From large deviations to transport semidistances: coherence analysis for finite Lagrangian data

Péter Koltai*   D.R. Michiel Renger†

Abstract

One way to analyze complicated non-autonomous flows is through trying to understand their transport behavior. In a quantitative, set-oriented approach to transport, finite time coherent sets play an important role. These are time-parametrized families of sets with unlikely transport to and from their surroundings under small or vanishing random perturbations of the dynamics. Here we propose, as a measure of transport, (semi)distances that arise under vanishing perturbations in the sense of large deviations. Analogously, for given finite Lagrangian trajectory data we derive a discrete-time and space semidistance that comes from the "best" approximation of the randomly perturbed process conditioned on this limited information on the deterministic flow. It can be computed as shortest path in a graph with time-dependent weights. Furthermore, we argue that coherent sets are regions of maximal farness in terms of transport, hence they occur as extremal regions on a spanning structure of the state space under this semidistance—in fact, under any distance measure arising from the physical notion of transport. Based on this notion we develop a tool to analyze the state space (or the finite trajectory data at hand) and identify coherent regions. We validate our approach on idealized prototypical examples and well-studied standard cases.

1. Introduction

Transport in dynamical systems. Instrumental to understanding the essential behavior of complicated non-autonomous flows is to grasp how transport is happening in them. This leads on a qualitative level to objects that prohibit transport, commonly named *transport barriers*; often originating from the geometric picture for autonomous systems and that trajectories are unable to cross co-dimension 1 invariant manifolds [28, 29, 31, 30]. For periodically-forced systems, invariant manifolds enclose regions called "lobes" that get transported across these periodically varying manifolds [36, 45].

On a quantitative level, one searches for surfaces of small flux [4, 32, 22], so-called partial barriers [54, 37]. Instead of characterizing regions that do not mix with one another via enclosing them by boundaries of low flux, there are approaches that aim to describe these sets

---

*Institute of Mathematics, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany. E-mail: peter.koltai@fu-berlin.de

†Weierstraß-Institut, Mohrenstraße 39, 10117 Berlin, Germany. E-Mail: d.r.michiel.renger@wias-berlin.de

directly. Such *set-oriented* concepts are strongly interwoven with the theory of transfer operators (Perron–Frobenius and Koopman operators), and comprise *almost-invariant sets* [10], *ergodic partitions* [39] in autonomous, and *coherent sets* [26, 17] in the non-autonomous cases.

Distinctive attention has been given to coherent sets, which are a (possibly time-dependent) family of sets having little or no exchange with their surrounding in terms of transport, and are robust to small diffusion over a finite time of consideration [17, 18]. Natural examples include moving vortices in atmospheric [47, 26], oceanographic [52, 9, 19], and plasma flows [43]. In such applications, one would like to be able to find coherent sets even in the cases when a dynamical model that can be evaluated arbitrarily often is not available, only a finite set of *Lagrangian trajectory data* (passive tracers moving with flow with positions sampled at discrete time instances). This problem has received lot of attention in recent years, and a diverse collection of tools has been developed to tackle it [8, 2, 50, 25, 1, 55, 27, 5, 49, 44, 48, 14, 21].

While other current methods aim at collecting trajectories into coherent sets, in [5] it has been proposed to go one step further and analyze the connectivity structure of the state space under transport with "transport coordinates" and the "skeleton of transport". Very similar observations have been made earlier in [8] in the infinite-time limit for periodically-forced systems. While coherent sets (and transport barriers) aim at *partitioning* the state space, the skeleton is aiming at *"spanning"* the state space with respect to transport. In this respect, coherent sets can be associated with distinct "extremal regions" of the skeleton. Here we will only use this idea of extremality, more precisely that coherent sets are "maximally far" from one another, as measured by transport. To this end we will need to measure "farness" of dynamical trajectories.

Several *dynamical distance measures* have been put forward already to measure the "distance" or "dissimilarity" of trajectories or initial states in dynamical systems [38, 8, 25, 27, 14, 33]. The majority of them are shown to serve their purpose well in revealing coherent structures efficiently and reliably. However, they are either heuristic in the sense that they are not derived from the physical notion of transport, or no discretizations to finite scattered trajectory data have been developed.

The purpose of this paper is thus twofold. On the one hand, we develop a distance measure (a semidistance) between trajectories that is derived from the physical notion of transport subject to diffusion of vanishing strength, and we also derive a discretized distance measure for finite (also possibly sparse and incomplete) Lagrangian data that is consistent with its continuous counterpart in the limit of infinite data. On the other hand, we construct a tool to analyze with such distances the structure of the state space under transport, especially to find coherent sets. This tool makes use of the idea that coherent sets are some sort of extremal regions on a spanning structure with respect to transport, although in this work we will not investigate this "skeleton" in its entirety.

Finite time coherent sets. Let us consider the ordinary differential equation (ODE)

$$\dot{x}_t = v(t, x_t) \tag{1}$$

on some bounded $X \subset \mathbb{R}^d$ and on a finite time interval $[0, T]$ for some $T > 0$. Throughout the paper we will assume that $v : [0, T] \times X \to \mathbb{R}^d$ is a continuous velocity field tangential at the boundary, such that the flow of (1), denoted by $\phi_{s,t}[\cdot]$, $0 \leq s, t \leq T$, is a diffeomorphism on appropriate subsets of X. For $t < s$ we flow backward in time: $\phi_{s,t} = \phi_{t,s}^{-1}$.

Many different notions to characterise coherent sets have been proposed in the literature. Central to all of these notions is the idea that coherent sets should be *robust under noise*; without such requirement any non-intersecting characteristic of a singleton could be considered

a coherent set. To this end one typically perturbs the ODE (1) by a random noise [12, 22, 33], leading to the Itô stochastic differential equation (SDE)[1]

$$d\boldsymbol{x}_t^{(\varepsilon)} = v(t, \boldsymbol{x}_t^{(\varepsilon)})dt + \sqrt{\varepsilon}d\boldsymbol{w}_t \,, \tag{2}$$

where $\{\boldsymbol{w}_t\}_{t\in[0,T]}$ is a Wiener process (Brownian motion) with generator $\phi \mapsto \frac{1}{2}\Delta\phi$, reflecting boundaries, and starting from $\boldsymbol{w}_0 = 0$ (deterministically) and $\varepsilon > 0$ is, at least for now, a given small constant. In fact, the rigorous mathematical formulation of an SDE with reflecting boundaries can be quite subtle, see [3]. We ignore this issue as it doesn't affect our analysis.

According to the definition of *finite time coherent pairs* [26, 17, 35], two sets $A, B \subset X$ are coherent for times 0 and T if most mass from set A is likely to end up in set B, and most mass ending up in set B is likely to originate from set A, that is,

$$\mathbb{P}\big(\boldsymbol{x}_T^{(\varepsilon)} \in B \mid \boldsymbol{x}_0^{(\varepsilon)} \in A\big) \approx 1, \qquad \text{and} \qquad \mathbb{P}\big(\boldsymbol{x}_0^{(\varepsilon)} \in A \mid \boldsymbol{x}_T^{(\varepsilon)} \in B\big) \approx 1. \tag{3}$$

Naturally, for practical purposes one would need to choose how small ε and how large these probabilities should be. As the systems we are dealing with are often deterministic by nature, and there is no "physically straightforward" choice of the diffusion strength ε, our first aim is to remove some of this indeterminacy by quantifying what it means for probabilities to be close to 1 for small ε, in terms of *large deviations* as we explain below[2]. However, it turns out that the forward and backward conditions (3) are essentially equivalent in the large-deviation regime, see Appendix A. Therefore they do not need to be treated separately, and we will consider "forward probabilities" only. Moreover, the large-deviation limits of (3) hardly gives any quantitative information about how coherent two sets might be, as discussed in Appendix A. To conclude, the large-deviations of conditions (3) do not yield sensible conditions for coherence.

Large-deviation based semidistances. In the current paper we take a different approach. We study semidistances that quantify how unlikely it is for mass to flow from one point to another. These are a semidistances in the sense that they satisfy all properties of a metric except for the triangle inequality. In the first part of this paper, sections 2 and 3, we show how such semidistances can arise naturally from probabilistic arguments via large-deviation principles, as we explain below. In the second part, sections 4 and 5, we discuss how (general) semidistances can be used to analyse coherent sets, and we apply the concepts of this paper to a number of examples.

In Section 2 we derive two different semidistances from the *large deviations* of two probabilities. The first one is related to the probability that the random trajectory ends in $\boldsymbol{x}_T^{(\varepsilon)} \asymp \phi_{0,T}[y]$, given that it starts in $\boldsymbol{x}_0^{(\varepsilon)} = x$, for any two initial positions $x, y \in X$. As $\varepsilon \to 0$, the process can no longer deviate from the deterministic flow of (1), and hence this probability will converge to 0 whenever $x \neq y$. In fact, it converges exponentially fast, i.e.

$$\mathbb{P}\big(\boldsymbol{x}_T^{(\varepsilon)} \asymp \phi_{0,T}[y] \mid \boldsymbol{x}_0^{(\varepsilon)} = x\big) \sim \mathrm{e}^{-\frac{1}{\varepsilon}\mu_T(x \to y)} \tag{4}$$

for some function $\mu_T(x \to y) \geq 0$. Such exponential convergence results are called *large-deviation principles*, and $\mu_T(x \to y)$ is the *large-deviation rate*. The less probable it is to reach one point

---

[1]We denote random variables by boldface symbols.

[2]A different way of factoring out diffusion to obtain coherent sets for deterministic flows appeared in the set-oriented transfer-operator based characterization in [18, 23], leading to the notion of the *dynamic Laplacian*. See also our concluding remarks in Section 6.1.

from another, the larger the rate between them. The first semi-distance is then obtained via symmetrisation:

$$\mu_T^{\mathrm{cross}}(x,y) := \mu_T(x{\to}y) + \mu_T(y{\to}x). \tag{5}$$

We call this the *cross* semidistance, since it arises from mass flowing from x to $\phi_{0,T}[y]$ and mass flowing from y to $\phi_{0,T}[x]$ simultaneously and independently, see Figure 1.

The second semidistance arises as the large deviations of the probability for two independent random trajectories $\boldsymbol{x}^{(\varepsilon)}, \boldsymbol{y}^{(\varepsilon)}$, starting from x and y respectively, to meet at any point, within the time interval $(0, T)$, see Figure 2:

$$\mathbb{P}\big(\boldsymbol{x}_T^{(\varepsilon)} \asymp \boldsymbol{y}_T^{(\varepsilon)} \mid \boldsymbol{x}_0^{(\varepsilon)} = x, \boldsymbol{y}_0^{(\varepsilon)} = y\big) \sim \mathrm{e}^{-\frac{1}{\varepsilon}\mu_T^{\mathrm{meet}}(x,y)}, \tag{6}$$

where the *meeting* semi-distance is given by

$$\mu_T^{\mathrm{meet}}(x,y) := \inf_{z \in X} \mu_T(x{\to}z) + \mu_T(y{\to}z).$$
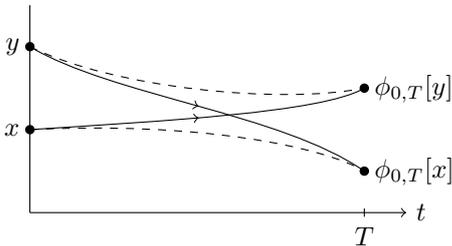
Figure 1: $\mu_T^{\mathrm{cross}}(x,y)$ is the cost to move from x to $\phi_{0,T}[y]$ and from y to $\phi_{0,T}[x]$.
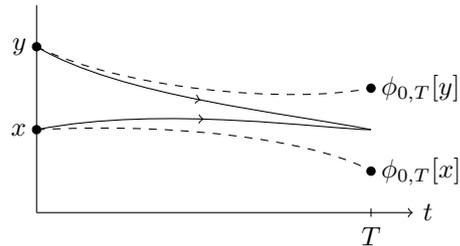
Figure 2: $\mu_T^{\mathrm{meet}}(x,y)$ is the cost for two trajectories to meet.

By this procedure we find two semidistances μ_T^{cross} and μ_T^{meet} that can be used as a measure of "nearness" of points x, y, which will be low for points in the same coherent set, and high otherwise. Since both arise from large-deviation principles, they have a nice additional interpretation as a probabilistic cost or free energy that needs to be paid in order to deviate from the expected flows; such interpretation is common in statistical physics, see for example [42].

Nevertheless, we will see that in order to calculate these costs explicitly, the velocity field v needs to be known. As discussed above, this is in practice seldom the case; mostly one can only assume to have discrete-time snapshots of the positions of a limited number of floaters. With this in mind, we derive similar cost functions as above, that are based on such a finite data set *only*. This will be the content of Section 3. First, the *dynamics* is discretized in time and space by conditioning a usual time-stepping method for the SDE (2) on the event that the random continuous trajectories are to be found in the set of known floater positions at the $K \in \mathbb{N}$ given time instances. This is an optimal choice in the sense that there can be no higher order[3] approximation of the SDE given the limited information about the dynamics. As above, we then derive two large-deviation semidistances $\nu_K^{\mathrm{cross}}(x,y)$ and $\nu_K^{\mathrm{meet}}(x,y)$ that have a clear probabilistic interpretation, that can be used to characterise coherent sets, and that are based on the finite data set rather than on the explicit velocity field. In fact, we will show that these discrete-space-time semidistances are really specific discretisations of the continuous-space-time semidistances μ_T^{cross} and μ_T^{meet}. They can be computed as shortest path lengths in a time-dependent weighted graph. We give an algorithm in Appendix B.

---

[3]In the sense of consistency order of numerical time-stepping methods for ODEs and SDEs [34].

Coherence analysis with semidistances. In Section 4 we describe how in general a semidistance on finite Lagrangian data can be used to analyse coherence. Key to our method is the notion of *cornerstone*: a point that is furthest away from all other points. Cornerstones are though of as "endpoints" of a spanning structure, and ideally each cornerstone is in some sense the center of a coherent set. As a next step, trajectories can be clustered around cornerstones to yield coherent sets. Of course, this approach is very close to the *k-means-* and *fuzzy c-means clustering* of trajectories with respect to dynamical distances in [27, 25], with the important difference that the centers are not chosen by the heuristics of these clustering approaches, but with regard to the properties of coherent sets in the light of transport.

To exemplify the usefulness of the theory put forth in this paper, in sections 4 and 5 we test our approach on a number of standard test cases. Finally, Section 6 discusses possible combinations of this work with other concepts.

2. Large-deviation semidistances in continuous time and space

In this section we study large deviations of the forms (4) and (6). In large-deviations theory it is ofter easier to first study large deviations in a larger space. In our setting, we first study the large deviations of trajectories in Section 2.1 before transforming to the large deviations of the end-points in Section 2.2. We end with a discussion of the resulting semidistances $\mu_T^{\mathrm{cross}}, \mu_T^{\mathrm{meet}}$ in Section 2.3.

2.1. Large deviations of trajectories

We denote trajectories by $w_{(\cdot)}$ to distinguish them from points w. Let $\mathbb{P}$ be the Wiener measure, i.e. the probability that a Brownian trajectory lies in a set $U \subset C(0, T; \mathbb{R}^d)$ is $\mathbb{P}[\boldsymbol{w}_{(\cdot)} \in U]$. Recall that there does not exist a canonical probability measure on the space of trajectories, and so the Wiener measure can not be identified with a meaningful density. This means that one always needs to consider sets rather than particular realisations of the Brownian trajectory. Nevertheless, large-deviation rates are always local, in the sense that they depend on one realisation only (the most likely one in the set U under consideration). This motivates writing $\boldsymbol{w}_{(\cdot)} \asymp f_{(\cdot)}$ if $\boldsymbol{w}_{(\cdot)}$ lies in an infinitesimal neighborhood U of the trajectory $f_{(\cdot)}$. [4]

The large deviations for the SDE (2) are a standard result by Freidlin–Wentzell [16]. This result can be derived via a combination of *Schilder's Theorem* and a *Contraction Principle* as we now explain.

We first consider the noise part $\sqrt{\varepsilon}\boldsymbol{w}_t$, which clearly converges to the constant trajectory 0 as $\varepsilon \to 0$. The corresponding large-deviation principle is given by Schilder's Theorem [11, Th. 5.2.3]:

$$-\varepsilon \log \mathbb{P}\left[\sqrt{\varepsilon}\boldsymbol{w}_{(\cdot)} \asymp w_{(\cdot)}\right] \xrightarrow[\varepsilon \to 0]{} \frac{1}{2}\int_0^T |\dot{w}_t|^2 dt, \tag{7}$$

for differentiable trajectories $w_{(\cdot)}$ starting from $w_0 = 0$ (otherwise the limit will be ∞).

Let us assume that the velocity field $v(t, \cdot)$ is Lipschitz, so for each realisation of the Brownian trajectory $\boldsymbol{w}_{(\cdot)} = w_{(\cdot)}$ corresponds a unique solution $\boldsymbol{x}_{(\cdot)}^{(\varepsilon)}$ of the SDE, starting from some given $\boldsymbol{x}_0^{(\varepsilon)} = x$, see [41, Th. 5.2.1]. The Contraction Principle [11, Th. 4.2.1] then states that the large-deviations rate for a trajectory $\boldsymbol{x}_{(\cdot)}^{(\varepsilon)}$ is given by the minimum of (7) over all realisations of

---

[4]See [11] for the rigorous definition and theoretical background of large-deviation principles.

the noise that give rise to the trajectory $\boldsymbol{x}^{(\varepsilon)}_{(\cdot)}$, i.e.:

$$-\varepsilon \log \mathbb{P}\left[\boldsymbol{x}^{(\varepsilon)}_{(\cdot)} \asymp x_{(\cdot)}\right] \xrightarrow[\varepsilon \to 0]{} \inf_{w_{(\cdot)}\,:\,\dot{x}_t = v(t,x_t)+\dot{w}_t} \frac{1}{2}\int_0^T |\dot{w}_t|^2\,dt = \frac{1}{2}\int_0^T |\dot{x}_t - v(t,x_t)|^2\,dt, \quad (8)$$

for differentiable trajectories $x_{(\cdot)}$ starting from $x_0 = x$.

2.2. Large deviations of end points

We now derive the large-deviation principle of the type (4) as discussed in the introduction. In a sense, the trajectory large deviations (8) encode more information than is needed if we are only interested in the end point $\boldsymbol{x}^{(\varepsilon)}_T \asymp \phi_{0,T}(y)$ of the random trajectory. Another application of the Contraction Principle then states that the large-deviation rate for the end-point is the minimum of (8) over all trajectories starting from x end ending in that given end point $\phi_{0,T}[y]$, i.e.:

$$-\varepsilon \log \mathbb{P}\big(\boldsymbol{x}^{(\varepsilon)}_T \asymp \phi_{0,T}[y] \mid \boldsymbol{x}^{(\varepsilon)}_0 = x\big) \xrightarrow[\varepsilon \to 0]{} \inf_{x_{(\cdot)}\,:\,x_0=x,x_T=\phi_{0,T}[y]} \frac{1}{2}\int_0^T |\dot{x}_t - v(t,x_t)|^2\,dt =: \mu_T(x \to y).$$
$$(9)$$

This defines the 'one-way' rate that we are after.

The sum (5) then defines the cross semidistance $\mu^{\text{cross}}_T(x,y)$, and has a natural interpretation in terms of large deviations: As mentioned in the introduction, it arises from two independent and simultaneous copies $\boldsymbol{x}^{(\varepsilon)}_t, \boldsymbol{y}^{(\varepsilon)}_t$. By independence, the probability that $(\boldsymbol{x}^{(\varepsilon)}_T, \boldsymbol{y}^{(\varepsilon)}_T) \asymp (y,x)$ given $(\boldsymbol{x}^{(\varepsilon)}_0, \boldsymbol{y}^{(\varepsilon)}_0) = (x,y)$ is a product of one-way probabilities, yielding the sum of two one-way rates in the large-deviations, see Figure 1.

A similar argument can be used to derive the meeting large deviations (6). Let $\boldsymbol{x}^{(\varepsilon)}_t$ and $\boldsymbol{y}^{(\varepsilon)}_t$ be two independent copies of the SDE (2), starting from given x and y, respectively. We consider the probability that both trajectories end in a given point, say $\phi_{0,T}[z]$ for some $z \in \mathbb{R}^d$, see Figure 2. Assuming independence of the two trajectories, we immediately get

$$-\varepsilon \log \mathbb{P}\big(\boldsymbol{x}^{(\varepsilon)}_T \asymp \phi_{0,T}[z], \boldsymbol{y}^{(\varepsilon)}_T \asymp \phi_{0,T}[z] \mid \boldsymbol{x}^{(\varepsilon)}_0 = x, \boldsymbol{y}^{(\varepsilon)}_0 = y\big)$$
$$= -\varepsilon \log \mathbb{P}\big(\boldsymbol{x}^{(\varepsilon)}_T \asymp \phi_{0,T}[z] \mid \boldsymbol{x}^{(\varepsilon)}_0 = x\big) - \varepsilon \log \mathbb{P}\big(\boldsymbol{y}^{(\varepsilon)}_T \asymp \phi_{0,T}[z] \mid \boldsymbol{y}^{(\varepsilon)}_0 = y\big)$$
$$\xrightarrow[\varepsilon \to 0]{(9)} \mu_T(x \to z) + \mu_T(y \to z).$$

However, we are only interested in the probability that the two trajectories meet, and not in the point where they meet. A final Contraction Principle thus yields:

$$-\varepsilon \log \mathbb{P}\big(\boldsymbol{x}^{(\varepsilon)}_T \asymp \boldsymbol{y}^{(\varepsilon)}_T \mid \boldsymbol{x}^{(\varepsilon)}_0 = x, \boldsymbol{y}^{(\varepsilon)}_0 = y\big) \xrightarrow[\varepsilon \to 0]{} \inf_{z \in X} \mu_T(x \to z) + \mu_T(y \to z) =: \mu^{\text{meet}}_T(x,y). \quad (10)$$

Observe the two trajectories could also meet earlier and subsequently follow the same trajectory up until time T with zero cost; the time T thus acts as a maximum time at which the trajectories should meet.

2.3. The semidistances

We now discuss some metric properties of the rate functionals. Recall from the introduction that we assumed that the flow is a diffeomorphism. Therefore $\mu_T(x \to y) = 0$ if and only if $x = y$. It is then easy to see that, for any x, y,

(i) $\mu_T^{\mathrm{cross}}(x,y) \geq 0$ and $\mu_T^{\mathrm{meet}}(x,y) \geq 0,$

(ii) $\mu_T^{\mathrm{cross}}(x,y) = 0 \iff x = y$ and $\mu_T^{\mathrm{meet}}(x,y) = 0 \iff x = y,$

(iii) $\mu_T^{\mathrm{cross}}(x,y) = \mu_T^{\mathrm{cross}}(y,x)$ and $\mu_T^{\mathrm{meet}}(x,y) = \mu_T^{\mathrm{meet}}(y,x).$

However, the triangle inequality can fail, and so μ_T^{cross} and μ_T^{meet} are semidistances only.

We point out the following useful relation between the two. Observe that by the definition, $\mu_T^{\mathrm{meet}}(x,y) \leq \mu_T(x{\to}y) + \mu_T(y,y) = \mu_T(x{\to}y)$, and similarly $\mu_T^{\mathrm{meet}}(x,y) \leq \mu_T(y{\to}x)$. Therefore,

$$\mu_T^{\mathrm{meet}}(x,y) \leq \max\left\{\mu_T(x{\to}y), \mu_T(y{\to}x)\right\} \leq \mu_T^{\mathrm{cross}}(x,y).$$

In order to investigate which semidistance is more suitable to study coherence, one would need to study in which setting the gap $\mu_T^{\mathrm{cross}}(x,y) - \mu_T^{\mathrm{meet}}(x,y)$ becomes large. This is beyond the scope of this paper, but we will show for several examples that both work as they should. The inequality does suggest a third candidate: the maximum. Since the maximum of two large-deviation rates does not have a clear large-deviation interpretation we do not study this candidate.

2.4. A simple example

To get an intuition for how the semidistances work, let us consider a simple example, where the domain of interest is the interval $[0, L]$, and there is no dynamics, i.e., $v \equiv 0$. The system is considered on the time interval $[0, T]$. One can then easily see that $\dot{x}_t \equiv L/T$ is optimal in (9), thus giving

$$\mu_T(0{\to}L) = \mu_T(L{\to}0) = \frac{1}{2}\int_0^T \left(\frac{L}{T}\right)^2 dt = \frac{L^2}{2T},$$

and so $\mu_T^{\mathrm{cross}}(0,L) = \frac{L^2}{T}$. Thus, also $\mu_T(0{\to}L/2) = \mu_T(L{\to}L/2) = \frac{L^2}{8T}$. In general, the one-way cost is proportional to the squared distance and inversely proportional to time. This also gives

$$\mu_T^{\mathrm{meet}}(0,L) = \frac{L^2}{4T},$$

so, in this symmetric situation the meeting distance is half of one-way cost and quarter of the cross semidistance.

3. Large-deviation semidistances in discrete time and space

As mentioned in the introduction, the cost functions μ_T^{cross} and μ_T^{meet} are difficult to calculate explicitly, and impossible if the velocity or flow field is not explicitly known. In this section we take a more practical approach. We will assume that the only information at hand is the position at finite times of a finite number I of floaters.

To be more specific, let $\{x_k^{(i)}\}_{k=0,\dots;K,i=1,\dots,I}$ be given positions of floaters $i = 1, \dots, I$ at time $k\tau$ for $k = 0, \dots, K$ for some $\tau > 0$. Assuming that the floaters sample from the deterministic flow field $\phi_{t,s}$, we know that for each floater i,

$$x_{k+1}^{(i)} = \phi_{k\tau,(k+1)\tau}[x_k^{(i)}]. \tag{11}$$

If we add noise to the system, we find random particles described by the set of SDE's,

$$d\boldsymbol{x}_t^{(i,\varepsilon)} = v(t, \boldsymbol{x}_t^{(i,\varepsilon)})dt + \sqrt{\varepsilon}d\boldsymbol{w}_t^{(i)}, \qquad \tilde{\boldsymbol{x}}_0^{(i,\varepsilon)} = x_0^{(i)}, \qquad \text{for } i = 1, \dots, I, \tag{12}$$

where $\boldsymbol{w}^{(i)}$ are now independent standard Brownian motions.

Our strategy is to study the probability that a random particle described by the SDE (2) deviates from the given floater trajectory, conditional to the given positions of all floaters. We first approximate the SDE's (12) by discrete-time, continuous-space Markov processes in Section 3.1. Next, in Section 3.2 we condition these discrete-time processes on the given floater positions. Next we calculate the large-deviation rate for trajectories in Section 3.3, and for end points in Section 3.4. We then end the section with a discussion of the metric properties of the resulting large-deviation rates in Section 3.5.

3.1. Discrete-time approximation

We first focus our attention to one time step $k\tau \to (k+1)\tau$ of one trajectory i, and temporarily drop the superindex for brevity. Since the noise process is a standard Brownian motion, we know its density,

$$\frac{d\mathbb{P}[\sqrt{\varepsilon}\boldsymbol{w}_{(k+1)\tau} \in dy \mid \sqrt{\varepsilon}\boldsymbol{w}_{k\tau} = x]}{dy} = (2\pi\varepsilon\tau)^{-d/2}\exp\left(-\frac{|x-y|^2}{2\varepsilon\tau}\right). \tag{13}$$

Hence, we have precise information on the purely deterministic part of the SDE by (11), and on the purely noise part by (13). We combine this information by using the following time-stepping approximation for the SDE (2).

Fix an $\alpha \in [0,1]$, and let $(\boldsymbol{\xi}_k^+, \boldsymbol{\xi}_k^-)_{k=0,\dots,K}$ be independent normally distributed random variables, each with probability (13) for $x = 0$. Given the approximated random position $\tilde{\boldsymbol{x}}_k$ at time $k\tau$, we iterate

$$\begin{aligned}
\tilde{\boldsymbol{x}}_k^+ &:= \tilde{\boldsymbol{x}}_k + \sqrt{\alpha\tau\varepsilon}\boldsymbol{\xi}_k^+ \\
\tilde{\boldsymbol{x}}_{k+1}^- &:= \phi_{k\tau,(k+1)\tau}[\tilde{\boldsymbol{x}}_k^+] \\
\tilde{\boldsymbol{x}}_{k+1} &:= \tilde{\boldsymbol{x}}_{(k+1)\tau}^- + \sqrt{(1-\alpha)\tau\varepsilon}\boldsymbol{\xi}_{k+1}^-
\end{aligned} \tag{14}$$

Here, $\tilde{\boldsymbol{x}}_k^+$ and $\tilde{\boldsymbol{x}}_{k+1}^-$ are only auxiliary (intermediate) steps. The method (14) is a special case of a *splitting method*, since the deterministic evolution and purely noise parts of the SDE (2) are handled separately in the distinct steps; here it would be "noise-flow-noise". For $\alpha = 1/2$ it is equivalent to the *Strang-splitting* [51], and thus has weak convergence order two [34].

3.2. Conditioning on finite data

Recall that we considered one discrete-time process $(\tilde{\boldsymbol{x}}_k)_{k=1,\dots,K}$ with initial condition $\tilde{\boldsymbol{x}}_0 = x_0^{(i)}$, and that we supressed the dependency on i. For each $k = 0, \dots, K$, we introduce the set

$$\mathcal{A}_k := \{x_k^{(j)}\}_{j=1,\dots I}.$$

of available points at time $t = k\tau$. We now condition the random process on the event that for each realisation $\tilde{\boldsymbol{x}}_k \in \mathcal{A}_k$ *and* for each intermediate point $\tilde{\boldsymbol{x}}_k^+ \in \mathcal{A}_k$. This automatically implies conditioning of the other intermediate points $\tilde{\boldsymbol{x}}_{k+1}^- \in \mathcal{A}_{k+1}$ due to (11). We choose to condition on the intermediate points for practical reasons; otherwise we would not be able to perform the second step in (14), since the discrete trajectories are our only information about the flow, cf Remark 3.1.

The conditioning on the finite data set results in replacing the discrete-time continuous-space process by a fully discrete-time discrete-space Markov chain that hops between the given trajectories. Therefore, the state of the new Markov chain can be represented by the labels

$j = 1, \ldots, I$; this is particularly useful since the deterministic flow (11) will change the positions but not the labels. Since the resulting process is still Markovian, we can fully characterise its behaviour through its transition probabilities for one time-step $k \to k+1$. We now calculate these transition probabilities, dealing with each step in (14) separately. See Figure 3 for a sketch.

For the transition from $\tilde{\boldsymbol{x}}_k$ to $\tilde{\boldsymbol{x}}_k^+$, note that we are in fact conditioning on a null set, so that the conditional probabilities are sensibly defined as the limits over balls $B_r(\cdot)$ of small radii $r \to 0$ around these points. We thus obtain, for any $j, \ell = 1, \ldots I$:

$$
\begin{aligned}
p_k^+(j, \ell) &:= \mathbb{P}\left[\tilde{\boldsymbol{x}}_k^+ = x_k^{(\ell)} \,\big|\, \tilde{\boldsymbol{x}}_k = x_k^{(j)} \text{ and } \tilde{\boldsymbol{x}}_k^+ \in \mathcal{A}_k\right] \\
&= \lim_{r \to 0} \frac{\mathbb{P}\left[\tilde{\boldsymbol{x}}_k^+ \in B_r(x_k^{(\ell)}) \,\big|\, \tilde{\boldsymbol{x}}_k = x_k^{(j)}\right]}{\mathbb{P}\left[\tilde{\boldsymbol{x}}_k^+ \in B_r(\mathcal{A}_k) \,\big|\, \tilde{\boldsymbol{x}}_k = x_k^{(j)}\right]} \\
&= \frac{\exp\left(-\left|x_k^{(\ell)} - x_k^{(j)}\right|^2 / (2\alpha\varepsilon\tau)\right)}{\sum_{\hat{\ell}=1}^I \exp\left(-\left|x_k^{(\hat{\ell})} - x_k^{(j)}\right|^2 / (2\alpha\varepsilon\tau)\right)},
\end{aligned}
\tag{15}
$$

and similarly for the transition from $\tilde{\boldsymbol{x}}_{k+1}^-$ to $\tilde{\boldsymbol{x}}_{k+1}$:

$$
\begin{aligned}
p_{k+1}^-(\ell, m) &:= \mathbb{P}\left[\tilde{\boldsymbol{x}}_{k+1} = x_{k+1}^{(m)} \,\big|\, \tilde{\boldsymbol{x}}_{k+1}^- = x_{k+1}^{(\ell)} \text{ and } \tilde{\boldsymbol{x}}_{k+1} \in \mathcal{A}_{k+1}\right] \\
&= \frac{\exp\left(-\left|x_{k+1}^{(m)} - x_{k+1}^{(\ell)}\right|^2 / (2(1-\alpha)\varepsilon\tau)\right)}{\sum_{\hat{m}=1}^I \exp\left(-\left|x_{k+1}^{(\hat{m})} - x_{k+1}^{(\ell)}\right|^2 / (2(1-\alpha)\varepsilon\tau)\right)}.
\end{aligned}
\tag{16}
$$

Since the transition from $\tilde{\boldsymbol{x}}_k^+$ to $\tilde{\boldsymbol{x}}_{k+1}^-$ is deterministic (middle equation in (14)), we have that,

$$
P_k(j, m) := \mathbb{P}\left[\tilde{\boldsymbol{x}}_{k+1} = x_{k+1}^{(m)} \,\big|\, \tilde{\boldsymbol{x}}_k = x_k^{(j)} \text{ and } \tilde{\boldsymbol{x}}_k^+ \in \mathcal{A}_k, \tilde{\boldsymbol{x}}_{k+1}^- \in \mathcal{A}_{k+1}\right] = \sum_{\ell=1}^I p_k^+(j, \ell) p_{k+1}^-(\ell, m).
\tag{17}
$$

In words, the process performs the following three subsequent steps for *one* time step (see Figure 3):

1. Start in $x_k^{(j)}$, and perform a jump to some $x_k^{(\ell)}$ with probability $p_{j,\ell}^{(k,+)}$,

2. Perform a deterministic jump from $x_k^{(\ell)}$ to $x_{k+1}^{(\ell)}$,

3. Perform a jump from $x_{k+1}^{(\ell)}$ to $x_{k+1}^{(m)}$ with probability $p_{\ell,m}^{(k+1,-)}$.

The transition probabilities $P_k(j, m)$ define our new, discrete-time Markov chain $(\boldsymbol{i}_k)_{k=0,\ldots,K}$ on the discrete space $\{1, \ldots, I\}$. To shorten notation, we will write $i_{(\cdot)} := (i_k)_{k=0,\ldots,K}$ for a discrete trajectory, analogous to the continuous-time setting. By the Markov property, the probability that the Markov chain realises such trajectory is simply

$$
\mathbb{P}\left[\boldsymbol{i}_{(\cdot)} = i_{(\cdot)}\right] = \prod_{k=0}^{K-1} P_k(i_k, i_{k+1}),
\tag{18}
$$

where we assumed that the chain starts (deterministically) from i_0.
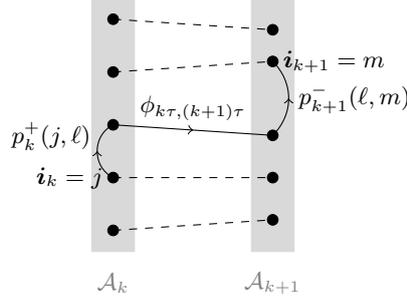
Figure 3: One time-step of the discrete-time discrete space Markov chain $\boldsymbol{i}_k$.

3.3. Large deviations of discrete trajectories

We now study the large deviations of the discrete Markov chain $\boldsymbol{i}_k$. Similarly to the continuous setting from Section 2 we start from the large deviations of trajectories. First we calculate the large deviations for $p_k^+(j,\ell)$ and $p_{k+1}^-(\ell,m)$. By the Laplace principle (26),

$$-\varepsilon \log p_k^+(j,\ell) \overset{(15)}{=} \varepsilon \log \sum_{\hat{\ell}=1}^I \exp \left(\frac{\left|x_k^{(\ell)} - x_k^{(j)}\right|^2 - \left|x_k^{(\hat{\ell})} - x_k^{(j)}\right|^2}{2\alpha\varepsilon\tau} \right)$$

$$\xrightarrow[\varepsilon \to 0]{} \max_{\hat{\ell}=1,\dots,I} \frac{\left|x_k^{(\ell)} - x_k^{(j)}\right|^2 - \left|x_k^{(\hat{\ell})} - x_k^{(j)}\right|^2}{2\alpha\tau} = \frac{|x_k^{(\ell)} - x_k^{(j)}|^2}{2\alpha\tau}.$$

We will make this simplification again below. Similarly, we obtain

$$-\varepsilon \log p_{k+1}^-(\ell,m) \xrightarrow[\varepsilon \to 0]{(16)} \frac{|x_{k+1}^{(m)} - x_{k+1}^{(\ell)}|^2}{2(1-\alpha)\tau}.$$

Using these two exponential approximations, we can again use the Laplace principle (26) to find for the jump probability of one time step:

$$\lim_{\varepsilon \to 0} -\varepsilon \log P_k(j,m) \overset{(17)}{=} \lim_{\varepsilon \to 0} -\varepsilon \log \sum_{\ell=1}^I p_k^+(j,\ell) p_{k+1}^-(\ell,m)$$

$$= \min_{\ell=1,\dots,I} \lim_{\varepsilon \to 0} \left(-\varepsilon \log p_k^+(j,\ell) - \varepsilon \log p_{k+1}^-(\ell,m) \right)$$

$$= \min_{\ell=1,\dots,m} \frac{|x_k^{(\ell)} - x_k^{(j)}|^2}{2\alpha\tau} + \frac{|x_{k+1}^{(m)} - x_{k+1}^{(\ell)}|^2}{2(1-\alpha)\tau}.$$

Finally the large-deviation rate of a discrete trajectory is

$$-\varepsilon \log \mathbb{P}\left[\boldsymbol{i}_{(\cdot)} = i_{(\cdot)}\right] \overset{(18)}{=} -\varepsilon \log \prod_{k=0}^{K-1} P_k(i_k, i_{k+1})$$

$$\xrightarrow[\varepsilon \to 0]{} \sum_{k=0}^{K-1} \min_{\ell=1,\dots,I} \frac{|x_k^{(\ell)} - x_k^{(i_k)}|^2}{2\alpha\tau} + \frac{|x_{k+1}^{(i_{k+1})} - x_{k+1}^{(\ell)}|^2}{2(1-\alpha)\tau} := \mathcal{J}(i_{(\cdot)}), \qquad (19)$$

Remark 3.1: Recall that we conditioned on the event that all $\tilde{\boldsymbol{x}}_k$ as well as the intermediate points $\tilde{\boldsymbol{x}}_k^+$ lie in the set $\mathcal{A}_k$ of available points. One might argue that in practice only the points

$\tilde{\boldsymbol{x}}_k$ are measured to lie in $\mathcal{A}_k$, while the other two are mathematical constructs that may lie anywhere. However, if we would relax this conditioning and follow the calculations as above, we would find:

$$-\varepsilon \log P_k(j,m) \xrightarrow[\varepsilon \to 0]{} \min_{x \in \mathbb{R}^d} \left\{ \frac{|x - x_k^{(j)}|^2}{2\alpha\tau} + \frac{\left|x_{k+1}^{(m)} - \phi_{t_k, t_{k+1}}[x]\right|^2}{2(1-\alpha)\tau} - \min_{\hat{m}=1,\ldots,I} \frac{\left|x_{k+1}^{(\hat{m})} - \phi_{t_k, t_{k+1}}[x]\right|^2}{2(1-\alpha)\tau} \right\}.$$

Since this large-deviation rate still depends on the unknown flow field ϕ, it can not be used if only the data of a finite number of floaters is available.

Remark 3.2 (Missing data and non-uniform time-sampling)**:** Note that the construction works exactly as described above even if information about trajectories is partially missing. The conditioning on the set $\mathcal{A}_k$ works identically, but now these sets might have different cardinalities smaller or equal I. Observe that our only information about the deterministic flow for times in $[k\tau, (k+1)\tau)$ comes from those trajectories that are available both in $\mathcal{A}_k$ and $\mathcal{A}_{k+1}$. If this intersection is empty, we need to skip that time slice completely. This is not a problem, since our choice of sampling time uniformly by the step size τ was solely in order to ease presentation. As the reader has probably observed, the extension for varying time steps τ_k is straightforward.

3.4. Large deviations of end points

Analogously to the continuous setting, we study the large deviations of the one-way probability to hop from i to j in discrete time K, and the meeting probability that two independent chains, starting from i and j respectively, meet in discrete time K. Since the trajectories (19) encode more information than the end points, we can now easily derive the large deviations of the one-way probability by a Contraction Principle. Indeed, for any two indices $i, j = 1, \ldots, I$,

$$-\varepsilon \log \mathbb{P}[\boldsymbol{i}_K = j \mid \boldsymbol{i}_0 = i] \xrightarrow[\varepsilon \to 0]{} \min_{i_{(\cdot)}: i_0 = i, i_K = j} \mathcal{J}(i_{(\cdot)}) =: \nu_K(i \to j), \tag{20}$$

where $\mathcal{J}$ is the discrete-trajectory large-deviation rate (19). Note that $\mathcal{J}$ is the shortest path length in a graph with time-dependent edge weights

$$w_k(i,j) = \min_{\ell=1,\ldots,I} \frac{|x_k^{(\ell)} - x_k^{(i)}|^2}{2\alpha\tau} + \frac{|x_{k+1}^{(j)} - x_{k+1}^{(\ell)}|^2}{2(1-\alpha)\tau}.$$

Again, the sum $\nu_K^{\mathrm{cross}}(i,j) := \nu_K(i \to j) + \nu_K(j \to i)$ can be given an interpretation in terms of large deviations as in Section 2.2. Moreover, following the same argument as in (10), if we take two independent trajectories $\boldsymbol{i}_{(\cdot)}$ and $\boldsymbol{j}_{(\cdot)}$, then

$$-\varepsilon \log \mathbb{P}[\boldsymbol{i}_K = \boldsymbol{j}_K \mid \boldsymbol{i}_0 = i, \boldsymbol{j}_0 = j] \xrightarrow[\varepsilon \to 0]{} \min_{\ell=1,\ldots,I} \nu_K(i \to \ell) + \nu_K(j \to \ell) =: \nu_K^{\mathrm{meet}}(i,j).$$

3.5. The semidistances

It is easily checked that in the discrete setting the properties of a semidistance are also satisfied:

(i) $\nu_K^{\mathrm{cross}}(i,j) \geq 0$ and $\nu_K^{\mathrm{meet}}(i,j) \geq 0$,

(ii) $\nu_K^{\mathrm{cross}}(i,j) = 0 \iff i = j$ and $\nu_K^{\mathrm{meet}}(i,j) = 0 \iff i = j$,

(iii) $\nu_K^{\mathrm{cross}}(i,j) = \nu_K^{\mathrm{cross}}(j,i)$ and $\nu_K^{\mathrm{meet}}(i,j) = \nu_K^{\mathrm{meet}}(j,i)$.

11

Furthermore, the triangle inequality fails, but we again have the following estimate:

$$\nu_K^{\mathrm{meet}}(i,j) \leq \max\left\{\nu_K(i{\to}j), \nu_k(j{\to}i)\right\} \leq \nu_K^{\mathrm{cross}}(i,j).$$

Both semidistances have the form of a shortest-path cost, where the cost of a path is given by (9). We stress that this expression is fairly simple, and depends on the flow field through the *known* positions of the floaters $x_k^{(\ell)}$ only. Because of this: 1) these costs can be used in practice if the velocity field is unknown (Section 4 and Section 5); 2) these costs can even be applied to cases where there may not be an underlying velocity field, as for example in discrete-time dynamical system (Section 4.1).

These semidistances can be computed by first computing the one-way rates $\nu_K(i \to j)$ using Algorithm 1, see Appendix B. From these rates one readily obtains the semidistances via

$$\nu_K^{\mathrm{cross}}(i,j) = \nu_K(i \to j) + \nu_K(j \to i) \quad \text{and} \quad \nu_K^{\mathrm{meet}}(i,j) = \min_{\ell=1,\dots,I} \nu_K(i \to \ell) + \nu_K(j \to \ell).$$

Remark 3.3: Other large-deviation-based semidistances are also possible. If one considers the "noise-flow" (i.e., $\alpha = 1$) time-stepping scheme for the SDE rather then "noise-flow-noise", expression (19) simplifies a bit. Then, as another example of a large-deviation-based semidistance between two given discrete paths $\{x_k^{(i)}, x_k^{(j)}\}_{k=0,\dots K}$, one could consider the probability to hop back and forth between the two trajectories, see Figure 4. In that case we find in the large-deviations scaling for $\alpha = 1$:

$$-\varepsilon \log \mathbb{P}[\boldsymbol{i}_1 = j_1, \boldsymbol{i}_2 = i_2, \cdots \mid \boldsymbol{i}_0 = i_0] \xrightarrow[\varepsilon \to 0]{} \sum_{k=0}^{K-1} \frac{|x_k^{(i)} - x_k^{(j)}|^2}{2\tau}, \tag{21}$$

for $\alpha = 0$ the sum would go from $k = 1$ to K. Naturally, this is simply the L^2-distance between two trajectories, as considered earlier in [25]. Although this construction is very easy to calculate and its square root is a genuine metric, it is less interpretable as a transport cost.
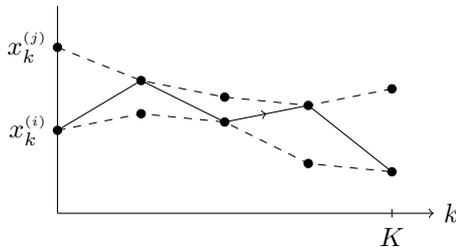
Figure 4: Hopping back and forth (solid line) between two given trajectories (dashed lines).

Remark 3.4: It should be noted that the semidistances $\nu_K^{\mathrm{meet}}, \nu_K^{\mathrm{cross}}$ scale quadratically in space; this becomes even more apparent in the example considered in Section 3.7. In the case of the L^2-distance (21), the cost becomes a genuine distance after taking the square root. However, if we take the square roots of ν_K^{meet} and ν_K^{cross}, the triangle inequality still fails. We therefore stick to the quadratic scaling as this has the most direct interpretation as large-deviation costs.

3.6. Discretisation of the continuous semidistances

We now show that the one-way discrete space-time cost ν_K can also be obtained by discretising the continuous space-time cost μ_T. We will not be precise about the discretisation error; of

course one needs to assume that the number of floaters is sufficiently large. We first divide the time interval into subintervals $[0, T) = \bigcup_{k=0}^{K-1} [k\tau, (k+\alpha)\tau) \cup [(k+\alpha)\tau, (k+1)\tau)$. Recall that $\phi_{t_0,t}$ is the flow associated to $v(t, \cdot)$, that is, for any t_0, x,

$$\partial_t \phi_{t_0,t}[x] = v\big(t, \phi_{t_0,t}[x]\big).$$

In each interval $[k\tau, (k+\alpha)\tau)$ we approximate:

$$\dot{x}_t \approx \frac{x_{(k+\alpha)\tau} - x_{k\tau}}{\alpha\tau} \qquad \text{and} \qquad v(t, x_t) \approx \frac{x_{(k+\alpha)\tau} - \phi_{(k+\alpha)\tau, k\tau}[x_{(k+\alpha)\tau}]}{\alpha\tau}.$$

In each interval $[(k+\alpha)\tau, (k+1)\tau)$ we approximate:

$$\dot{x}_t \approx \frac{x_{(k+1)\tau} - x_{(k+\alpha)\tau}}{(1-\alpha)\tau} \qquad \text{and} \qquad v(t, x_t) \approx \frac{\phi_{(k+\alpha)\tau, (k+1)\tau}[x_{(k+\alpha)\tau}] - x_{(k+\alpha)\tau}}{(1-\alpha)\tau}.$$

Because of the assumption that the flow is one-to-one, we can always write $x_{(k+\alpha)\tau} = \phi_{k\tau, (k+\alpha)\tau}[\hat{x}_k]$ for some $\hat{x}_k$. Note that $x_{(\cdot)}$ is some path, not necessarily a trajectory of the flow. We thus obtain:

$$\frac{1}{2} \int_0^T \big|\dot{x}_t - v(t, x_t)\big|^2 \, dt \approx \sum_{k=0}^{K-1} \frac{|x_{k\tau} - \hat{x}_k|^2}{2\alpha\tau} + \frac{|x_{(k+1)\tau} - \phi_{k\tau, (k+1)\tau}[\hat{x}_k]|^2}{2(1-\alpha)\tau}.$$

Since the number of floaters $\{x_k^{(i)}\}_{k=0,\dots,K; i=1,\dots,I}$ is large,

$$\mu_T(x_0^{(i)} \to x_K^{(j)}) \approx \inf \bigg\{ \frac{1}{2} \int_0^T |\dot{x}_t + v(t, x_t)|^2 \, dt : x_0 = x_0^{(i)}, x_T = x_K^{(j)}, x_{k\tau} \in \mathcal{A}_k,$$

$$x_k^{(\ell)} := \phi_{(k+\alpha)\tau, k\tau}[x_{(k+\alpha)\tau}] \in \mathcal{A}_k \bigg\}$$

$$\approx \min_{i_{(\cdot)}: i_0 = i, i_K = j} \sum_{k=0}^{K-1} \min_{\ell=1,\dots,I} \frac{|x_k^{(i_k)} - x_k^{(\ell)}|^2}{2\alpha\tau} + \frac{|x_{k+1}^{(i_{k+1})} - x_{k+1}^{(\ell)}|^2}{2(1-\alpha)\tau}$$

$$= \nu_K(i \to j).$$

This shows that we can either derive the large-deviations rate function in continuous space and discretize this to finite trajectories (as done here), or we can restrict the continuous dynamics to finite trajectory data and derive a large-deviations rate function for that (as done above); we obtain consistent results whichever route we take.

3.7. The simple example revisited

Let us now demonstrate how the results of this section apply to the example of Section 2.4.

Discrete time and continuous space. Let us first suppose we are given infinitely many "trajectories" of the system, one starting at each point $x \in [0, L]$, and they are sampled at discrete time points $k\tau$, $k = 0, 1, \dots, K$, with $\tau = \frac{T}{K}$. From Section 3.3 with $\alpha = 1/2$ we obtain, by writing $\Delta x = \frac{L}{K}$, that

$$\nu_K(0 \to L) = \frac{1}{2} \sum_{k=1}^{K} \frac{(\Delta x)^2}{\tau} = \frac{1}{2} K \cdot \frac{(L/K)^2}{T/K} = \frac{L^2}{2T},$$

where we used that the optimal discrete trajectory in (20) is the one making jumps of equal lengths Δx. Note that the rate function is identical to that in the fully continuous case. Analogously, $\nu_K(0 \to L/2) = \nu_K(L/2 \to L) = \frac{L^2}{8T}$, and generally, if $|x - y| = \delta$, then $\nu_K(x \to y) = \frac{\delta^2}{2T}$. The derived semidistances scale similarly. Note that the semidistances converge to zero as $T \to \infty$.

13

Discrete time and space. If we are given a finite number I of equispaced trajectories of this system sampled at the same times as in the previous paragraph, the virtual random walker cannot make arbitrarily small jumps as in the continuous state case, thus

$$\nu_K(0{\to}L) = \frac{1}{2}\sum_{k=1}^{K}\frac{(\Delta x_k)^2}{\tau} \approx \frac{L^2}{2T}\,, \qquad \text{if } K \le I\,,$$

since we can take $\Delta x_k \approx L/K$ with error $\mathcal{O}(I^{-1})$ as I grows. However, if $K > I$, the smallest jumps are $\Delta x_k = \frac{L}{I}$, thus

$$\nu_K(0{\to}L) = \frac{1}{2}I \cdot \frac{(L/I)^2}{\tau} = \frac{L^2 K}{2I\,T}\,.$$

Thus, if the observation time of trajectories grows and they are still observed at the same rate (i.e., τ stays constant), the semidistances saturate at $\frac{L^2}{I\tau}$ and do not converge to zero as in the continuous time case. Moreover, to reach $y = L/2$ from $x = 0$, we still cannot make smaller jumps than $\Delta x = \frac{L}{I}$, but now we only require only $I/2$ of them, such that we obtain $\nu_K(0{\to}L/2) = \nu_K(L/2{\to}L) = \frac{I}{2} \cdot \frac{(L/I)^2}{\tau} = \frac{L^2}{4I\tau}$ (for even I, and vanishing error for odd I as I grows). The main lesson is, that while in the continuous space case halving the Euclidean distance makes the semidistance scale by $\frac{1}{4}$, if the spatial resolution of trajectories is coarse, the discrete semidistance scales only by $\frac{1}{2}$. In general, if $|x - y| = \delta$, then on a coarse resolution grid it takes about $\frac{\delta}{\Delta x}$ jumps to travel between these two points, and we obtain $\nu_K(x{\to}y) \approx \frac{\delta}{\Delta x} \cdot \frac{(\Delta x)^2}{\tau} = \delta \cdot \frac{\Delta x}{\tau}$. Note that Δx and τ are constant quantities, and thus the one-way cost scales linearly in the Euclidean distance between the two points, as opposed to quadratic scaling in the continuous space case.

4. Coherence analysis with semidistances

Suppose we are given a set of discrete time and space trajectories $\{x_k^{(j)}\}_{j=1,\dots,I, k=0,\dots,K}$, and a (semi)distance d. We now describe how such semidistances can be used to distinguish and analyse coherent sets from the finite data. We shall work with an unspecified semidistance d, but of course the semidistances that we have in mind are ν_K^{cross} and ν_K^{meet} that we derived in the previous section. Other—not large-deviations based—distance measures could be used just as well, as we discuss below. Nevertheless, the semidistances should not be completely arbitrary; we assume that they share the behavior of ν_K^{meet} and ν_K^{cross} that we discuss in Section 4.1.

To illustrate the ideas we first analyze the behavior of two one-dimensional prototypical examples. These examples show the difference between two types of regions: "mixing" and "static" (also known as "regular"). In these one-dimensional and simple examples, one can easily determine the regions and whether they are mixing or static from the semidistances. One example has two invariant sets under the dynamics, which is (measure-theoretically and topologically) mixing on both of them. The other has two static regions, where the mutual physical distance of trajectories does not change under the dynamics, and these regions are separated by a third, mixing region.

After this we proceed with a more involved model: a two-dimensional periodically forced double gyre flow, where the boundaries of the separate regions are no longer as clear-cut as in the one-dimensional example. Nevertheless we will show that one can identify the separate regions via the tools that we present next.

First we introduce the notion of cornerstones, representing possible coherent sets or mixing regions, how to find them and when to stop searching for them. Finally, to obtain coherent

sets, we assign the trajectories to cornerstones. The notion of fuzzy affiliations will be used to express the uncertainty whether a trajectory close to the boundary of a set belongs to it or not. Note that in the case of finite data such an uncertainty is always present.

4.1. Two illustrative model cases

Two invariant, mixing subdomains. As mentioned in Section 3.5, we may also apply the techniques developed in this paper to a discrete-time dynamical system. To gain some intuition for the behavior of the semidistances at mixing regions, we consider the discrete-time system on the unit interval $X = [0, 1]$ and one-step flow map, see Figure 5 (left),

$$\phi(x) = \begin{cases} 4x \mod \frac{1}{2}, & x < \frac{1}{2} \\ \left(4(x - \frac{1}{2}) \mod \frac{1}{2}\right) + \frac{1}{2}, & x \geq \frac{1}{2}. \end{cases}$$
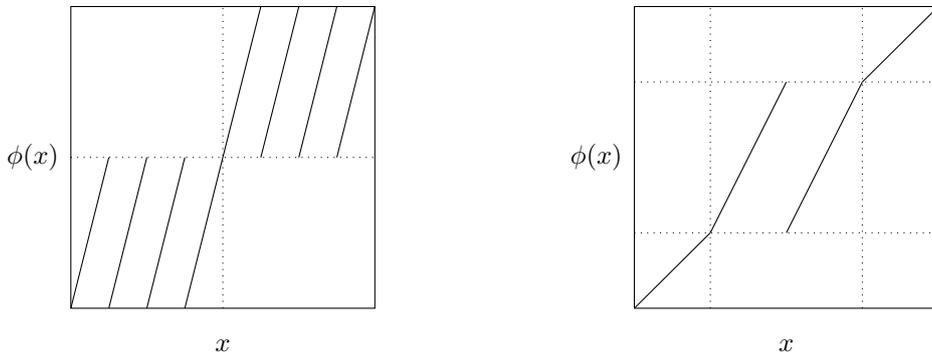
Figure 5: Left: The time-discrete flow map with two invariant mixing subdomains. Right: The time-discrete flow map with two static regions and an invariant mixing subdomains between them.

The sets $X_1 = [0, \frac{1}{2}]$, $X_2 = (\frac{1}{2}, 1]$ are invariant, i.e., $\phi^{-1}(X_1) = X_1$ and $\phi^{-1}(X_2) = X_2$, and ϕ is simply the circle-quadrupling map on each of these sets, i.e., it is mixing on the single components. Consequently[5],

$$\liminf_{t \in \mathbb{N}, \, t \to \infty} \left| \phi^t(x) - \phi^t(y) \right| = 0 \tag{22}$$

for (Lebesgue-)almost every pair $x, y \in X_i$, $i = 1, 2$.

Thus, $\nu_K(i \to j) \to 0$ as $K \to \infty$, because if x_i, x_j are both in X_1 or both in X_2, then (22) shows that their trajectories get arbitrarily close eventually. If the trajectories start in different halves of $[0, 1]$, then a similar reasoning as in Footnote 5 shows that $(\phi^t(x_i), \phi^t(x_j))$ gets arbitrarily close to $(\frac{1}{2}, \frac{1}{2})$ eventually, thus the cost of a jump from one trajectory to another gets arbitrarily cheap. See Figure 6.

The transport semidistances between any two points within the same region are very small, at least if the time window is large enough. This behaviour is typical for mixing regions. In fact, since the two mixing regions are only separated by one point, it is relatively cheap to move from one region to the other, and so the semidistances between two points in separate regions converge with increasing time to zero. Nevertheless, the semidistances still detect a difference

---

[5]If a system (X, ϕ, μ) is mixing, then $(X \times X, \phi \times \phi, \mu \times \mu)$ is ergodic [53, Theorem 1.24]. Thus, for $\mu \times \mu$-almost every pair (x, y), the trajectory $(\phi \times \phi)^k(x, y)$ will enter every set of non-zero measure for some $k \geq 0$, thus also $\{(x, y) \mid |x - y| < \varepsilon\}$ for any fixed $\varepsilon > 0$. This shows (22).
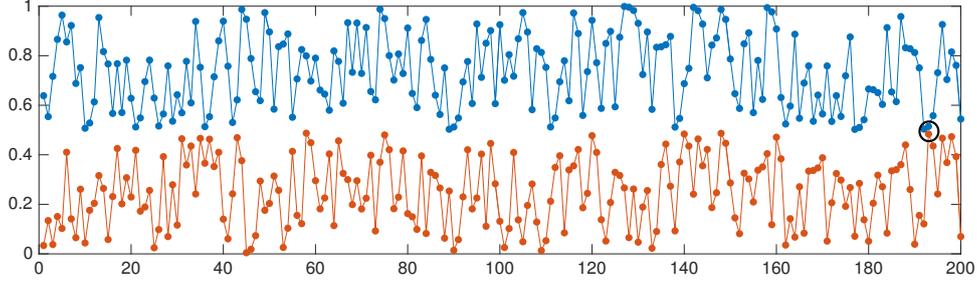
Figure 6: Two trajectories of the map ϕ of length 200 steps, starting in X_1 and X_2, respectively. Theory shows that they come arbitrary close, eventually. Here they get the closest at time step 193, shown by a circle.

between the two invariant sets: the semidistance between two trajectories in the same invariant component goes to zero quicker than the one between two from different components, as shown in Figure 7 for $I = 100$ initially equispaced trajectories. Thus, it is the relative difference between the semidistances that is relevant for the transport-structure of the state space, and not the absolute values.
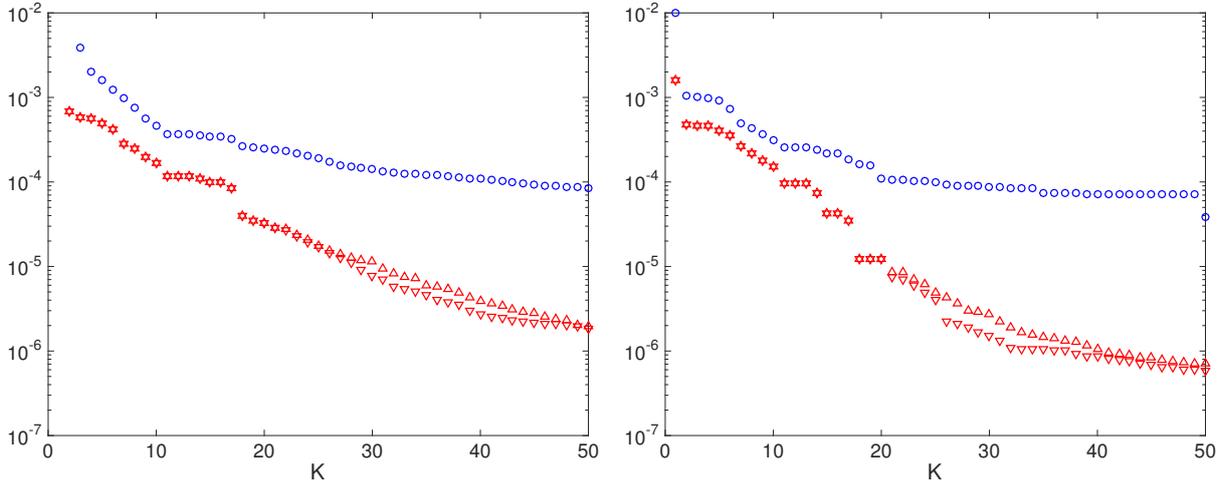
Figure 7: Semidistances $\nu_K^{\mathrm{cross}}(i \to j)$ (left) and $\nu_K^{\mathrm{meet}}(i, j)$ (right) for increasing maximal time K, averaged over $x_i, x_j \in X_1$ (downward-pointing triangles), $x_i, x_j \in X_2$ (upward-pointing triangles), and $x_i, \in X_1$, $x_j \in X_2$ (circles), respectively. Note that the decrease of the distance is much slower for trajectories taken from different invariant sets.

Two static regions divided by a mixing one. To gain some intuition about static regions, let us now consider the discrete-time system on $X = [0, 1]$ given by

$$\phi(x) = \begin{cases} x, & x \in [0, \frac{1}{4}) \cup (\frac{3}{4}, 1] \\ \left(2(x - \frac{1}{4}) \mod \frac{1}{2}\right) + \frac{1}{4}, & x \in [\frac{1}{4}, \frac{3}{4}], \end{cases}$$

see Figure 5 (right). This map has three invariant sets. The left and right ones are static, such that the mapping restricted to them is the identity, and are meant to model regions of the state

16

space in complicated flows, that are "static" in the sense that the mutual distance of points is not changed (or just barely) by the dynamics. We will consider these as one kind of prototype for coherent sets. The third region is mixing, and physically separates the other two.

We take $I = 100$ initially equispaced trajectories and compute the one-way costs $\nu_K(i{\to}\cdot)$ with $K = 50$ for $i = 1$ and $i = 51$, respectively, shown in Figure 8. From our analysis in
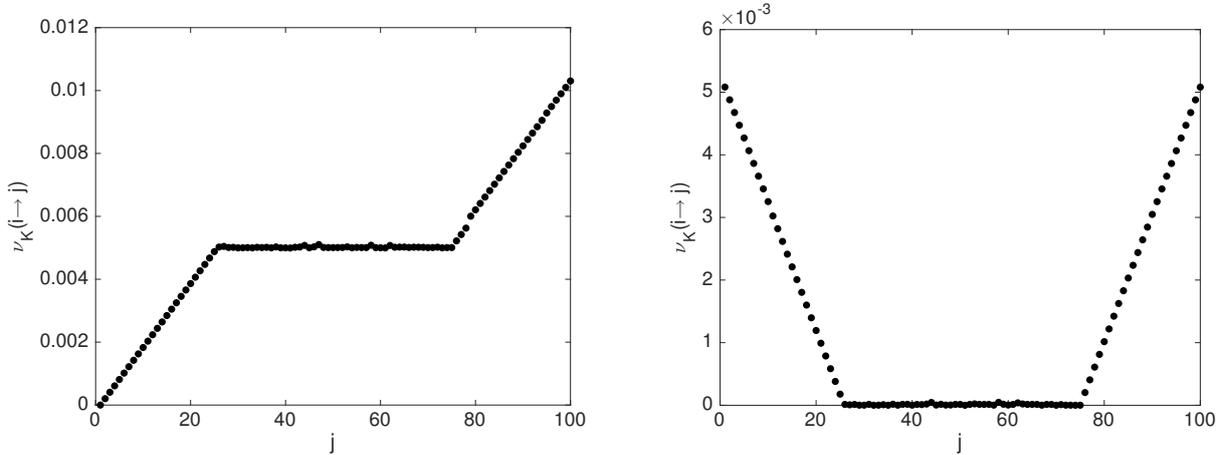
Figure 8: One-way cost $\nu_K(i{\to}\cdot)$ for $i = 1$ (left) and $i = 51$ (right) for the map with two static and one mixing region.

Section 3.7 we would have expected to see quadratic growth of the one-way cost with respect to physical distance in the static regions, but we only observe linear growth. This is due to the finite number of considered trajectories, as also explained in the second paragraph of Section 3.7. The points of the mixing region have all almost constant cost from points in the static regions, and approximately zero cost from one another. To obtain the cost between two points of different static regions, one has to consider the cost to go to the boundary of the static and mixing regions (linear cost in Euclidean distance), travel on a trajectory from there to the boundary of the other static region (at zero cost), and then go from there to the desired point (again, linear cost in Euclidean distance that needs to be covered). Thus, the cost (and semidistance) between these two points is the sum of their one-way cost (and semidistance) to the mixing region, provided the time of consideration is sufficiently large for the mixing to take place.[6] Since our fictive random walker uses trajectories (essentially only one of them) of the mixing region to travel from one static region to the other, we will also call it *transition region* henceforth.

To conclude, from Figure 8 we can easily identify three separate regions, and from the steepness of the slopes (linear/quadratic or flat), we can determine wether a region is static or mixing. As the next example shows, this distinction is usually not as clear as in these one-dimensional examples, but the main ideas will be based on this observation.

---

[6]Note that for this argument, ergodicity of the dynamics in the "mixing" region would be sufficient, since one only needs to travel "from one static region to the other". The crucial additional property we get from mixingness is that the mutual semidistances of points in this region go to zero.

4.2. The periodically forced double gyre

Let us now consider the non-autonomous system $\dot{x}_t = v(t, x_t)$ on $X = [0, 2] \times [0, 1]$ with [24]

$$v(t, x) := \begin{bmatrix} -\pi A \sin(\pi f(t, x)) \cos(\pi y) \\ \pi A \cos(\pi f(t, x)) \sin(\pi y) \frac{df}{dx}(t, x) \end{bmatrix}, \tag{23}$$

where $f(t, x) = \beta \sin(\omega t) x^2 + (1 - 2\beta \sin(\omega t)) x$. We fix the parameter values $A = 0.25$, $\beta = 0.25$ and $\omega = 2\pi$, hence the vector field has time period 1. The system preserves the Lebesgue measure on X. Equation (23) describes two counter-rotating gyres next to each other (the left one rotates clockwise), with the vertical boundary between the gyres oscillating periodically, see Figure 9.
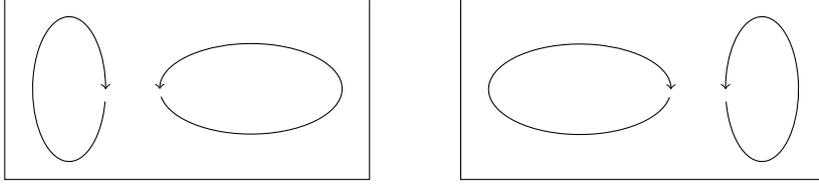
Figure 9: Sketch of the velocity field of the periodically forced double gyre flow at two different times.

We choose a uniform 50×25 grid as initial conditions for the floaters at time $t = 0$; i.e., $I = 1250$. We sample the trajectories of these floaters at times $t_k = k\tau$, $k = 0, 1, \ldots, K$, where $K = 100$ and $\tau = 0.2$. That means, the length of trajectories in consideration is 20 times the period of the forcing.

Employing our large deviations based distance computations on this data set using Algorithm 1 and $\alpha = 1/2$, we get the one-way costs $\nu_K(i \to j)$, $i, j = 1, \ldots, I$, from which we compute $\nu_K^{\text{cross}}(i, j)$ and $\nu_K^{\text{meet}}(i, j)$.

As a first simple analysis, we can order the points by their semidistances to the center of one gyre, see Figure 10. Here and in the following, the rates and semidistances will be always given in units $1/\tau$. On a log-log scale, the slope $1/2$ indicates that most trajectories in the gyre are approximately concentric circular regions around the center.[7] Since the semidistances grow linearly in the Euclidean distance inside the gyre, we see that we are in the low-resolution regime discussed in Section 3.7. Analogously to Figure 8, we can again (vaguely) distinguish three regions: a steep (quadratic) region, a flat region, and another steep (quadratic) region. As before, the flat region is typically strongly mixing, and the steep regions are static. We shall make this distinction more precise in the next sections.

4.3. Cornerstones

To start the analysis of the state space under a semidistance d, we randomly choose a trajectory, represented by a label $c_0 \in \{1, \ldots, I\}$, and compute the trajectory furthest from it, i.e, we set

$$c_1 = \underset{i=1,\ldots,I}{\arg\max} \, d(i, c_0).$$

---

[7]As on a regular grid there are $\mathcal{O}(\delta^2)$ points not further than Euclidean distance δ from a reference point, the r-th closest point to the reference point has distance $\mathcal{O}(r^{1/2})$.
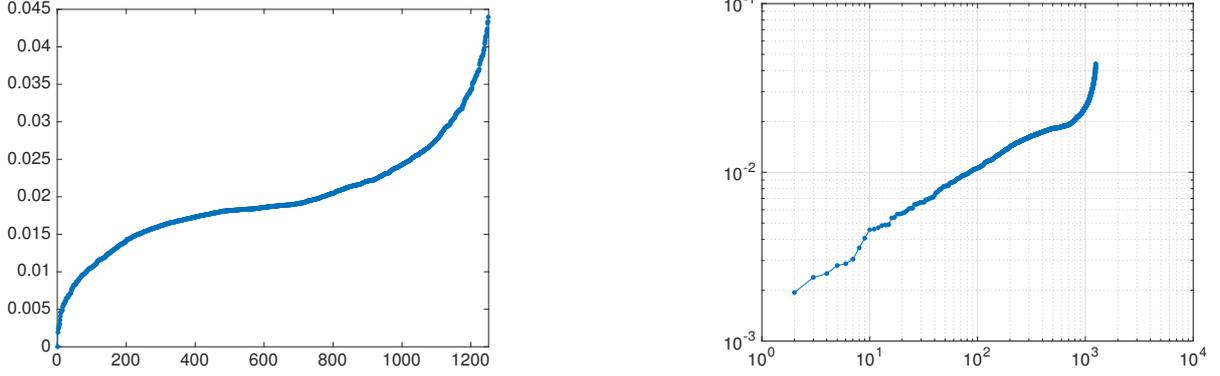
Figure 10: Left: $\nu_K^{\mathrm{meet}}(c_1, \cdot)$ sorted in ascending order. Right: the same as left, on a log-log scale.

To find a set of points that "spans" the state space, we identify successively further trajectories that are far away from *all* the other already identified "cornerstones" $\{c_q\}_{q=1,\ldots,Q}$, as in [46]:

$$c_{Q+1} = \underset{i=1,\ldots,I}{\arg \max} \ \underset{q=1,\ldots,Q}{\min} \ d(i, c_q). \tag{24}$$

Observe that in this optimalisation problem we ignore the first, randomly chosen trajectory c_0; hence the set of cornerstones $\{c_q\}_{q=1,\ldots,Q}$ will be less dependent on this randomness. Moreover, even if the first trajectory c_0 would represent a coherent set, the algorithm will eventually provide a new cornerstone in that set, which lies closer to the semidistance center of that set.

For the double gyre and the meeting distance, we identified three cornerstones. The objective function of the maximisation problem (24) is plotted in Figure 11; this yields a similar but more detailed picture as Figure 10. Note that the chaotic, well mixed transition region appears as flat region in these distance graphs, and the gyres appear as steep regions towards the maxima of the respective graphs. That the chaotic region is well mixed, and has no stratification (invariant rings as the gyres), can be seem from its flat behavior towards its maximum. The forth cornerstone is part of a gyre, it starts to stratify it. Nevertheless, its distance to the other corners is much smaller.
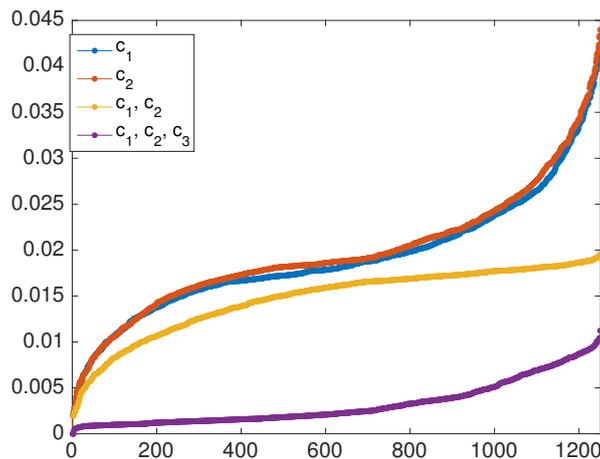
Figure 11: The objective functions of the maximisation problem in (24), sorted in ascending order (yellow and purple). Blue and red: $\nu_K^{\mathrm{meet}}(c_i, \cdot)$, $i = 1, 2$.

19

To get a first glimpse of the separate regions in the state space we have plotted the semidistances from each cornerstone in Figure 12, both at the initial and final times. Note that since we work with trajectory labels rather than physical positions, the semidistances are invariant in time, whereas the physical positions of the floaters change over time. From these figures, one can approximately identify the two static (gyre) regions, being very close and very far from c_1 and c_2 respectively, and the chaotic transition region in between, having approximately constant distance from c_1 and c_2, cf [24, Figure 1].
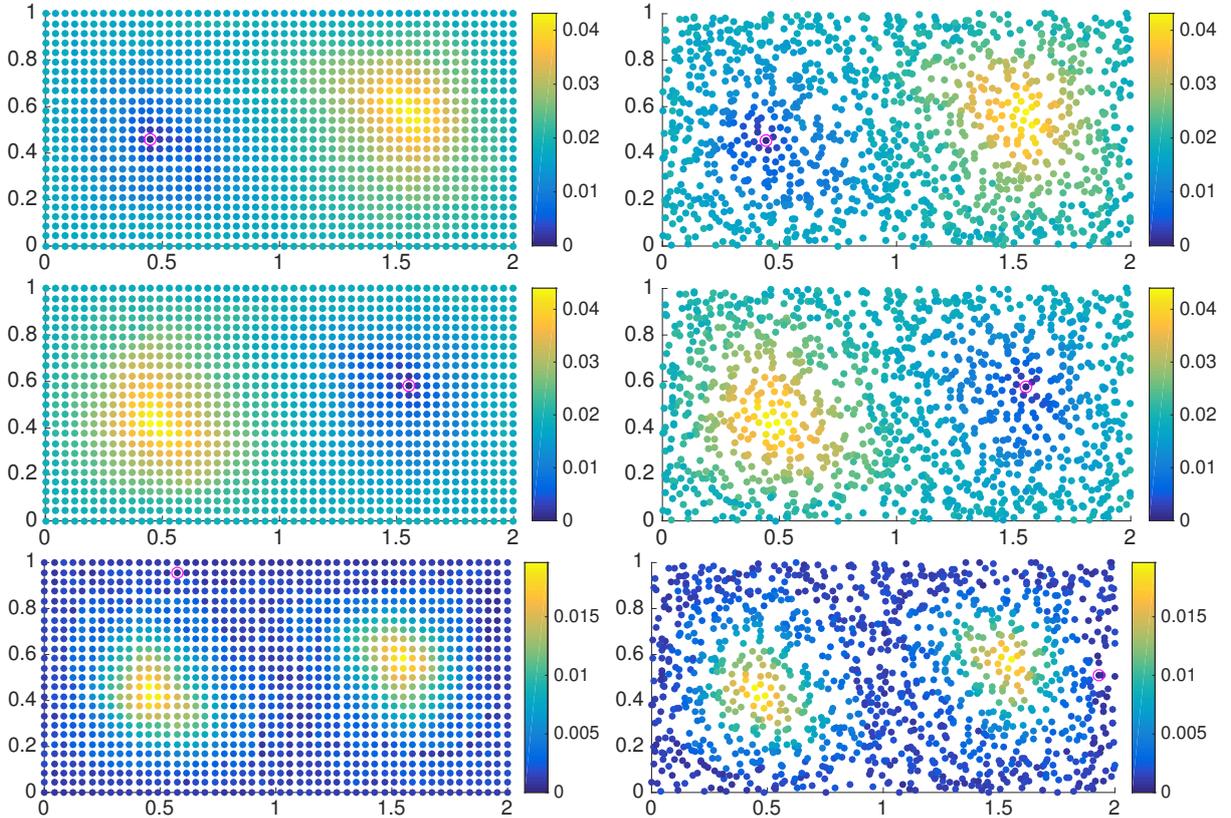
Figure 12: Distances ν_K^{meet} from trajectory c_1 (top), c_2 (middle) and c_3 (bottom), marked by the magenta circle, at initial (left) and final times (right). The semidistances are given in units $1/\tau$.

4.4. Number of cornerstones

How to determine the number of cornerstones that should be used? Is there an optimal number, or is it up to our liking? In the case of the double gyre, as noted above, a fourth cornerstone would be part of one of the gyres, and assigning affiliations would thus split one gyre into two sets. If the gyres would consist of a continuum of periodic orbits, then we could proceed and split them this way into as many rings as we like. The same situation in an idealized framework appears in Section 4.1 for the static regions: since they are static, arbitrary subsets are perfectly coherent (even invariant in this case).

A good place to stop searching for further cornerstones would be when they would start to subdivide "maximal coherent" sets, as the gyres in the double-gyre example, or the static sets in the second, and the invariant sets in the first example of Section 4.1. To this end we make

an idealized assumption: *Coherent sets appear as for the second example in Section 4.1, i.e., multiple static regions divided by one mixing region.* Here, "static" is meant in the sense that the mutual distances between points in the set barely change. Such an assumption was also utilized in [27].

Note that if there are $C \geq 2$ coherent sets, the first C corner stones are going to be in them, one in each. This is due to the fact that to move from the center of one static region to another, the shortest path in (20) needs to move out of one set, travel in the transition region to the other set, and move to its center, hence maximizing the minimal distance to all other cornerstones. After finding all static regions, the next cornerstone is to be found in the transition region, if all static regions are approximately of the same size—which we assume here. The crucial observation is, that this $(C + 1)$-st cornerstone is half as far from the other cornerstones[8], as they are from one another. In other words, $d(c_i, c_{C+1}) + d(c_j, c_{C+1}) \approx d(c_i, c_j)$, $i, j \leq C$.

To summarize, our simple check when to stop searching for cornerstones is going to be, when the value of the objective function in (24) drops by at least a factor two compared with the previous value. Observe how nicely this works in the periodically forced double gyre case: the rightmost points of the curves in Figure 11 are the optimizers, and the corresponding value of the yellow curve is less than half of the values for the first two cornerstones. This indicates to stop with three cornerstones, as they will represent both the gyres and the transition region.

4.5. Clustering and fuzzy affiliations

To get an even more crisp picture of the subdivision of the state space into regions which are far away in terms of the semidistance d, we assign to each cornerstone c_1, c_2, c_3 the trajectories that are closer to them than to the two other cornerstones, respectively. For the periodically forced double gyre and the meeting distance this is shown in Figure 13.
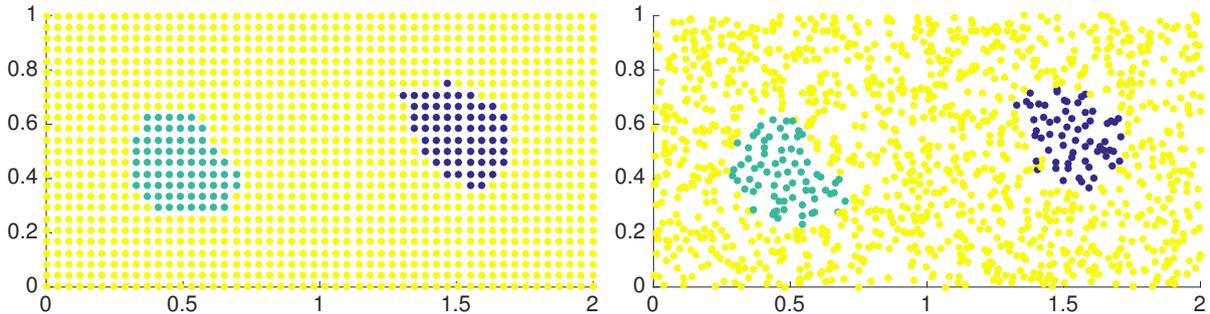
Figure 13: The trajectories closest in terms of ν_K^{meet} to one of the cornerstones than to the others. Left: initial time, right: final time.

Comparing this picture with the typical trajectories of the time-1 Poincaré map of the system (again, see [24, Figure 1]), it appears that the gyre regions in our figure are smaller. This is due to the nature of the transport distance at hand: the gyres are partly made up of so-called "regular regions" of the Poincaré map, meaning that typical trajectories move on periodic orbits that are approximately concentric circular lines. Transport between these trajectories is only possible through diffusion, and the price one has to pay for this transport in radial direction is reflected by the rate function (recall, this is what we model by the static regions in Section 4.1).

---

[8]Here we assume that we are in the coarse spatial resolution case, where the semidistances scale linearly and not quadratically, cf Section 3.7. Otherwise, the drop in the distance is more than a factor two (towards factor four).

The cost to get from the center of the gyre (the cornerstones c_1 and c_2) to a regular trajectory in the same gyre is proportional to the "radial distance" between them (compare with the static part of the second example in Section 4.1).

This is underlined in Figure 10. If we were to take a regular grid of points, and plot their Euclidean distances from a fixed point x ranked in ascending order, the distances would follow an $r^{1/2}$ asymptotics, where r denotes the rank. This is because there are $\mathcal{O}(\delta)$ points δ-far from x, or, equivalently, there are $\mathcal{O}(\delta^2)$ points closer to x than any point that has distance δ to x, and thus $\delta^2 \sim r$. This shows up as a slope $1/2$ on the double-logarithmic plot Figure 10 (right) for the close-by points, hence the distance around the cornerstone c_1 grows proportionally to the radius, indicating concentric periodic trajectories. The same holds for c_2 (now shown). This behavior is not characteristic for the well-mixed transition region, because the dynamics (eventually) brings any two trajectories close to each other. This effect is most prominent if the time frame of consideration grows infinitely large, and on our finite time horizon it appears as a flattening of the curve. This brings us back to why the blue and green regions in Figure 13 are smaller than gyres in the Poincaré map. The answer is simply, because the outer periodic orbits are closer to the transition region than to the center of the gyre, hence also closer to the cornerstone c_3 that is in the transition region, because the points in the transition region have very small distance from one another.

Instead of a hard clustering we can assign the trajectories to the cornerstones by *fuzzy affiliations* $q_{c_i}(\cdot)$, to obtain more refined information on coherence. For instance, let $m > 1$, and minimizing the affiliation-weighted penalty function

$$\sum_{j=1}^{I} \sum_{i=1}^{\ell} q_{c_i}(j)^m d(c_i, j)^2$$

subject to the constraints $0 \leq q_{c_i}$ for $i = 1, \ldots, \ell$ and $\sum_{i=1}^{\ell} q_{c_i}(j) = 1$ for every $j = 1, \ldots, I$, yields

$$q_{c_i}(j) = \frac{1}{\sum_{k=1}^{\ell} \left(\frac{d(c_i, j)}{d(c_k, j)} \right)^{\frac{2}{m-1}}} . \tag{25}$$

This is the affiliation function in the *fuzzy c-means algorithm* [6], giving $q_{c_i}(j) = 1 \Leftrightarrow d(c_i, j) = 0$, i.e., affiliation is maximal if the distance is minimal. Further, the parameter m controls the fuzziness of the clustering: large m gives soft clusters, while m approaching 1 gives more and more "crisp" clusters as the affiliations converge either to 0 or to 1 [7]. The resulting affiliations (indicated at initial time) for $m = 2$ are shown in Figure 14. For m close to 1 we obtain affiliations very similar to the hard clusters in Figure 13.
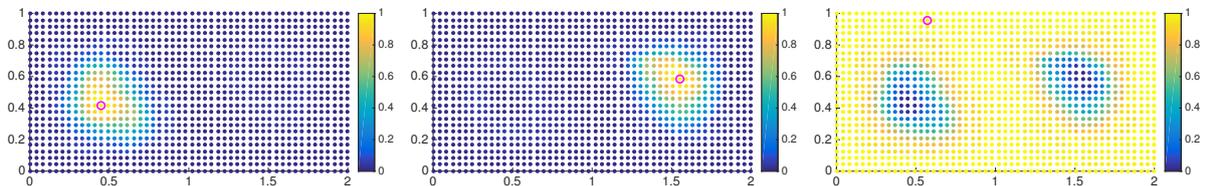
Figure 14: Fuzzy affiliations $q_{c_i}(\cdot)$ of the trajectories to the three cornerstones, c_1, c_2, c_3 (from left to right) for fuzziness exponent $m = 2$, shown at initial time.

5. Numerical results

After having introduced our analysis tools for the double gyre system, we consider here two other well-analyzed test cases: the perturbed Bickley Jet and the rotating (transitory) double gyre. In turns out that the choice between the two semidistances ν_K^{cross} or ν_K^{meet} has only marginal influence on the results. In this section we shall mostly work with the cross semidistance.

5.1. The Bickley Jet
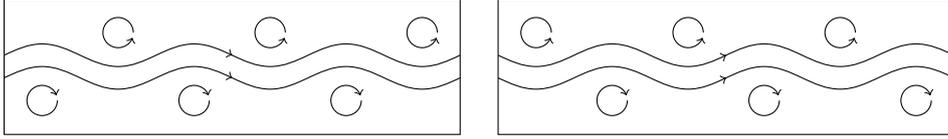
Figure 15: Sketch of the Bickley Jet flow field at two different times. The flow pattern travels from left to right on the horizontally periodic domain.

We consider a perturbed Bickley Jet as described in [47]. This is an idealised zonal jet approximation in a band around a fixed latitude, assuming incompressibility, on which three travelling Rossby waves are superimposed, see Figure 15. The dynamics is given by $\dot{x}_t = v(t, x_t)$ with $v(t, x) = (-\frac{\partial \Psi}{\partial x_2}, \frac{\partial \Psi}{\partial x_1})$ and stream function

$$\Psi(t, x_1, x_2) = -U_0 L \tanh\left(x_2/L\right) + U_0 L \operatorname{sech}^2\left(x_2/L\right) \sum_{n=1}^{3} A_n \cos\left(k_n\left(x_1 - c_n t\right)\right).$$

The constants are chosen as in [47, Section 4]. In particular, we set $k_n = 2n/r_e$ with $r_e = 6.371$, $U_0 = 5.414$, and $L = 1.77$. The phase speeds c_n of the Rossby waves are $c_1 = 0.1446 U_0$, $c_2 = 0.205 U_0$, $c_3 = 0.461 U_0$, their amplitudes $A_1 = 0.0075$, $A_2 = 0.15$, and $A_3 = 0.3$, as in [27]. The system is considered on a state space $X = [0, \pi r_e] \times [-3, 3]$ which is periodic in the horizontal x_1 coordinate.

We choose a uniform 60×18 grid as initial conditions for the floaters at time $t = 0$; i.e., $I = 1080$. We sample the trajectories of these floaters at times $t_k = k\tau$, $k = 0, 1, \ldots, K$, where $K = 80$ and $\tau = 0.5$. In this time interval, typical trajectories cross the cylindrical state space horizontally 4-5 times, trajectories in the jet core (the wavy structure in Figure 15) up to 9 times.

Employing our large deviations based distance computations on this data set using Algorithm 1 and $\alpha = 1/2$, we get the rates $\nu_K(i{\to}j)$, $i, j = 1, \ldots, I$. From these rates we readily obtain the $\nu_K^{\text{cross}}(i, j)$ via

$$\nu_K^{\text{cross}}(i, j) = \nu_K(i{\to}j) + \nu_K(j{\to}i).$$

We repeat the cornerstone finding analysis from the previous section. The optimal values of the objective function in the cornerstone finding problem (24) are for 8 cornerstones, in order:

$$2.06, \ 3.21, \ 2.45, \ 2.33, \ 2.30, \ 2.14, \ 1.42, \ 0.70.$$

Recall, that the first value is with respect to a random cornerstone c_0 that we discard. These numerical values with our previous analysis shed light on the topological structure of the state space with respect transport and mixing. Note, that our assumption from Section 4.4, that all coherent sets are divided by *one* mixing region, is not satisfied: the jet core is a coherent

23

set itself, dividing two mixing regions (below and above it), each containing 3 further coherent sets (the gyres). Thus, c_1 and c_2 have maximal distance ($\nu_K^{\mathrm{cross}}(c_1, c_2) = 3.21$), because the random walker needs to cross the jet core. Every further cornerstone $c_3, \ldots, c_6$ can be reached from either c_1 or c_2 through one of the mixing regions, and thus have a very similar cost. The deviation of these costs, $2.14 - 2.45$, shows that we did not reach the state of full mixing on the chosen time interval.

Now, the seventh cornerstone lies in the jet core, which has to be crossed if traveling between cornerstones that are below and above it, respectively. The corresponding cost (1.42) is a bit larger than half of the previous cost, because c_7 does not lie on the shortest path between cornerstones below and above the jet core. Intuitively, the "center line" of the jet core should be equally far from all cornerstones $c_1, \ldots, c_6$, if the time interval is large enough such that the regions around the gyres are truly mixing. Since it is not, there are points on the boundary of the jet core which are easier to reach from them, and thus easier to cross there. The cornerstone c_7 represents the position where it is the hardest to cross. The eighth cornerstone has truly half the semidistance to the closest one than c_7, and lies in one of the mixing regions.

We show our results for seven cornerstones[9]. The semidistances are shown in Figure 16. The
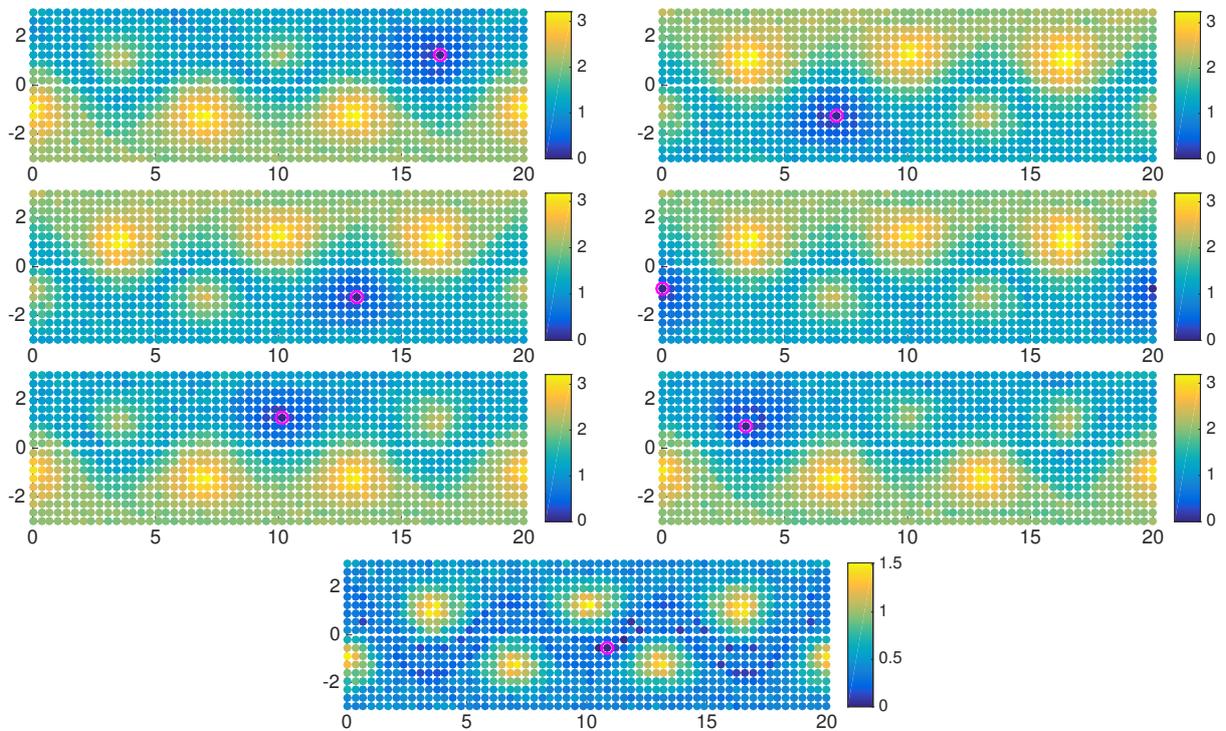
Figure 16: Row-wise from top left to bottom: the identified corner stores c_i, $i = 1, \ldots, 7$, (magenta circles) and their distances $\nu^{\mathrm{cross}}(c_i, \cdot)$ to the other trajectories, at initial time. The cornerstones are located in the six gyres and the central jet region. The distances are given in units $1/\tau$.

corresponding fuzzy affiliations from (25) for $m = 1.1$ are shown in Figure 17. They show a very crisp distinction of the six gyres from the rest of the state space. The bottom right figure shows the affiliation $q_{c_7}(\cdot)$ for $m = 1.9$, which suggests that the region around the gyres

---

[9]If we include additional cornerstones, results tend to deteriorate due to the low resolution and because the chosen time interval is not giving full mixing.

could still be partitioned into coherent sets itself: the jet core appears more strongly affiliated to this cornerstone than the other trajectories. It is not surprising that we could not see this for $m = 1.1$, since the closer m is to 1, the more "crisp" the affiliation function is forced to be, and the mixing region is more easily reached from the thin jet core than from the gyres.
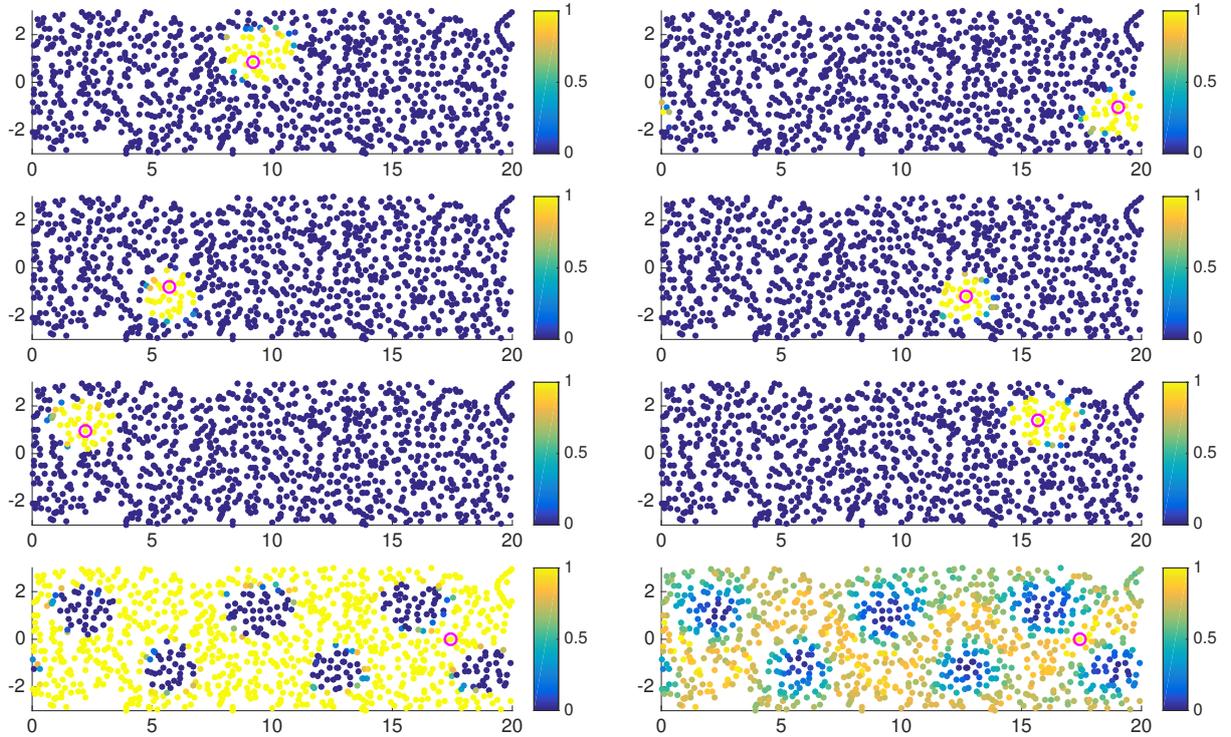
Figure 17: Row-wise from top left to bottom: the fuzzy affiliations (25) of the trajecories at time $t = 5$ to the cornerstones $c_1, \ldots, c_7$, respectively (magenta circles). Bottom right: affiliation $q_{c_7}(\cdot)$ for $m = 1.9$, which suggests that the 7th coherent region could contain a coherent set itself: the jet core.

5.2. The rotating double gyre
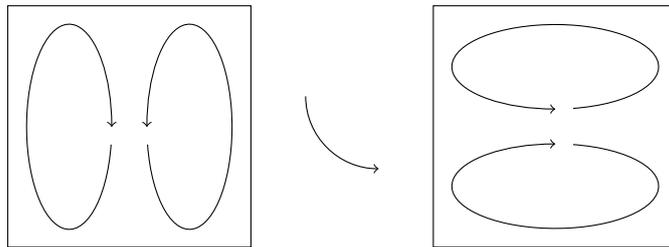
Figure 18: Sketch of the flow field of the rotating double gyre at initial (left) and final (right) times.

Let us consider a prototype for a system, where transport is considered only on a limited time interval. The rotating double gyre system [40] is given by the stream function $\psi(t, x, y) = (1 - s(t))\psi_P(x, y) + s(t)\psi_F(x, y)$, with $s(t) = t^2(3 - 2t)$, $\psi_P(x, y) = \sin(2\pi x)\sin(\pi y)$, and $\psi_F(x, y) =$

25

$\sin(\pi x)\sin(2\pi y)$, and is considered on the state space $X = [0,1]^2$ and time interval $t \in [0,1]$. The two gyres, which initially occupy the left and right halves of the unit square, turn during this time by $\pi/2$ to occupy the top and bottom halves at final time, see Figure 18.

We choose a uniform 30×30 grid as initial conditions for the floaters at time $t = 0$; i.e., $I = 900$. We sample the trajectories of these floaters at times $t_k = k\tau$, $k = 0, 1, \ldots, K$, where $K = 100$ and $\tau = 0.01$. We employ the cross-semidistance, and start our cornerstone search. The first three values of the optimization problem (24) are

$$0.0274, \ 0.0474, \ 0.0262.$$

We identify the significant drop after two corner stones, hence we expect two coherent sets with one mixing region dividing them. The drop in the distance is by a factor 0.55, which is not below one half, the reason for this being again that the time interval of consideration is not sufficient for perfect mixing of the transition region. The semidistances from the three identified
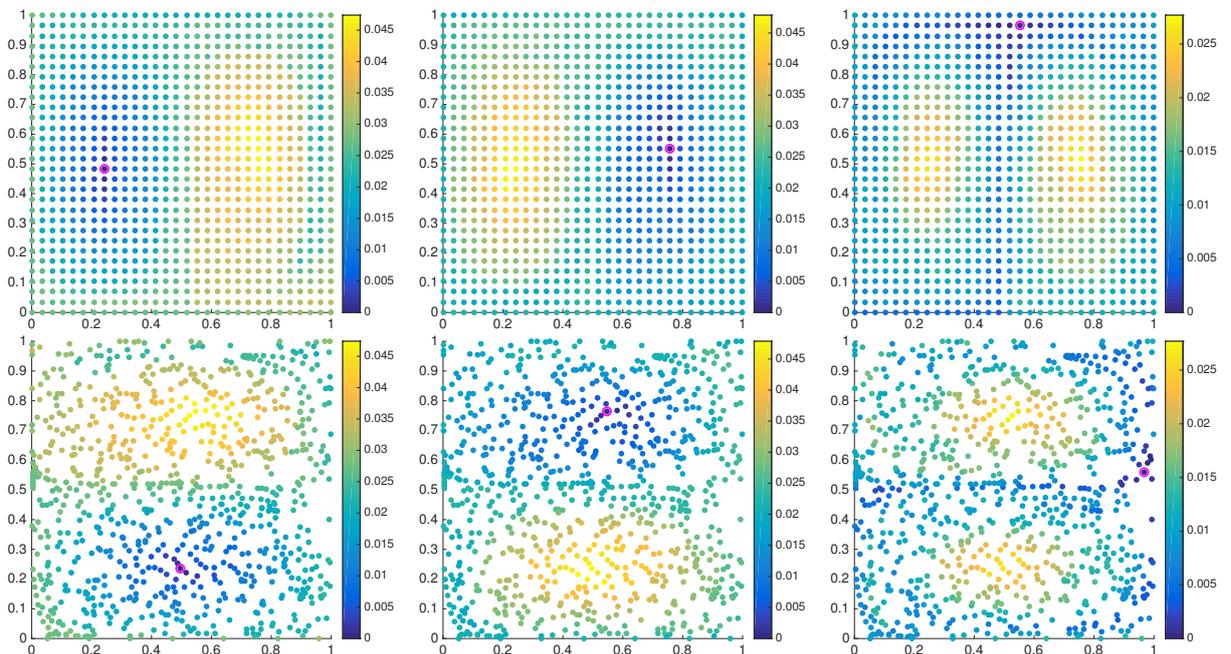
Figure 19: From left to right: the identified corner stores c_i, $i = 1, \ldots, 3$, (magenta circles) and their distances $\nu^{\mathrm{cross}}(c_i, \cdot)$ to the other trajectories, at initial time (top) and final time (bottom). The distances are given in units $1/\tau$.

cores and the affiliations to these cores for exponent $m = 1.2$ are shown in figures 19 and 20, respectively. Although both ν^{cross} and ν^{meet} have been shown to be able to detect coherent sets, we demonstrate their different nature by showing the shortest paths in the respective distance in Figure 21.

Finally, we demonstrate the approach for a scattered set of sparse data points, taking 400 initial points randomly distributed in X, and repeating the analysis for their trajectories. We show the resulting fuzzy affiliations in Figure 22.
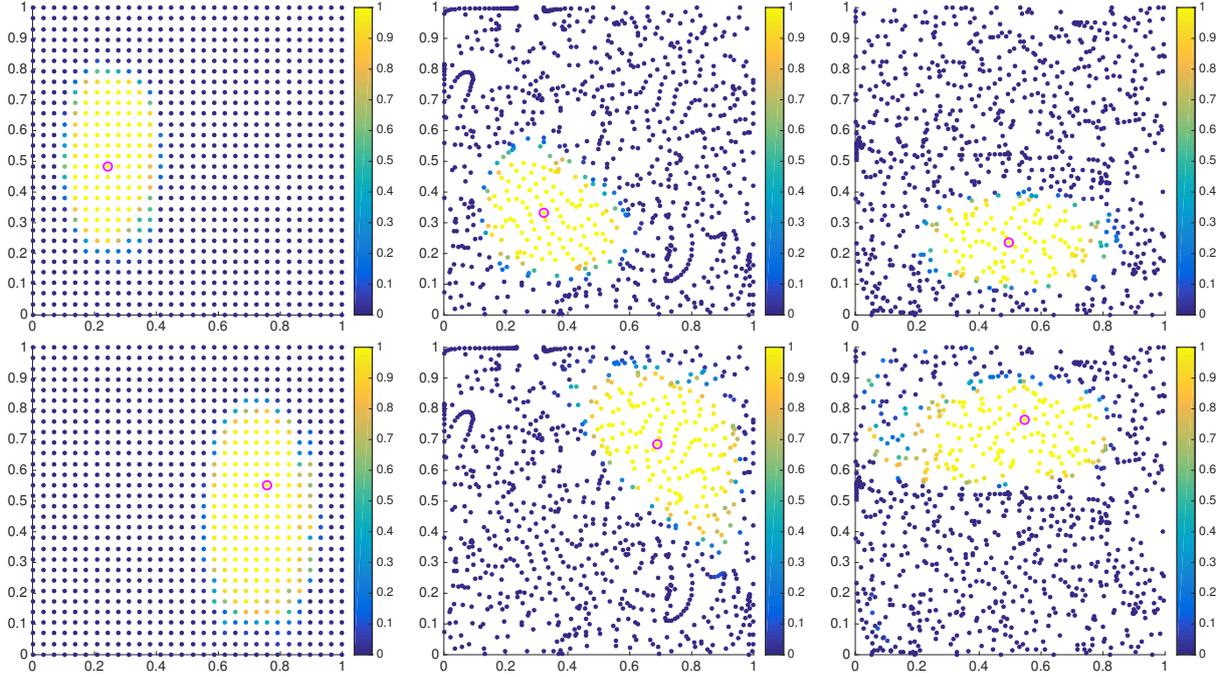
Figure 20: The fuzzy affiliations computed with $m = 1.2$ to the cornerstones c_1 (top) and c_2 (bottom), at times $t = 0, 0.5, 1$, from left to right, respectively.

6. Discussion and outlook

6.1. The dynamic Laplacian

Froyland [18] has introduced the *dynamic Laplacian* as a transport-related tool to find coherent sets. Similarly to our approach, it makes use of a small random perturbation of size ε, then ε is driven to zero.

Numerical methods so far discretize directly the dynamic Laplacian [20, 21]. In light of our analysis, which can be used both ways (derive the large-deviations principle in continuous space, then discretize it to finite trajectories, cf Section 3.6, or discretize the dynamics to finite trajectories, then derive the large-deviations principle on them, cf Section 3.3), we ask whether there is a discrete dynamic Laplacian that can be derived from a discretization of the perturbed dynamics?

Mimicking the construction in [18] and sketching the idea while skipping details, one should construct a discrete, ε-dependent transfer operator $T_\varepsilon \in \mathbb{R}^{I \times I}$, that represents transition probabilities of a forward-backward process, then obtain a discrete dynamic Laplace operator $L_{\mathrm{dyn}} := \frac{d}{d\varepsilon}\big|_{\varepsilon=0} T_\varepsilon$. A discrete transfer operator T_ε that is a consistent approximation of the continuous dynamics can be obtained by a construction as in Section 3.2, by using the transition probabilities (17). Technical details aside, we see that the probabilities are linear combinations of terms of the form $\mathrm{e}^{-\Delta x/\varepsilon}$, where Δx here is a formal distance term that appears in the formulas. Differentiation with respect to ε immediately yields that all off-diagonal entries (basically, where $\Delta x > 0$) of L_{dyn} are zero, in fact the matrix is the identity.

Thus, this approach of discretizing the dynamics first, and then factoring out the ε-small stochastic perturbation does not give a dynamically meaningful result. In analytic terms the very same problem occurred in a different attempt to introduce a discrete dynamic Laplacian
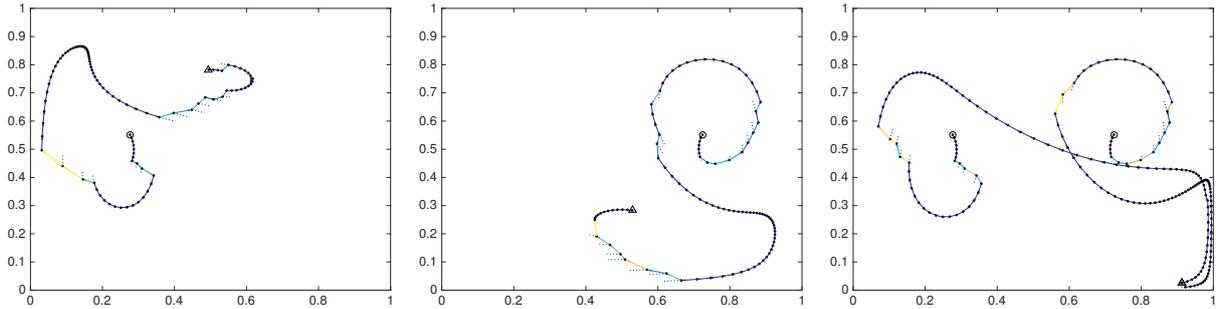
Figure 21: Shortest paths from cornerstone c_1 to c_2 (left), from c_2 to c_1 (middle), and the meeting paths with the shortest joint length (right). The brighter the color of a path segment, the larger the cost of that transition. For those segments for which the path crosses from one trajectory to another we show with a dashed segment how the former trajectory would have continued. The starting point of the path is indicated by a circle, and the end point by a triangle.

from a discrete transfer operator, see [5, Section IV]. In general, it would be desirable to understand when and how can the "first discretize, then factor out ε" methods work, such that they can complement the methods that directly discretize the (continuous) dynamic Laplace operator.

6.2. Other distance measures

The time-dependent shortest path problem used to compute our semidistances is computationally demanding in our current algorithmic realization, which theoretically limits the number of trajectories that can be handled. Moreover, they do not satisfy the triangle inequality, hence they are not a metric. Although numerical efficiency is not the main focus of this paper, and we demonstrated the usefulness of our semidistances in unraveling the underlying dynamical structure of the example systems, a more cheaply computable metric would enhance the utility and significance of the analysis methods presented here.

Ultimately, one would like to understand the intrinsic, possibly low-dimensional geometric organization of the state space with respect to transport (or mixing), as pioneered in [5]. Employing proper metrics would allow, e.g., the usage of low-dimensional embedding techniques, such as multidimensional scaling, to represent and better understand this geometric organization. One canonical candidate would be the metric structure related to the dynamic Laplacian, considered in [33]. This will be subject of future studies.

To summarize, although other distance measures could be used, we showed that the semidistances we derived in this paper are natural in the sense of transport in the vanishing diffusion setting.

Acknowledgments
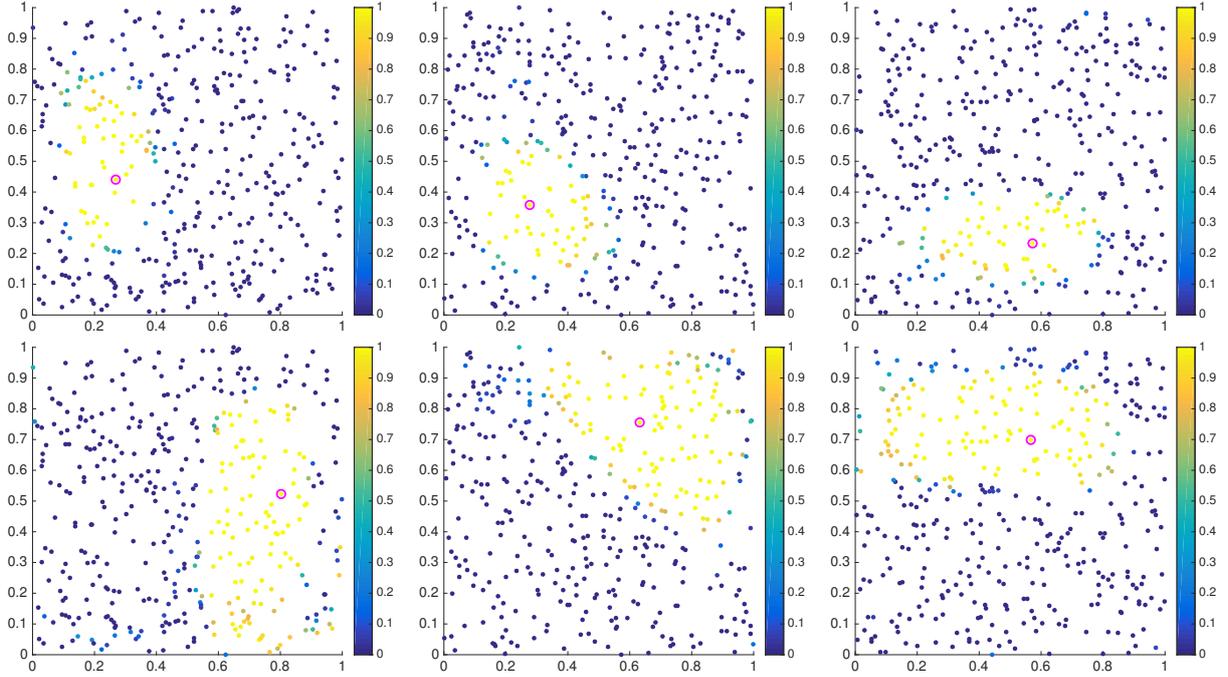
Figure 22: The rotating double gyre with randomly chosen 400 initial points. The fuzzy affiliations computed with $m = 1.2$ to the cornerstones c_1 (top) and c_2 (bottom), at times $t = 0, 0.5, 1$, from left to right, respectively.

A. Large deviations of the forward-backward conditions

In this appendix we explore the conditions (3) in the large-deviations regime. The argument is based on the *Laplace Principle*, which states that for any measure ρ and function f:

$$\lim_{\varepsilon \to 0} -\varepsilon \log \int e^{-\frac{1}{\varepsilon} f(x)} \rho(dx) = \inf_{x \in \operatorname{supp} \rho} f(x). \tag{26}$$

As in (9),

$$-\varepsilon \log \mathbb{P}\big(\boldsymbol{x}_T^{(\varepsilon)} \asymp y \mid \boldsymbol{x}_0^{(\varepsilon)} = x\big) \xrightarrow[\varepsilon \to 0]{} \inf_{x_{(\cdot)} \,:\, x_0 = x, x_T = y} \frac{1}{2} \int_0^T |\dot{x}_t - v(t, x_t)|^2 \, dt =: \lambda_T(x \to y), \tag{27}$$

where, contrary to (9), the symbol y now denotes a position at time T, that is, $\mu_T(x \to \tilde{x}) = \lambda_T(x \to \phi_{0,T}[\tilde{x}]) = \lambda_T(x \to y)$.

Fix an ε-independent initial probability measure $\rho_0(dx) = \mathbb{P}(\boldsymbol{x}_0 \in dx)$. For the large deviations of the forward condition in (3), it follows from the Laplace principle that

$$\mathcal{J}_T^{\mathrm{fw}}(B|A) := \lim_{\varepsilon \to 0} -\varepsilon \log \mathbb{P}(\boldsymbol{x}_T^{(\varepsilon)} \in B \mid \boldsymbol{x}_0 \in A)$$

$$= \lim_{\varepsilon \to 0} -\varepsilon \log \int_B \int_A \mathbb{P}\big(\boldsymbol{x}_T^{(\varepsilon)} \in dy \mid \boldsymbol{x}_0 = x\big) \rho_0(dx)$$

$$\overset{(27)}{=} \lim_{\varepsilon \to 0} -\varepsilon \log \int_B \int_A e^{-\frac{1}{\varepsilon} \lambda_T(x \to y)} \rho_0(dx)$$

$$\overset{(26)}{=} \inf_{y \in B} \inf_{x \in A \cap \operatorname{supp} \rho_0} \lambda_T(x \to y).$$

Observe that since the initial distribution ρ_0 is independent of ε, it only appears in the large deviations through its support $\operatorname{supp} \rho_0$.

The large deviations of the backward conditions in (3) can be calculated analogously, but now the conditioning does depend on ε. By Bayes' rule,

$$
\begin{aligned}
\mathcal{J}_T^{\text{bw}}(A|B) &:= \lim_{\varepsilon \to 0} -\varepsilon \log \mathbb{P}(\boldsymbol{x}_0 \in A \mid \boldsymbol{x}_T^{(\varepsilon)} \in B) \\
&= \lim_{\varepsilon \to 0} -\varepsilon \log \mathbb{P}(\boldsymbol{x}_T^{(\varepsilon)} \in B \mid \boldsymbol{x}_0 \in A) \frac{\mathbb{P}(\boldsymbol{x}_0 \in A)}{\mathbb{P}(\boldsymbol{x}_T^{(\varepsilon)} \in B)} \\
&= \lim_{\varepsilon \to 0} -\varepsilon \log \int_B \int_A \mathbb{P}(\boldsymbol{x}_T^{(\varepsilon)} \in dy \mid \boldsymbol{x}_0 = x) \rho_0(dx) \\
&\quad + \varepsilon \log \int_B \int \mathbb{P}(\boldsymbol{x}_T^{(\varepsilon)} \in dy \mid \boldsymbol{x}_0 = x) \rho_0(dx) - \varepsilon \log \rho_0(A) \\
&\overset{(26,27)}{=} \inf_{y \in B} \inf_{x \in A \cap \operatorname{supp} \rho_0} \lambda_T(x \to y) - \inf_{y \in B} \inf_{x \in \operatorname{supp} \rho_0} \lambda_T(x \to y) \\
&= \mathcal{J}_T^{\text{fw}}(B|A) - \mathcal{J}_T^{\text{fw}}(B|X).
\end{aligned}
$$

If we assume that that there is at least one admissible path $x_{(\cdot)}$ that starts in $\operatorname{supp} \rho_0$ and ends in B, then in fact $\mathcal{J}_T^{\text{fw}}(B|X) = 0$, and so $\mathcal{J}_T^{\text{fw}}(B|A) = \mathcal{J}_T^{\text{bw}}(A|B)$.

These calculations have two important implications. First, observe that the while the forward and backward probabilities $\mathbb{P}(\boldsymbol{x}_T^{(\varepsilon)} \in B \mid \boldsymbol{x}_0 \in A)$ and $\mathbb{P}(\boldsymbol{x}_0 \in A \mid \boldsymbol{x}_T^{(\varepsilon)} \in B)$ are not equal in general, the forward and backward rate functions are. The same argument even holds if we shrink the sets A ad B down to single points x and y; in that case we obtain for the "backward rates" that

$$
\lambda_T(x \leftarrow y) := \lim_{\varepsilon \to 0} -\varepsilon \log \mathbb{P}(\boldsymbol{x}_0 \asymp x \mid \boldsymbol{x}_T^{(\varepsilon)} = y) = \lambda_T(x \to y).
$$

Apparently, in the large-deviation scaling it does not matter whether we consider the forward or the backward process. Since the forward condition $\mathcal{J}_T^{\text{fw}}(B|A) \approx 0$ in itself does not hold enough information to characterise coherence and the backward condition $\mathcal{J}_T^{\text{bw}}(A|B) \approx 0$ does not add information, these conditions are not helpful to characterise coherence.

Secondly, we see that $\mathcal{J}_T^{\text{fw}}(B|A) = 0$ as soon as $\phi_{0,T}[A \cap \operatorname{supp} \rho_0] \cap B \neq \emptyset$. Naturally, there are many such pairs A, B, and the set function $\mathcal{J}_T^{\text{fw}}$ does not give any quantitative information about which pairs are more coherent than others. Because of this, the large deviations of the forward and backward conditions (3) are even less usefull to identify coherent sets.

We start to gain useful information about coherence, if there are at least two coherent pairs, say A_1, B_1 and A_2, B_2. Then the rates $\mathcal{J}_T^{\text{fw}}(B_2|A_1)$ and $\mathcal{J}_T^{\text{fw}}(B_1|A_2)$ are in general large, since coherence of the respective set pairs dictate that it is very unlikely to encounter paths from pair #1 to pair #2. Using these rates as measures of farness is one idea this paper exploits.

B. Algorithm: shortest path in time-dependent graphs

Since we could not find an algorithm suited to our purpose, we describe in this appendix a solution we came up with to solve the problem of finding shortest paths in a graph with time-dependent non-negative edge weights. Transition is only possible between nodes that are connected by an edge of positive weight.

There are several solutions to the shortest path problem for *time-independent* graphs, such as *Dijkstra's algorithm* [13] or the *Floyd–Warshall algorithm* [15]. Each of them use in some sense a "monotony" argument, namely, that sub-paths of shortest paths are shortest paths

themselves. This does not hold for time-dependent graphs, because at every step that we make the environment might change completely, and the number of steps we can make is limited by the number of time instances of the graph.

We propose the following algorithm to compute shortest paths from a specific node s to all other nodes. Note that we can stay in a node for any time at zero cost. The weight of the transition $i \to j$ at time t is denoted by $w_t(i \to j)$.

Algorithm 1 Shortest distance in time-dependent graphs

1: $R_{old} = \{s\}, R_{new} = \emptyset$ (reached states at times 0 and 1)
2: $\text{dist}(s) = 0$, $\text{dist}(i) = \infty$ for $i \neq s$
3: **for** $t = 1, \ldots, T$ **do**
4:      **while** $R_{old} \neq \emptyset$ **do**
5:          $v = \arg\max_{i \in R_{old}} \text{dist}(i)$
6:          $R_{new} \leftarrow R_{new} \cup \{v\}$
7:          $R_{old} \leftarrow R_{old} \setminus \{v\}$
8:          **for** $j : w_t(v \to j) < \infty$ **do**
9:              **if** $\text{dist}(v) + w_t(v \to j) < \text{dist}(j)$ **then**
10:                 $\text{dist}(j) = \text{dist}(v) + w_t(v \to j)$
11:      $R_{old} = R_{new}, R_{new} = \emptyset$

It is important to have the max on line 5, since if we don't start the update procedure at the node which has the maximal distance, then we might erroneously cut off nodes that could still be reached from it.

Algorithm 1 can clearly be extended to keep track of the shortest path as well. The distance of a node j is updated to a smaller one, whenever there is a path through some other node v that is shorter than the previous one (line 10). Hence, the new candidate shortest path is the one leading to v and then jumping to j in the current time step. This is implemented in Algorithm 2 (line 12). Herein, $\text{path}(i \to j)$ is the shortest path from node i to node j, such that $\text{path}_t(i \to j)$ is the node the walker resides in at time $t = 0, 1, \ldots, T$ while going through the shortest path, and $\text{path}_0(i \to j) = i$. If there is no path from i to j, then $\text{path}(i \to j)$ is the zero vector. We use $1 : k$ to denote the index set $1, 2, \ldots, k$.

Algorithm 2 Shortest path in time-dependent graphs

1: $R_{old} = \{s\}, R_{new} = \emptyset$ (reached states at times 0 and 1)
2: $\text{dist}(s) = 0$, $\text{dist}(i) = \infty$ for $i \neq s$
3: $\text{path}(s \to j) = 0 \in \mathbb{R}^{T+1}$ for all j, $\text{path}_0(s \to s) = s$
4: **for** $t = 1, \ldots, T$ **do**
5:      **while** $R_{old} \neq \emptyset$ **do**
6:          $v = \arg\max_{i \in R_{old}} \text{dist}(i)$
7:          $R_{new} \leftarrow R_{new} \cup \{v\}$
8:          $R_{old} \leftarrow R_{old} \setminus \{v\}$
9:          **for** $j : w_t(v \to j) < \infty$ **do**
10:              **if** $\text{dist}(v) + w_t(v \to j) < \text{dist}(j)$ **then**
11:                 $\text{dist}(j) = \text{dist}(v) + w_t(v \to j)$
12:                 $\text{path}_{1:t-1}(s \to j) = \text{path}_{1:t-1}(s \to v), \text{path}_t(s \to j) = j$
13:      $\text{path}_t(s \to s) = s$
14:      $R_{old} = R_{new}, R_{new} = \emptyset$

References

[1] M. R. Allshouse and T. Peacock. Lagrangian based methods for coherent structure detection. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(9):97617, 2015.

[2] M. R. Allshouse and J.-L. Thiffeault. Detecting coherent structures using braids. *Physica D: Nonlinear Phenomena*, 241(2):95–105, 2012.

[3] S. Andres. *Diffusion processes with reflection*. PhD thesis, TU Berlin, 2009.

[4] S. Balasuriya, G. Froyland, and N. Santitissadeekorn. Absolute flux optimising curves of flows on a surface. *Journal of Mathematical Analysis and Applications*, 409(1):119–139, 2014.

[5] R. Banisch and P. Koltai. Understanding the geometry of transport: Diffusion maps for Lagrangian trajectory data unravel coherent sets. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(3):035804, 2017.

[6] J. C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Kluwer Academic Publishers, 1981.

[7] J. C. Bezdek, R. J. Hathaway, M. J. Sabin, and W. T. Tucker. Convergence theory for fuzzy c-means: counterexamples and repairs. *IEEE Transactions on Systems, Man, and Cybernetics*, 17(5):873–877, 1987.

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

[9] M. Dellnitz, G. Froyland, C. Horenkamp, K. Padberg-Gehle, and A. S. Gupta. Seasonal variability of the subpolar gyres in the Southern Ocean: a numerical investigation based on transfer operators. *Nonlinear Processes in Geophysics*, 16:655–664, 2009.

[10] M. Dellnitz and O. Junge. On the approximation of complicated dynamical behavior. *SIAM J. Numer. Anal.*, 36:491–515, 1999.

[11] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume 38. Springer, New York, NY, USA, 2nd edition, 1987.

[12] A. Denner, O. Junge, and D. Matthes. Computing coherent sets using the Fokker–Planck equation. *Journal of Computational Dynamics*, 3(2):163–177, 2016.

[13] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.

[14] A. Fabregat, I. Mezic, and A. C. Poje. Finite-time partitions for Lagrangian structure identification in Gulf Stream eddy transport. *arXiv preprint arXiv:1606.07382*, 2016.

[15] R. W. Floyd. Algorithm 97: shortest path. *Communications of the ACM*, 5(6):345, 1962.

[16] M. Freidlin and A. Wentzell. *Random Perturbations of Dynamical Systems*. Springer, Berlin, Germany, second edition, 1998.

[17] G. Froyland. An analytic framework for identifying finite-time coherent sets in time-dependent dynamical systems. *Physica D: Nonlinear Phenomena*, 250:1–19, 2013.

[18] G. Froyland. Dynamic isoperimetry and the geometry of Lagrangian coherent structures. *Nonlinearity*, 28(10):3587, 2015.

[19] G. Froyland, C. Horenkamp, V. Rossi, and E. van Sebille. Studying an Agulhas ring's long-term pathway and decay with finite-time coherent sets. *Chaos*, 25(8):083119, 2015.

[20] G. Froyland and O. Junge. On fast computation of finite-time coherent sets using radial basis functions. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(8):087409, 2015.

[21] G. Froyland and O. Junge. Robust FEM-based extraction of finite-time coherent sets using scattered, sparse, and incomplete trajectories. *arXiv preprint arXiv:1705.03640*, 2017.

[22] G. Froyland and P. Koltai. Estimating long-term behavior of periodically driven flows without trajectory integration. *Nonlinearity*, 30(5):1948, 2017.

[23] G. Froyland and E. Kwok. A dynamic Laplacian for identifying Lagrangian coherent structures on weighted Riemannian manifolds. *arXiv preprint arXiv:1610.01128*, 2016.

[24] G. Froyland and K. Padberg-Gehle. Almost-invariant and finite-time coherent sets: directionality, duration, and diffusion. In *Ergodic Theory, Open Dynamics, and Coherent Structures*, pages 171–216. Springer, 2014.

[25] G. Froyland and K. Padberg-Gehle. A rough-and-ready cluster-based approach for extracting finite-time coherent sets from sparse and incomplete trajectory data. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(8):087406, 2015.

[26] G. Froyland, N. Santitissadeekorn, and A. Monahan. Transport in time-dependent dynamical systems: Finite-time coherent sets. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(4):043116, 2010.

[27] A. Hadjighasem, D. Karrasch, H. Teramoto, and G. Haller. Spectral-clustering approach to Lagrangian vortex detection. *Physical Review E*, 93(6):063107, 2016.

[28] G. Haller. Finding finite-time invariant manifolds in two-dimensional velocity fields. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 10(1):99–108, 2000.

[29] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D*, 149(4):248–277, 2001.

[30] G. Haller and F. Beron-Vera. Coherent Lagrangian vortices: The black holes of turbulence. *Journal of Fluid Mechanics*, 731, 2013.

[31] G. Haller and F. J. Beron-Vera. Geodesic theory of transport barriers in two-dimensional flows. *Physica D: Nonlinear Phenomena*, 241(20):1680–1702, 2012.

[32] D. Karrasch. Lagrangian transport through surfaces in volume-preserving flows. *SIAM Journal on Applied Mathematics*, 76(3):1178–1190, 2016.

[33] D. Karrasch and J. Keller. A geometric heat-flow theory of Lagrangian coherent structures. arXiv Preprint 1608.05598, 2016.

[34] P. E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer-Verlag, 3 edition, 2010.

[35] P. Koltai, G. Ciccotti, and C. Schütte. On metastability and Markov state models for non-stationary molecular dynamics. *The Journal of Chemical Physics*, 145(17):174103, 2016.

[36] R. MacKay, J. Meiss, and I. Percival. Transport in Hamiltonian systems. *Physica D: Nonlinear Phenomena*, 13(1-2):55–81, 1984.

[37] J. Meiss. Symplectic maps, variational principles, and transport. *Reviews of Modern Physics*, 64(3):795, 1992.

[38] I. Mezić and A. Banaszuk. Comparison of systems with complex behavior. *Physica D: Nonlinear Phenomena*, 197(1):101–133, 2004.

[39] I. Mezić and S. Wiggins. A method for visualization of invariant sets of dynamical systems based on the ergodic partition. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 9(1):213–218, 1999.

[40] B. A. Mosovsky and J. D. Meiss. Transport in transitory dynamical systems. *SIAM Journal on Applied Dynamical Systems*, 10(1):35–65, 2011.

[41] B. Øksendal. *Stochastic Differential Equations - An Introduction with Applications*. Springer, Berlin, Germany, sixth edition, 2003.

[42] L. Onsager and S. Machlup. Fluctuations and irreversible processes. *Phys. Rev.*, 91(6):1505–1512, Sep 1953.

[43] K. Padberg, T. Hauff, F. Jenko, and O. Junge. Lagrangian structures and transport in turbulent magnetized plasmas. *New Journal of Physics*, 9:400, 2007.

[44] K. Padberg-Gehle and C. Schneide. Network-based study of Lagrangian transport and mixing. *Nonlinear Processes in Geophysics Discussions*, 2017:1–14, 2017.

[45] V. Rom-Kedar and S. Wiggins. Transport in two-dimensional maps. *Archive for Rational Mechanics and Analysis*, 109(3):239–298, 1990.

[46] S. Rüdrich, M. Sarich, and C. Schütte. Utilizing hitting times for finding metastable sets in non-reversible markov chains. To appear in Journal of Computational Dynamics. Preprint: https://opus4.kobv.de/opus4-zib/frontdoor/index/index/docId/5120, 2017.

[47] I. Rypina, M. Brown, F. Beron-Vera, H. Kocak, M. Olascoaga, and I. Udovydchenkov. On the Lagrangian dynamics of atmospheric zonal jets and the permeability of the stratospheric polar vortex. *Journal of the Atmospheric Sciences*, 64(10):3595–3610, 2007.

[48] I. I. Rypina and L. J. Pratt. Trajectory encounter volume as a diagnostic of mixing potential in fluid flows. *Nonlinear Processes in Geophysics*, 24(2):189, 2017.

[49] K. L. Schlueter-Kuck and J. O. Dabiri. Coherent structure colouring: identification of coherent structures from sparse data using graph theory. *Journal of Fluid Mechanics*, 811:468–486, 2017.

[50] E. Ser-Giacomi, V. Rossi, C. López, and E. Hernández-García. Flow networks: A characterization of geophysical fluid transport. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(3):036404, 2015.

[51] G. Strang. On the construction and comparison of difference schemes. *SIAM Journal on Numerical Analysis*, 5(3):506–517, 1968.

[52] A.-M. Treguier, O. Boebel, B. Barnier, and G. Madec. Agulhas eddy fluxes in a 1/6 degrees Atlantic model. *Deep Sea Research Part II*, 50(1):251–280, 2003.

[53] P. Walters. *An introduction to ergodic theory*, volume 79. Springer Science & Business Media, 2000.

[54] E. P. Wigner. Calculation of the rate of elementary association reactions. *Journal of Chemical Physics*, 5:720–725, 1937.

[55] M. O. Williams, I. I. Rypina, and C. W. Rowley. Identifying finite-time coherent sets from limited quantities of Lagrangian data. *Chaos*, 25(8), 2015.