coherent cluster.

The superimposed means and standard deviations for $K = 6, 7$, and 8 are presented in Figure 6. For $K = 6$, Figure 6a,b reveals that different $(\sigma/l_x, m)$ free-parameter choices can result in different vortices not being identified. Notably, vortices 3, 5, and 6 are missed in some of the realizations, which reduces their mean probability and increases uncertainty for those clusters. Figure 6c,d, for $K = 7$, presents six vortices identified with similar probability distributions to the ones obtained by varying σ alone in Section 3.1 (in Figure 3c,d). Such similarity between these cases highlights the dominance of the σ -sensitivity over the choice of m , and the fact that these two parameters do not compound each other. Finally, Figure 6e,f present the results for $K = 8$, highlighting the consistent identification of the

jet as a coherent cluster in the system. The jet cluster has a higher sensitivity to σ and m , with higher standard deviations than the vortex cores in Figure 6b.

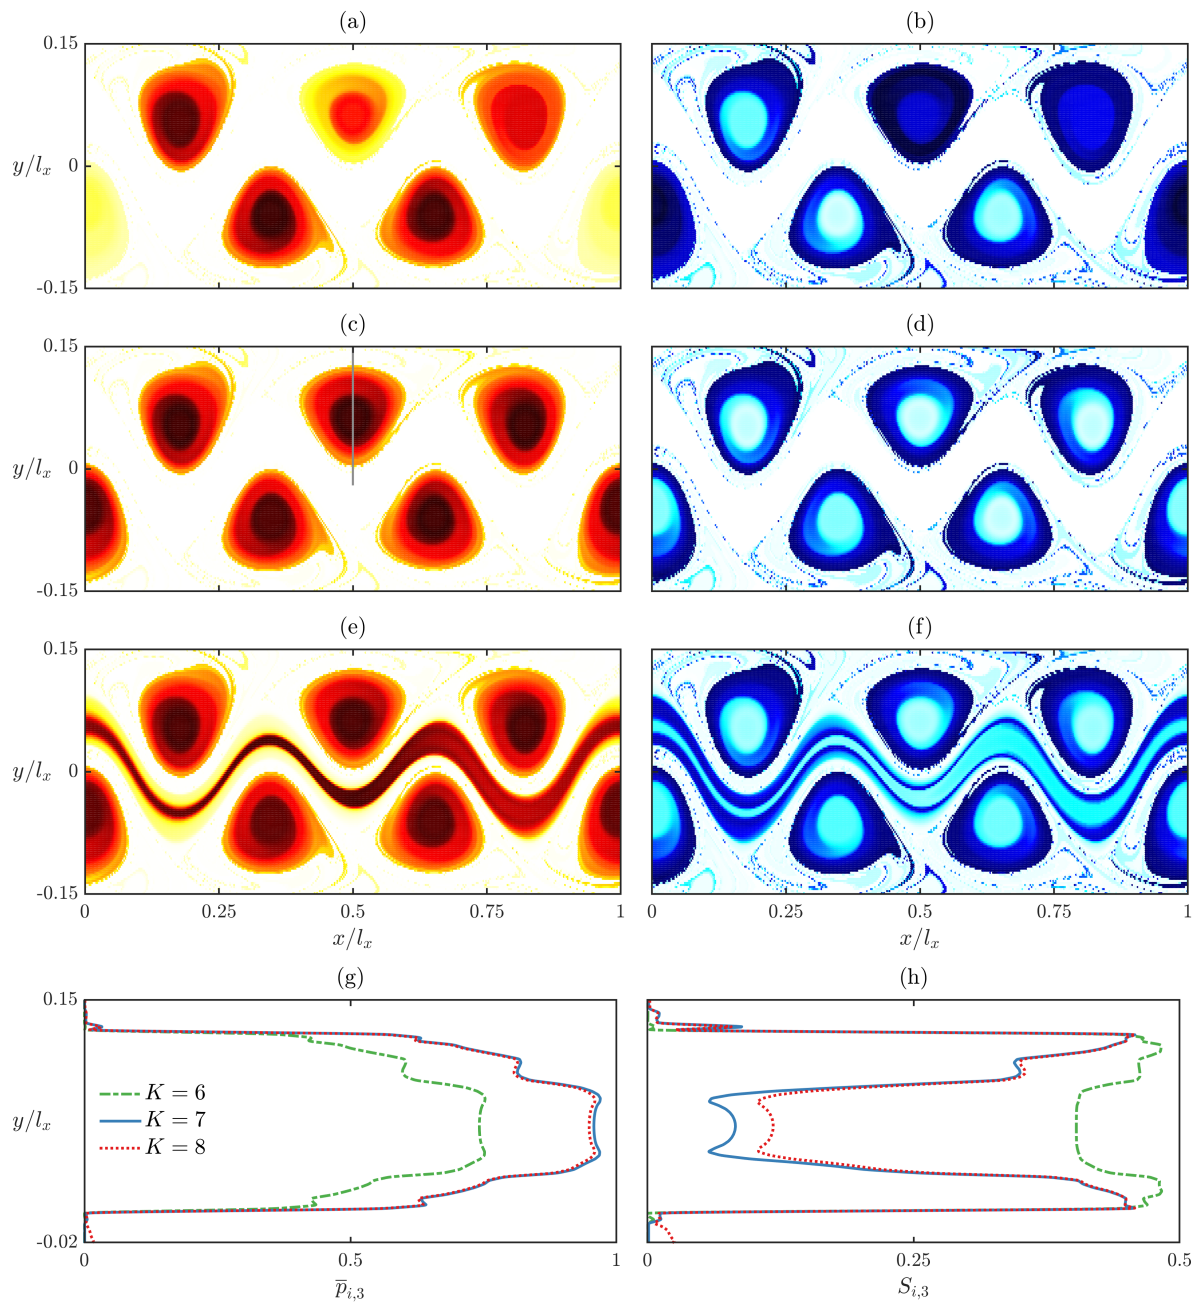


Figure 6. Clustering statistics varying σ and m , for different numbers of clusters K , with $(\sigma/l_x, m)$ sampled from a uniform distribution over $[0.005, 0.040] \times [1.05, 3.00]$. Each row for (a–f) presents the superimposed mean $\bar{p}_{i,\hat{k}}$ (left) and the corresponding superimposed standard deviation $S_{i,\hat{k}}$ (right), for (a,b) $K = 6$, (c,d) $K = 7$, and (e,f) $K = 8$. Membership probability profiles corresponding to the cross-section $x/l_x = 0.5$ (solid gray line in (c)) are interpolated at $t = t_0$ and presented for (g) mean $\bar{p}_{i,3}$ and (h) standard deviation $S_{i,3}$ corresponding to vortex $k = 3$, for variable K . Corresponding colorbars presented in Figure 3.

Figure 6g,h present a direct comparison between mean and standard deviation profiles along the cross-section $x/l_x = 0.5$, passing through the center of vortex 3 at $t = t_0$, for the different K values. For $K = 6$, the peak mean probabilities of vortex 3 are less than 0.75, and there is high standard

deviations throughout. The mean membership probability at the core goes up to over 0.95 for $K = 7$, corresponding to a drop in the standard deviation. As another cluster is added for $K = 8$ and the jet is identified, Figure 6g,h reveals a modest drop in the mean values compared to $K = 7$, associated to an increase in the standard deviation for the vortex core from 0.08 to 0.12.

While these ensemble statistics could be used as a basis for setting the method free-parameters, a thorough investigation on this is beyond the focus of the parameter sensitivity study presented in this paper. Based on the results for $K = 6, 7$, and 8, it would be reasonable to state that $K = 7$ could be chosen over the other options, as it is the case leading to the smallest space-averaged uncertainty (see impact of choosing $K = 8$ compared to $K = 7$ in Figure 6h). At the same time, however, the jet is a structure that differentiates itself from the rest of the flow, by behaving in a distinct and coherent way compared to the remaining particles in the chaotic sea. It could be argued that the choice of $K = 8$ for this system is an equally possible way of clustering the system, and provides extra information about the jet cluster, at the price of increasing the result sensitivity to other method free-parameters. Further work is necessary to automate this selection for a general system.

4. Ensemble Realizations and Uncertainty to Model Parameters and System Dynamics

In the previous section, a single set of model parameters of the Bickley jet system, with fixed physical parameters, was used to demonstrate how the choice of the method free-parameters impact the clustering results. Here, we focus on model parameters that influence the dynamics of the system by changing the velocity field and the resulting trajectories. Multiple realizations of the system are used to determine the clustering sensitivity to model parameters, allowing for a characterization of the robustness of the identified clusters. Section 4.1 explains how system parameters are randomized to generate multiple dynamically different realizations. Section 4.2 presents the statistics over the described realizations, leading to an uncertainty quantification of the clustering results to model parameters.

4.1. Perturbing the Bickley Jet Dynamics

While system parameters such as U , L , c_2 , and c_3 are set by physical arguments (as discussed in Section 3), the wave amplitudes A_n and phases ϕ_n are not, despite exerting a major influence on the system dynamics [37]. We use these values as unknown model parameters to introduce variability in the system dynamics. The amplitudes and phases are sampled from normal distributions centered around the values used in Section 3 (and [15]). The realization presented in Section 3 is hereafter referred to as the central realization, and corresponds to amplitudes $A_n = \bar{A}_n$, with $\bar{A}_1 = 0.0075$, $\bar{A}_2 = 0.15$, and $\bar{A}_3 = 0.30$, and phases $\phi_n = 0$. The parameters for each realization I in this section are generated by

$$A_n^{(I)} / \bar{A}_n \sim \mathcal{N} \left(1, \left(\frac{1}{2} \right)^2 \right) \quad \text{and} \quad \phi_n^{(I)} / l_x \sim \mathcal{N} \left(0, \left(\frac{1}{24} \right)^2 \right), \quad (8)$$

where $\mathcal{N}(\mu, \Sigma^2)$ denotes a normal distribution of mean μ and standard deviation Σ . The standard deviations for $A_n^{(I)}$ are scaled by the corresponding mean values, while the standard deviation for $\phi_n^{(I)}$ is chosen small enough so that the vortex centers for each realization are likely to be inside of the area covered by the vortices in the central realization.

For each of the studies presented, a total of $R = 1000$ realizations are generated from the distributions in (8), and the spectral clustering algorithm described in Section 2 is applied with method parameters fixed to $\sigma/l_x = 0.020$, $m = 2$ and $K = 7$. While it is possible that individual realizations may require a different method free-parameter selection, our goal in presenting this section is to separate the effects of method free-parameters from those of model parameters. By sampling model parameters together with method free-parameters, the cluster uncertainty resulting from the model would be obfuscated. We thus fix the method free-parameters based on the central realization only, while considering multiple realizations with varying model parameters.

Figure 7 presents the membership probabilities for four realizations with variable A_n and ϕ_n . While the realizations do modify the position, shape, and dynamics of the vortices (see videos in Supplementary Materials), their presence and number, as well as the presence of the shear jet, mostly remain unchanged. Figure 7a,b presents how variable the initial cluster sizes and shapes can be, as well as the effect of the wave phases on the nonuniform spacing between the identified vortices at t_0 . While for most realizations all six vortices are identified, there are cases where some of the expected vortices are not identified. Figure 7c presents a case in which the jet is identified and one of the vortices is missed. For that case, vortex 5 gets identified as part of the chaotic sea (not plotted), while a highly asymmetric jet is identified as another coherent cluster in the system, and trajectories are assigned membership probabilities $p_{i,jet}$. Figure 7d presents a case for which a more symmetric jet is identified as a cluster and two of the vortices (2 and 4) are merged into a single cluster. Other realizations, not illustrated here, result in cases where only a few (or even none) of the vortices are identified. For those realizations, there is no clear Eulerian signature of the six vortices. In what follows, only clusters that identify one and only one vortex are considered for statistical purposes.

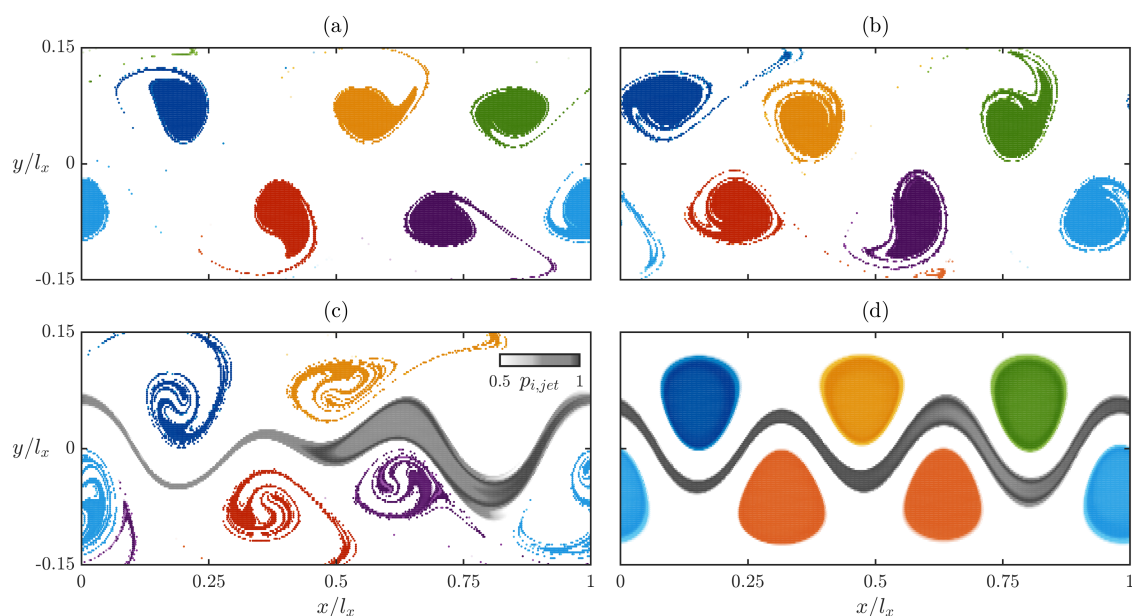


Figure 7. Examples of membership probabilities for the six identified clusters, plotted at $t = t_0$, for different model parameters $\{A_n\}$ and $\{\phi_n\}$ sampled from normal distributions. Cases correspond to parameters $(A_1, A_2, A_3, \phi_1/l_x, \phi_2/l_x, \phi_3/l_x)$ equal to (a) (0.0087, 0.19, 0.20, 0.01, 0.06, -0.03), (b) (0.0069, 0.26, 0.25, 0.00, -0.08 , 0.09), (c) (0.0102, 0.35, 0.25, 0.01, -0.01 , 0.01), where the jet is identified and one of the vortices missed, and (d) (0.0077, 0.11, 0.32, 0.00, -0.01 , 0.02), where the jet (grayscale) is identified and two vortices are merged. Corresponding colorbars for the vortices presented in Figure 1. (See Supplementary Materials for a video of the time evolution of the clustered trajectories.)

4.2. Uncertainty Quantification of Ensemble Simulations

Clustering results for all the realizations are analyzed to measure their uncertainty statistics. For each one of the vortices, the mean and standard deviation over the ensemble of parameter value sets of the Bickley jet are computed according to (5) and (6). Three independent studies, with $R = 1000$ realizations each, are performed to first isolate the effect of the amplitude variation, with variable A_n and $\phi_n = 0$ (Figure 8a,b), then the phase variation, with variable ϕ_n and $A_n = \bar{A}_n$ (Figure 8c,d), and finally the combined effect of varying both A_n and ϕ_n at the same time (Figure 8e,f).

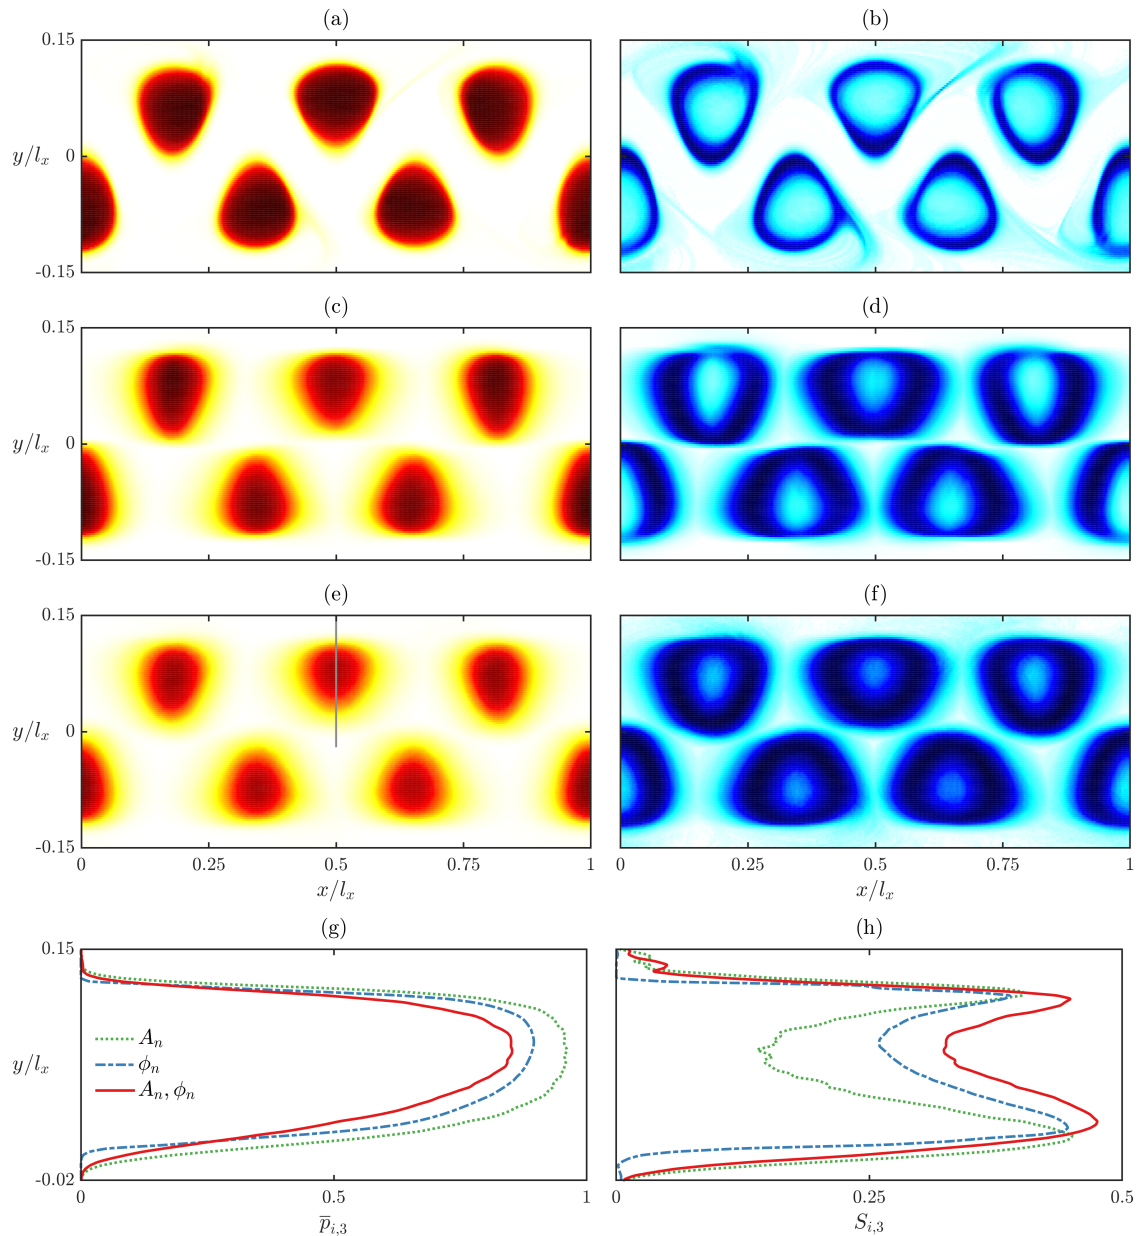


Figure 8. Clustering statistics, plotted at $t = t_0$, for three ensembles of $R = 1000$ realizations of the Bickley jet system with (a,b) variable amplitudes A_n and constant phases, (c,d) variable phases ϕ_n and constant amplitudes, and (e,f) variable amplitudes A_n and phases ϕ_n . Coefficient values are sampled from normal distributions around the central realization values. Each row in (a–f) presents the superimposed mean $\bar{p}_{i,\hat{k}}$ (left) and the corresponding superimposed standard deviation $S_{i,\hat{k}}$ (right). Membership probability profiles corresponding to the cross-section $x/l_x = 0.5$ (solid gray line in (e)) are interpolated at $t = t_0$ and presented for (g) mean $\bar{p}_{i,3}$ and (h) standard deviation $S_{i,3}$ corresponding to vortex $k = 3$, for the three studies. Corresponding colorbars presented in Figure 3.

Figure 8a,c,e presents the superimposed membership probability means $\bar{p}_{i,\hat{k}}$ and Figure 8b,d,f the corresponding superimposed sample standard deviations $S_{i,\hat{k}}$, for each of the ensembles. While the amplitude variation mostly affects the vortex size, the phase variation introduces uncertainty on the horizontal position of the identified vortices, and the combined variation leads to an increase in uncertainty around the vortex core positions. Even with a variety of behaviors observed, as illustrated in Figure 7, averaging the ensemble smooths the vortices, resulting in cluster shapes and positions that are similar to the ones obtained for the central realization (Figure 1a). However, the mean probabilities

in Figure 8a,c,e highlight how introducing uncertainty in the wave amplitudes and phases leads to smaller vortex cores with high membership probabilities. The mean probabilities for the combined A_n and ϕ_n variation study now peak at 0.91 rather than 1.00. The membership probability decay from the vortex cores to the perimeters is also more gradual than for the central realization, and the averaging clears out previously identified filaments that are realization-dependent.

Figure 8b,d,f highlights a more spread out distribution and overall higher magnitude for the superimposed standard deviations. Moreover, higher standard deviations are now observed at the vortex cores. In addition, most of the low uncertainty regions associated to the jet in Figure 3b have higher uncertainty in Figure 8f. With the varying model parameters and dynamics, the vortex positions, shapes and trajectories are more variable. All of these introduce uncertainties that are not observed for the central set of parameter values. Figure 8g,h presents a comparison between means and standard deviations along the cross-section $x/l_x = 0.5$, passing through the center of vortex 3 at $t = t_0$, for the three studies. Higher means and lower standard deviations are obtained for the amplitude-only variation, and the lowest means and highest uncertainty are obtained for the combined variation. The way these parameters contribute to the uncertainties and where mean and standard deviations are maximized, however, do not behave like a simple superposition of contributions. The nonlinear combination is highlighted, for example, by the location of the peaks in standard deviation in Figure 8h that occur closer to the vortex core for the combined variation than for any of the individual variations. The ensemble analysis, therefore, highlights structure sensitivities (and robustness) that are not apparent from the central realization alone.

To illustrate the ramifications of regions of high uncertainty, we demonstrate how different ensemble trajectories are when initialized at high and low uncertainty locations. At $t_0 = 0$, the orange particles in Figure 9a are initially located at the core of cluster 1 and correspond to $\bar{p}_{i,1} = 0.883$ and $S_{i,1} = 0.287$, while the red particles start at the perimeter of cluster 4 and correspond to $\bar{p}_{j,4} = 0.469$ and $S_{j,4} = 0.475$. Figure 9b–d presents snapshots of the particle positions for all 1000 realizations at three different times. For reference, the gray boundaries enclose vortices $k = 1$ and 4 for the central realization (for particles such that $p_{i,k} \geq 0.5$). While particles released in the position of lowest uncertainty (orange) remain concentrated, particles released in the position of highest uncertainty (red) quickly spread throughout the domain. The presence of the jet separating the top and bottom of the domain keeps most orange and red particles from moving to the opposite half. Particle concentration is more pronounced for the case of lower uncertainty, with 76.4% of the particles remaining inside the corresponding boundary after $t = 40$ days, as opposed to only 45.0% for the higher uncertainty case. This corresponds well to the dispersion statistics of the final particle positions, where the dispersion length is the standard deviation of the final position for the different realizations. This length is $0.38 l_x$ for the particles released inside the vortex core (orange) and $2.48 l_x$ for the particles released in the high uncertainty zone (red). Note that the periodicity has been disregarded in this calculation. The six fold increase in the dispersion length is characteristic of particles released in the chaotic sea and the high uncertainty regions. Within the vortex cores, the dispersion is consistently low because the trajectories remain close due to being stuck within the vortex that is highlighted by the clustering results.




Figure 9. Evolution of trajectories from the 1000 realizations with variable A_n and ϕ_n released from a low (orange, top) and high (red, bottom) uncertainty position. (a) Orange particles are initialized at the core of vortex 1, while red particles are initialized at the perimeters of vortex 4. Positions are plotted on top of the superimposed standard deviation $S_{i,k}$ computed using vortices 1 and 4, and plotted at time $t_0 = 0$. The positions for the multiple realizations and the central realization vortex boundaries are presented at (b) $t = 5$, (c) $t = 20$, and (d) $t_f = 40$ days. Transparency is used to highlight high or low concentration of particles. The vortex boundaries from the central realization are plotted in gray and enclose particles with membership probabilities greater than 0.5, for clusters 1 and 4. Corresponding colorbar for (a) presented in Figure 3. (See Supplementary Materials for a video of the time evolution.)

Regions of higher uncertainty, not characterized when considering the central realization only, are revealed by the ensemble analysis, and correspond to flow regions for which particle cluster membership is most unknown. Our analysis shows that the cores of the vortices are robust even if the model parameters are varied, but the narrow filaments identified in the central realizations should be viewed as less robust as demonstrated by the ensemble analysis. Further, while both particles in Figure 9 are initialized and remain inside of their respective clusters for the central realization, the same happens for less than half of the realizations considered, once model parameter uncertainty is incorporated and accounted for.

5. Martha's Vineyard Ensemble Forecast and Surface Drifter Trajectories

Having demonstrated the clustering uncertainty quantification method on an analytic system, we apply the technique to a real ensemble forecast of a nested primitive equation ocean model, with intrinsic model uncertainty. The model is used to forecast the three-dimensional velocity field for the coastal region near Martha's Vineyard, an island located south of Cape Cod, Massachusetts, USA. Trajectories from drifters deployed [36] for the corresponding day are then compared to the cluster behaviors. Section 5.1 presents characteristics of the Martha's Vineyard region, the model used to forecast the velocity field, and the results of the clustering uncertainty quantification analysis applied to the ensemble forecast. Section 5.2 details the drifter experiment and compares the drifter trajectories to the forecast trajectory clustering results and the associated uncertainty.

5.1. Velocity Model Ensemble Forecast and Uncertainty Quantification

The island of Martha's Vineyard, with an area of almost 250 km², is the largest island in New England, and lies 11 km off the coast of Cape Cod. The prevailing currents in the coastal region south of the island are associated with wind-driven coastal divergence, tidal forcing and a mean southward

drift [31]. During the summer months, the region experiences a mean westward surface current that reaches velocities of 15 cm s^{-1} [32]. The drifter deployment experiments targeting predicted coherent structures [36], presented in Section 5.2, took place around the 2.5 km^2 uninhabited island of Nomans Land, south of Martha's Vineyard. The channel between the two islands has a width of approximately 5 km and an average depth of 10 m.

We used the MIT Multidisciplinary Simulation, Estimation, and Assimilation Systems primitive equation (MSEAS-PE) ocean modeling system [39,40] to compute ocean surface velocity forecasts in the Martha's Vineyard coastal region during August 2018. The modeling system provided forecasts of the ocean state variable fields (three-dimensional velocity, temperature, salinity, and sea-surface height) every hour, with a spatial resolution of 200 m. More details about the model forecast initialization, tidal forcing, atmospheric flux forcing, and CTD data assimilation can be found in [25,36]. The deterministic two-way nesting ocean forecast initialized from the estimated ocean state conditions at a particular time is referred to as the central forecast, and ensemble forecasts were initialized using Error Subspace Statistical Estimation procedures [43]. The forecasts within the ensemble were initialized from perturbed initial conditions of all state variables and forced by perturbed tidal forcing, atmospheric forcing fluxes and lateral boundary conditions. These perturbations were created to represent the expected uncertainties in each of these quantities. Finally, parameter uncertainties (bottom drag, mixing coefficients, etc.) were also modeled by perturbing the values of parameters for each forecast.

A total of 71 forecasts were considered for the present study. Fluid particle trajectories confined to the surface were analyzed over a 6 h-time-window, between $t_0 = 16:00$ and $t_f = 22:00$ UTC on 7 August 2018. Such a short timing is critical in search-and-rescue operations, as after six hours the likelihood of rescuing people alive drops significantly [25]. The forecast velocity fields used for trajectory integration were generated the night before the experiment. Synthetic trajectories were computed using the web-based gateway Trajectory Reconstruction and Analysis for Coherent Structure Evaluation [44,45]. At time t_0 , particles are uniformly distributed on a 250-by-250 grid covering the domain $[70.65^\circ \text{ W}, 70.90^\circ \text{ W}] \times [41.15^\circ \text{ N}, 41.35^\circ \text{ N}]$, from which portions corresponding to land are removed. This grid is approximately 21 km by 22 km.

Figure 10a presents the initial particle distribution, superposed with the velocity field for the central forecast. Trajectories are integrated using an adaptive 7th-order Runge–Kutta–Fehlberg method and bicubic spline spatial interpolation, with free-slip boundary conditions applied near land. Particle positions are output every 5 minutes. Figure 10b presents the final positions of the particles for the central forecast. Darker regions correspond to particles collecting, which is mostly observed along the coast. The model velocity field captures the reversal of the tide, as can be observed from the flipping in the average flow direction between t_0 in Figure 10a and t_f in Figure 10b. Trajectories obtained in each of the 71 forecasts are presented in Figure 10c in contrast with the central forecast, for particles initialized at distinct positions, to demonstrate the degree of trajectory variability between forecasts.

A similar study to the one in Section 3 was performed for the central forecast, to determine the spectral clustering sensitivity to method free-parameters. We select the following method free-parameters: similarity radius $\sigma = 1 \text{ km}$, fuzziness $m = 2$, and number of clusters $K = 4$, which are used for all $R = 71$ forecasts. While it is possible to break the domain into fewer clusters, four clusters minimize the space-averaged standard deviation of the results. As a result that the graph is fully connected with this choice of σ , $\lambda_1 = 0$ and the components of \mathbf{q}_1 are constant [18]. Therefore, the clustering (in the absence of a chaotic sea) into $K = 4$ clusters is performed using the information from eigenvectors \mathbf{q}_2 , \mathbf{q}_3 and \mathbf{q}_4 only. The membership probabilities for the central forecast are presented in Figure 11a for trajectories with $p_{i,k} \geq 0.5$. The domain is partitioned into 4 quadrants of similar size, and gaps between clusters correspond to particles with membership probabilities lower than 0.5.

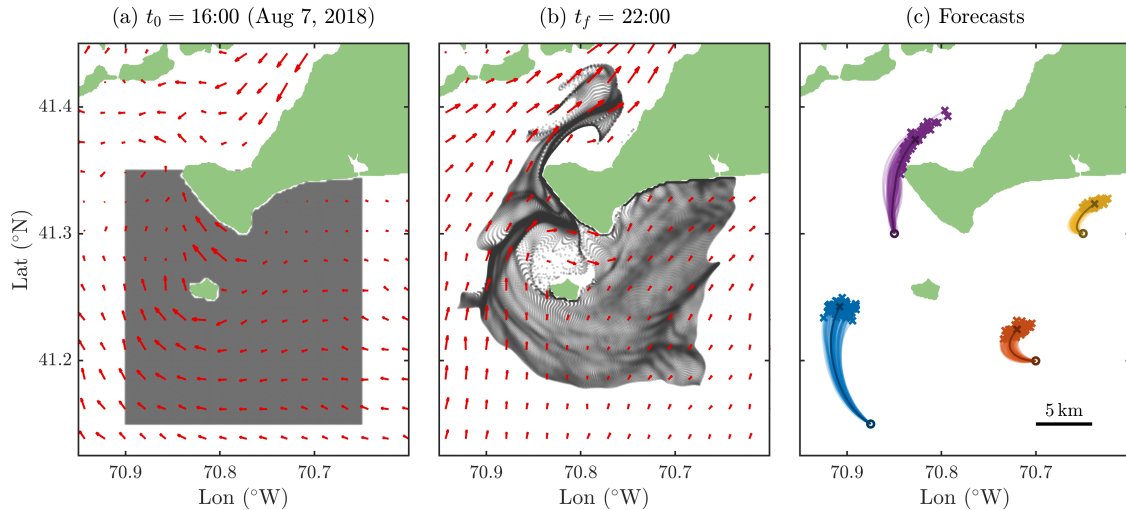


Figure 10. Martha's Vineyard coastal area, central forecast particle distribution, and superimposed forecast model velocity field at (a) the initial time $t_0 = 16:00$, (b) the final time $t_f = 22:00$ UTC. (c) Trajectories obtained in each of the 71 forecasts for 4 different initial positions. Circles represent the initial positions and crosses the final positions, after 6 hours. The darker trajectories correspond to the central forecast.

Different forecasts are used to compute the means $\bar{p}_{i,k}$ and sample standard deviations $S_{i,k}$ for $k \in \{1, \dots, 4\}$. The superimposed mean $\bar{p}_{i,k}$ and standard deviation $S_{i,k}$ are presented in Figure 11b,c. The parameterization used for the model produces only a modest variation in the trajectory outcomes over six hours as demonstrated in Figure 10c. Regions of highest uncertainty for this system correspond to identifying the edges of the clusters accurately, but this level of uncertainty is significantly lower than those observed in the Bickley jet example. The most pronounced uncertainty regions, in Figure 11c, correspond to the boundary between clusters 1 and 4, followed by that between clusters 1 and 2.

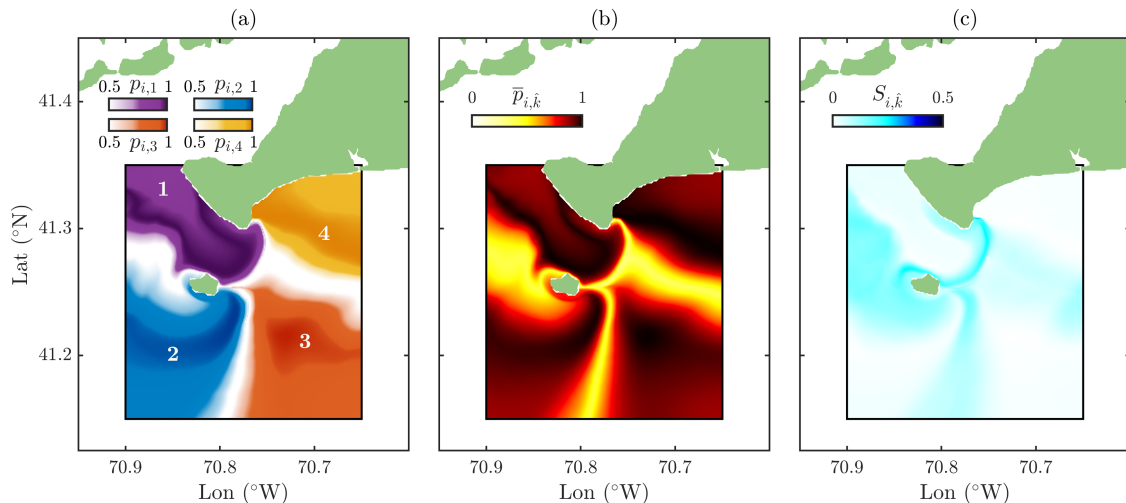


Figure 11. (a) Central forecast membership probabilities at $t = t_0$, for clusters $k \in \{1, \dots, 4\}$. The black box encloses the domain where trajectories are initialized. Clustering statistics for ensemble of $R = 71$ forecasts: superimposed membership probability (b) mean $\bar{p}_{i,k}$ and (c) standard deviation $S_{i,k}$. The color maps presented in (a) are also used in Figure 12.

5.2. Drifter Data and Forecast Cluster Dynamics

The experiment targeting predicted coherent structures consisted of surface drifter releases that took place on 7 August 2018, in the vicinity of Nomans Land [36]. The CODE drifters used in the experiment have technical specifications listed in [25,36]. Drifters of the same design are routinely used

by the U.S. Coast Guard in search-and-rescue operations, as well as in previous field experiments in the coastal region near Martha's Vineyard [31,32]. The drifters were equipped with GPS transmitters that provided positioning fixes every 5 min, with an accuracy of a few meters. An elliptical route around Nomans Land was used for the drifter deployment, employing two WHOI vessels to minimize ship time, so that all drifters were in water by the start of the interval of analysis. Eighteen drifters were deployed in the water around the predicted locations of the targeted coherent structures. The position of the drifters at the starting time of our analysis, $t_0 = 16:00$ UTC, is presented in Figure 12a.

Figure 12a–d presents the mean cluster trajectories for each of the 4 clusters, on which we superpose the drifter positions from [36]. To plot the evolution of $\bar{p}_{i,k}$, the entire domain is first split into 175×175 bins. Then, at each time instance, an average membership probability for each bin, over the forecasts, is calculated. This average is based on particles inside each bin at time t , and weighted by their membership probabilities $p_{i,k}$. The overall system dynamics are similar to what was observed for the central forecast in Figure 10a,b, with the clusters highlighting groups of different behavior. Using the labels from Figure 11a, while cluster 1 (purple) travels a longer distance to the northeast, ending north of Martha's Vineyard, cluster 2 (blue) moves a shorter distance northward, while clusters 3 and 4 (red and yellow) are less dynamic and remain mostly to the east of Nomans Land.

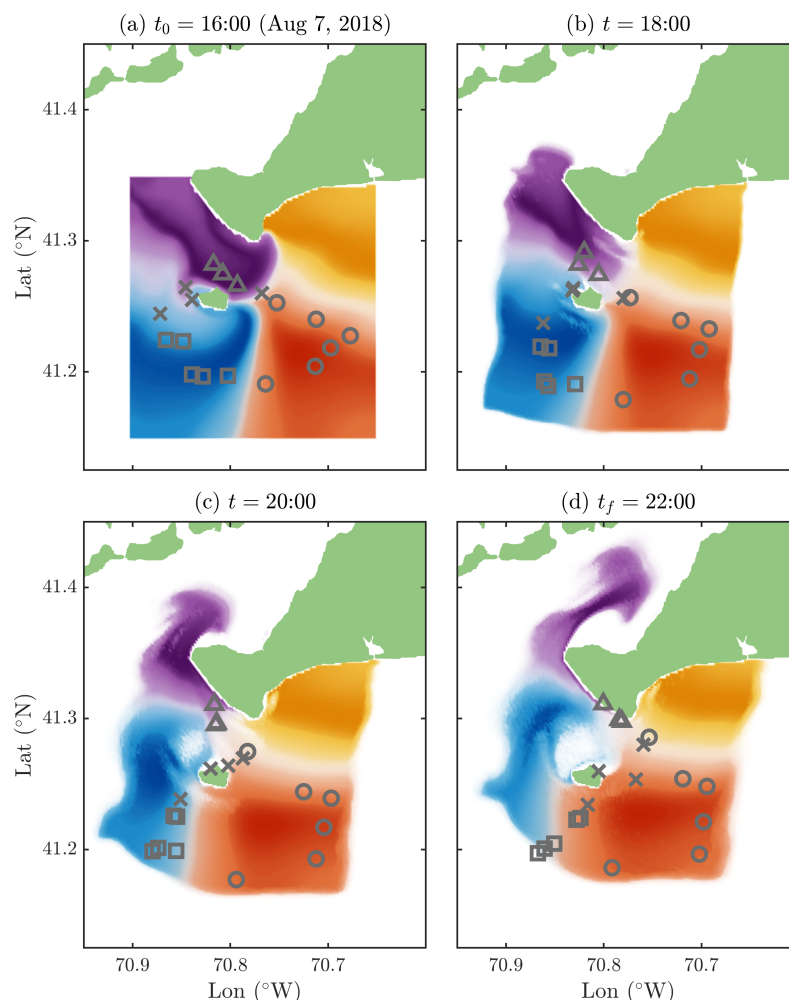


Figure 12. Time evolution of the average clusters in a binned domain. The 18 drifters are marker-coded depending on the cluster in which they started. Crosses represent the 4 drifters initialized at locations of higher uncertainty, with $S_{i,\hat{k}} > 0.1$. The times correspond to (a) $t_0 = 16:00$, (b) $t = 18:00$, (c) $t = 20:00$, and (d) $t_f = 22:00$ UTC. Triangles are initialized on the purple cluster (3 drifters), squares on the blue cluster (5 drifters), and circles on the orange cluster (6 drifters). Corresponding colorbars presented in Figure 11a. (See Supplementary Materials for a video of the time evolution.)

Drifters are marker-coded according to the mean membership probabilities $\bar{p}_{i,k}$ of their spatial position at t_0 . Triangles correspond to cluster 1, squares to cluster 2, circles to cluster 3, and no drifters were initialized in cluster 4. Drifters initialized in higher uncertainty regions (see Figure 11c) are plotted as crosses, and correspond to regions where $S_{i,\hat{k}} > 0.1$, or $\bar{p}_{i,\hat{k}} < 0.7$. The results demonstrate that the drifters predominately remain inside of the forecast model clusters during the first four hours of the time interval, and exceptions to this behavior are associated with higher uncertainty. Consider, for example, the three drifters represented as crosses, initially west of Nomans Land. Two of them move along the northern coast of the island, with one of them beaching and the other ending on cluster 2. The third drifter headed south first, then east, also toward cluster 2. This fact highlights the value of the uncertainty quantification analysis to understand different dynamical behaviors in the flow which are not captured by the central forecast alone. All drifters were eventually advected eastward after 20:00.

Some discrepancies between the clustering predicted behavior and the drifter trajectories were observed. Wind gusts that occurred between 16:00 and 20:00 significantly affected the drifter trajectories. These gusts had not been predicted by the forecasts from the National Centers for Environmental Prediction used by the MSEAS-PE model for the atmospheric forcing [36]. The study was therefore limited by differences between observed flows and model flows, and the coherent structure analysis accuracy is limited to the accuracy of the velocity field used for processing. Nonetheless, drifters released in the clusters with high membership probabilities tend to behave similarly to the clusters, even if they do not precisely match the predicted trajectory behavior. This fact highlights the coherent structure robustness, even in uncertain conditions. The clustering analysis partitions the domain into robust regions where a drifter released in the region will remain there or nearby over the period of analysis. The uncertainty quantification analysis helped identifying the key structures delimiting regions with different transport behaviors, further showing and expanding the applicability of trajectory clustering for studying oceanic flows, despite model imperfections.

6. Conclusions

Ensemble statistics of the trajectory clustering results provide a partitioning of the fluid domain that may provide critical information in emergency response situations, such as search-and-rescue operations, when operational decisions about optimal resource allocation need to be made quickly, accurately, and account for model uncertainties. We presented a modified version of the spectral clustering method with soft membership probabilities, and applied it to fluid particle trajectories to identify coherent structures first in an analytic flow model, then in forecast simulations provided by a coastal ocean model. Uncertainty quantification was applied to assess both the result sensitivity to the clustering method free-parameters and the cluster variability with unknown parameters of the model data. The method sensitivity study, performed on the analytic quasi-periodic Bickley jet system, identified the similarity radius as the free-parameter to which the clusters are most sensitive. To mimic model uncertainty, the Bickley jet parameters were varied to perform a model sensitivity study that highlights the robustness of vortex cores compared to the more uncertain vortex perimeters.

Finally, the method was applied to an ocean ensemble forecast of the coastal region of Martha's Vineyard, and the clustering results were compared to drifter trajectories from a drifter release experiment targeting predicted coherent structures. The forecast clusters from the ensemble analysis provided a good baseline for the drifter behavior, as drifters deployed within each cluster performed similar motions to their corresponding clusters. Ocean transport predictions are challenging due to the complexity of the underlying dynamics governing the flow, and while the Lagrangian approach ultimately depends on the accuracy of the available velocity fields and the quality of the model data, the results presented here demonstrate that the identified clusters are robust to uncertainties in the model and able to predict the main elements of the organization of flow transport. The coupling of the clustering approach to uncertainty quantification can provide a more complete and informative description of flow transport and areas of higher and lower uncertainty within different clusters.

Despite not having been the case for the drifter release presented here, a clustering forecast analysis could be used in planning a drifter deployment for future experiments.

Further refinement of the trajectory clustering method is highly desirable, in particular aiming to reduce parameter sensitivity. The uncertainty quantification with respect to method free-parameters could be used to select parameters that minimize cluster variability, in order to identify clusters that are physical structures, and not a byproduct of the system and parameters chosen. On a more fundamental level, while the method presented here provides robust clusters, it may be possible to improve the method by incorporating fundamental changes to the similarity measure, rather than addressing the sensitivity to free-parameters only. To quantify trajectory similarity, one could not only consider the similarity between the particle spatial coordinates in time, but also that of their velocity vectors, and propose a hybrid notion of similarity. Finally, applying the method to other oceanic forecasts and using the forecast clustering results to plan and execute drifter release experiments can be a promising path to more effective experiments, by increasing the likelihood of released drifters capturing targeted coherent structures and ocean transport barriers.

Supplementary Materials: The following are available online at <http://www.mdpi.com/2311-5521/5/4/184/s1>, Supplementary Material A: Similarity Matrix Sparsification, Supplementary Material B: Video Description, Supplementary Videos 2, 7, 9 and 12 matching the corresponding figures in the manuscript.

Author Contributions: Conceptualization, G.S.V. and M.R.A.; formal analysis, G.S.V.; funding acquisition, I.I.R.; methodology, G.S.V. and M.R.A.; software, G.S.V.; supervision, M.R.A.; visualization, G.S.V.; writing—original draft, G.S.V.; writing—review and editing, G.S.V., I.I.R., and M.R.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by National Science Foundation Division of Atmospheric and Geospace Sciences, award number 1520825.

Acknowledgments: We thank Margaux Filippi and Thomas Peacock for sharing the drifter data set, Pierre Lermusiaux for providing the forecast ensemble, H. M. Aravind and the anonymous referees for their helpful suggestions.

Conflicts of Interest: The authors declare no conflict of interest.

References

- McWilliams, J.C. The emergence of isolated coherent vortices in turbulent flow. *J. Fluid Mech.* **1984**, *146*, 21–43. [[CrossRef](#)]
- Provenzale, A. Transport by coherent barotropic vortices. *Annu. Rev. Fluid Mech.* **1999**, *31*, 55–93. [[CrossRef](#)]
- Haller, G. Lagrangian coherent structures. *Annu. Rev. Fluid Mech.* **2015**, *47*, 137–162. [[CrossRef](#)]
- Haller, G.; Beron-Vera, F.J. Geodesic theory of transport barriers in two-dimensional flows. *Physica D* **2012**, *241*, 1680–1702. [[CrossRef](#)]
- 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 Interdiscip. J. Nonlinear Sci.* **2015**, *25*, 087406. [[CrossRef](#)]
- Allshouse, M.R.; Peacock, T. Lagrangian based methods for coherent structure detection. *Chaos Interdiscip. J. Nonlinear Sci.* **2015**, *25*, 097617. [[CrossRef](#)]
- Allshouse, M.R.; Thiffeault, J.L. Detecting coherent structures using braids. *Phys. D Nonlinear Phenom.* **2012**, *241*, 95–105. [[CrossRef](#)]
- Froyland, G. Dynamic isoperimetry and the geometry of Lagrangian coherent structures. *Nonlinearity* **2015**, *28*, 3587. [[CrossRef](#)]
- Beron-Vera, F.J.; Wang, Y.; Olascoaga, M.J.; Goni, G.J.; Haller, G. Objective detection of oceanic eddies and the Agulhas leakage. *J. Phys. Oceanogr.* **2013**, *43*, 1426–1438. [[CrossRef](#)]
- Haller, G.; Beron-Vera, F.J. Coherent Lagrangian vortices: The black holes of turbulence. *J. Fluid Mech.* **2013**, *731*. [[CrossRef](#)]
- Boffetta, G.; Lacorata, G.; Redaelli, G.; Vulpiani, A. Detecting barriers to transport: A review of different techniques. *Phys. D Nonlinear Phenom.* **2001**, *159*, 58–70. [[CrossRef](#)]
- Peacock, T.; Dabiri, J. Introduction to focus issue: Lagrangian coherent structures. *Chaos Interdiscip. J. Nonlinear Sci.* **2010**, *20*, 017501. [[CrossRef](#)]

13. Jain, A.K.; Murty, M.N.; Flynn, P.J. Data clustering: A review. *ACM Comput. Surv. (CSUR)* **1999**, *31*, 264–323. [CrossRef]
14. Everitt, B.; Landau, S.; Leese, M.; Stahl, D. *Cluster Analysis*, 5th ed.; Wiley Series in Probability and Statistics; Wiley: Hoboken, NJ, USA, 2011.
15. Hadjighasem, A.; Karrasch, D.; Teramoto, H.; Haller, G. Spectral-clustering approach to Lagrangian vortex detection. *Phys. Rev. E* **2016**, *93*, 063107. [CrossRef]
16. Shi, J.; Malik, J. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2000**, *22*, 888–905.
17. Ng, A.Y.; Jordan, M.I.; Weiss, Y. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems*; 2002; pp. 849–856. Available online: <http://papers.nips.cc/paper/2092-on-spectral-clustering-analysis-and-an-algorithm.pdf> (accessed on 11 October 2020).
18. von Luxburg, U. A tutorial on spectral clustering. *Stat. Comput.* **2007**, *17*, 395–416. [CrossRef]
19. Filippone, M.; Camastra, F.; Masulli, F.; Rovetta, S. A survey of kernel and spectral methods for clustering. *Pattern Recognit.* **2008**, *41*, 176–190. [CrossRef]
20. Hadjighasem, A.; Farazmand, M.; Blazeovski, D.; Froyland, G.; Haller, G. A critical comparison of Lagrangian methods for coherent structure detection. *Chaos Interdiscip. J. Nonlinear Sci.* **2017**, *27*, 053104. [CrossRef]
21. Vieira, G.S.; Allshouse, M.R. Internal wave boluses as coherent structures in a continuously stratified fluid. *J. Fluid Mech.* **2020**, *885*, A35. doi:10.1017/jfm.2019.993. [CrossRef]
22. Lermusiaux, P.F. Uncertainty estimation and prediction for interdisciplinary ocean dynamics. *J. Comput. Phys.* **2006**, *217*, 176–199. [CrossRef]
23. Lermusiaux, P.F.; Chiu, C.S.; Gawarkiewicz, G.G.; Abbot, P.; Robinson, A.R.; Miller, R.N.; Haley, P.J.; Leslie, W.G.; Majumdar, S.J.; Pang, A.; et al. Quantifying uncertainties in ocean predictions. *Oceanography* **2006**, *19*, 92–105. [CrossRef]
24. Kratzke, T.M.; Stone, L.D.; Frost, J.R. Search and rescue optimal planning system. In Proceedings of the 2010 13th International Conference on Information Fusion, Edinburgh, UK, 26–29 July 2010; pp. 1–8.
25. Serra, M.; Sathe, P.; Rypina, I.; Kirincich, A.; Ross, S.D.; Lermusiaux, P.; Allen, A.; Peacock, T.; Haller, G. Search and rescue at sea aided by hidden flow structures. *Nat. Commun.* **2020**, *11*, 2525. [CrossRef] [PubMed]
26. Haller, G. Lagrangian coherent structures from approximate velocity data. *Phys. Fluids A* **2002**, *14*, 1851–1861. [CrossRef]
27. Olascoaga, M.J.; Beron-Vera, F.J.; Haller, G.; Trinanes, J.; Iskandarani, M.; Coelho, E.; Haus, B.K.; Huntley, H.; Jacobs, G.; Kirwan, A.; et al. Drifter motion in the Gulf of Mexico constrained by altimetric Lagrangian coherent structures. *Geophys. Res. Lett.* **2013**, *40*, 6171–6175. [CrossRef]
28. Jacobs, G.A.; Bartels, B.P.; Bogucki, D.J.; Beron-Vera, F.J.; Chen, S.S.; Coelho, E.F.; Curcic, M.; Griffa, A.; Gough, M.; Haus, B.K.; et al. Data assimilation considerations for improved ocean predictability during the Gulf of Mexico Grand Lagrangian Deployment (GLAD). *Ocean Model.* **2014**, *83*, 98–117. [CrossRef]
29. Beron-Vera, F.J.; Olascoaga, M.J.; Haller, G.; Farazmand, M.; Triñanes, J.; Wang, Y. Dissipative inertial transport patterns near coherent Lagrangian eddies in the ocean. *Chaos Interdiscip. J. Nonlinear Sci.* **2015**, *25*, 087412. [CrossRef] [PubMed]
30. Williams, M.O.; Rypina, I.I.; Rowley, C.W. Identifying finite-time coherent sets from limited quantities of Lagrangian data. *Chaos Interdiscip. J. Nonlinear Sci.* **2015**, *25*, 087408. [CrossRef]
31. Rypina, I.; Kirincich, A.; Limeburner, R.; Udovychenko, I. Eulerian and Lagrangian correspondence of high-frequency radar and surface drifter data: Effects of radar resolution and flow components. *J. Atmos. Ocean. Technol.* **2014**, *31*, 945–966. [CrossRef]
32. Rypina, I.I.; Kirincich, A.; Lentz, S.; Sundermeyer, M. Investigating the eddy diffusivity concept in the coastal ocean. *J. Phys. Oceanogr.* **2016**, *46*, 2201–2218. [CrossRef]
33. Rypina, I.I.; Pratt, L.J. Trajectory encounter volume as a diagnostic of mixing potential in fluid flows. *Nonlinear Process. Geophys.* **2017**, *24*, 189–202. [CrossRef]
34. Haza, A.; Griffa, A.; Martin, P.; Molcard, A.; Özgökmen, T.; Poje, A.; Barbanti, R.; Book, J.; Poulain, P.; Rixen, M.; et al. Model-based directed drifter launches in the Adriatic Sea: Results from the DART experiment. *Geophys. Res. Lett.* **2007**, *34*, 6. [CrossRef]
35. Haza, A.C.; Özgökmen, T.M.; Griffa, A.; Molcard, A.; Poulain, P.M.; Peggion, G. Transport properties in small-scale coastal flows: Relative dispersion from VHF radar measurements in the Gulf of La Spezia. *Ocean Dyn.* **2010**, *60*, 861–882. [CrossRef]

36. Filippi, M.; Rypina, I.I.; Hadjighasem, A.; Peacock, T. A parameter-free spectral clustering approach to coherent structure detection in geophysical flows. *Fluids* **2020**, submitted.
37. Rypina, I.; Brown, M.G.; Beron-Vera, F.J.; Koçak, H.; Olascoaga, M.J.; Udovydchenkov, I. On the Lagrangian dynamics of atmospheric zonal jets and the permeability of the stratospheric polar vortex. *J. Atmos. Sci.* **2007**, *64*, 3595–3610. [[CrossRef](#)]
38. Beron-Vera, F.J.; Olascoaga, M.J.; Brown, M.G.; Koçak, H.; Rypina, I.I. Invariant-tori-like Lagrangian coherent structures in geophysical flows. *Chaos Interdiscip. J. Nonlinear Sci.* **2010**, *20*, 017514. [[CrossRef](#)]
39. Haley, P.J.; Lermusiaux, P.F. Multiscale two-way embedding schemes for free-surface primitive equations in the “Multidisciplinary Simulation, Estimation and Assimilation System”. *Ocean Dyn.* **2010**, *60*, 1497–1537. [[CrossRef](#)]
40. Haley, P.J., Jr.; Agarwal, A.; Lermusiaux, P.F. Optimizing velocities and transports for complex coastal regions and archipelagos. *Ocean Model.* **2015**, *89*, 1–28. [[CrossRef](#)]
41. Nock, R.; Vaillant, P.; Henry, C.; Nielsen, F. Soft memberships for spectral clustering, with application to permeable language distinction. *Pattern Recognit.* **2009**, *42*, 43–53. [[CrossRef](#)]
42. Schlueter-Kuck, K.L.; Dabiri, J.O. Coherent structure colouring: Identification of coherent structures from sparse data using graph theory. *J. Fluid Mech.* **2017**, *811*, 468–486. [[CrossRef](#)]
43. Lermusiaux, P. On the mapping of multivariate geophysical fields: Sensitivities to size, scales, and dynamics. *J. Atmos. Ocean. Technol.* **2002**, *19*, 1602–1637. [[CrossRef](#)]
44. Ameli, S.; Shadden, S.C. A transport method for restoring incomplete ocean current measurements. *J. Geophys. Res. Ocean.* **2019**, *124*, 227–242. [[CrossRef](#)]
45. Ameli, S.; Shadden, S.C. Trajectory Reconstruction and Analysis for Coherent Structure Evaluation (TRACE). 2020. Available online: <http://transport.me.berkeley.edu/trace/> (accessed on 11 October 2020).

Publisher’s Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).