Almost-invariant and finite-time coherent sets: directionality, duration, and diffusion

Gary Froyland and Kathrin Padberg-Gehle

Abstract Regions in the phase space of a dynamical system that resist mixing over a finite-time duration are known as almost-invariant sets (for autonomous dynamics) or coherent sets (for nonautonomous or time-dependent dynamics). These regions provide valuable information for transport and mixing processes; almost invariant sets mitigate transport between their interior and the rest of phase space, and coherent sets are good transporters of "mass" precisely because they move about with minimal dispersion (for example, oceanic eddies are good transporters of water that is warmer/cooler/saltier than the surrounding water). The most efficient approach to date for the identification of almost-invariant and coherent sets is via transfer operators. In this chapter we describe a unified setting for optimal almost-invariant and coherent set constructions and introduce a new coherent set construction that is suited to tracking coherent sets over several finite-time intervals. Under this unified treatment we are able to clearly explain the fundamental differences in the aims of the techniques, and describe the differences and similarities in the mathematical and numerical constructions. We explore the role of diffusion, the influence of the finite-time duration, and discuss the relationship of time-directionality with hyperbolic dynamics. All of these issues are elucidated in detailed case studies of two well-known systems.

Gary Froyland
University of New South Wales, Sydney NSW 2052, Australia.
e-mail: g.froyland@unsw.edu.au

Kathrin Padberg-Gehle
Technische Universität Dresden, D-01069 Dresden, Germany.
e-mail: kathrin.padberg@tu-dresden.de

1 Introduction

The mathematical description of transport and mixing processes in dynamical systems has been the subject of intense research over the last two decades. Relevant applications include astrodynamics, molecular dynamics, geophysical flows, and biological systems, see e.g. [32, 46, 1, 47] for discussions and reviews of transport and mixing phenomena.

Much research has focussed on the detection, approximation and analysis of the *geometrical* structures that may explain transport barriers and the underlying transport mechanisms in autonomous and nonautonomous dynamical systems, see e.g. [39, 40, 46, 21, 22, 43, 47, 23] and references therein.

Probabilistic approaches provide a macroscopic view of the dynamics, studying the global evolution of densities. These techniques can be used to detect regions in phase space that remain *coherent* under the action of the dynamical system. In autonomous systems such regions are termed *almost-invariant* or metastable sets, introduced in the past 15 years in the context of dynamical systems [8, 9] and time-symmetric Markov processes [10, 42, 24]. These concepts rely on the Perron-Frobenius (or transfer) operator, a linear Markov operator. Subdominant eigenfunctions of this operator are heuristically used to estimate almost-invariant sets; see [14, 11, 12] for further extensions to this approach. Studies on the connections between eigenmodes of evolution operators and slow mixing in fluid flow can be found in [31, 37, 38, 44], numerical investigations of stochastically perturbed transfer operators include [4, 2], and a related series of work beginning with [33, 34] decomposes the phase space into ergodic components.

Building on the transfer operator framework, a mathematical definition and the corresponding numerical treatment of coherent sets in nonautonomous systems has only recently been proposed in the time-asymptotic [16, 17] and finite-time [19, 13] settings. The mathematical concepts introduced in [11] for the autonomous case and [19] for the finite-time case deal with finite-state Markov chains, i.e. discretized transfer operators, and are thus purely finitary. Froyland [13] recently proposed a transfer operator-based framework for identifying finite-time coherent sets, generalising the matrix-based approach of [19]. A transfer operator is defined via an appropriately chosen stochastic kernel and is shown to be a compact L^2-operator with suitable spectral properties. Of particular focus in [13] is a Perron-Frobenius operator pre- and post-composed with ε-diffusion; the influence of noise on the spectrum and singular vectors is studied. While [13] focusses on developing an analytic framework for the matrix setup used in [19], we will show in this paper that the construction of [13] is rather generally applicable. We will adapt it to verify the assumptions underlying the finitary almost-invariant sets framework in [11], and to make further theoretical extensions to finite-time coherent sets concepts. We also demonstrate how the duration of the finite-time interval under consideration and the strength of diffusion influence the structure of the regions in phase space that most resist mixing. Furthermore, we show how different concepts give rise to, or suppress, the impact of time-directionality on almost-invariant sets and finite-time coherent sets in strongly hyperbolic conditions.

The paper is organized as follows. Section 2 begins by giving brief background information on the Perron-Frobenius operator and is followed by three introductory sections that describe the three dynamical settings considered in this chapter and set up the associated problems of finding optimal almost-invariant and coherent sets. Section 3 introduces the two key tools we will use: firstly, a variant of the Perron-Frobenius operator developed in [13] that is adapted to the finite-time setting, and secondly, some background on the minimax properties of eigenvectors of compact, self-adjoint operators on Hilbert space. Section 4 contains the mathematical setups to handle the three problems described in Section 2. Section 5 discusses the differences and similarities of single- and bi-directional coherence, describes how one can create a sequence of finite-time coherent sets, and summarises some further mathematical properties of the framework. Section 6 details how one can numerically implement the theory in Section 4 in each of the three dynamical settings, and Section 7 contains the two main case studies, in which we explain the similarities and differences of the three dynamical setups, and demonstrate the influence of finite-time duration, diffusion amplitude, and connections with time-directions.

2 Transfer operators and three transport problems

Let $X, Y \subset \mathbb{R}^d$ be compact, let ℓ denote Lebesgue (or volume) measure on X and Y, and consider a map $T : X \to Y$, which is non-singular[1]. The map T may describe discrete time dynamics, or may be a time-t map of a continuous time flow. The *Perron-Frobenius operator* for T, denoted $\mathbf{P} : L^1(X,\ell) \to L^1(Y,\ell)$, describes the evolution of densities under T. Its action on an $f \in L^1(X,\ell)$ is defined by insisting that $\int_A \mathbf{P}f \, d\ell = \int_{T^{-1}A} f \, d\ell$ for all (Borel-)measurable $A \subset Y$. In the situation where T is differentiable, one has the equivalent definition of $\mathbf{P}$: $\mathbf{P}f(y) = \sum_{x \in T^{-1}y} f(x)/|\det DT(x)|$, where $DT(x)$ is the spatial linearisation of T at $x \in X$. If T is differentiable and invertible, as in the situation where T arises as a time-t map of a smooth flow, then the expression above simplifies to $\mathbf{P}$: $\mathbf{P}f(y) = f(T^{-1}y)/|\det DT(T^{-1}y)|$. The operator $\mathbf{P}$ is bounded; in fact $\|\mathbf{P}\|_1 = 1$, and preserves nonnegative functions; that is, $\mathbf{P}f \geq 0$ if $f \geq 0$. A density is invariant under T iff the density is a fixed point of $\mathbf{P}$. Further details on the Perron-Frobenius operator may be found in [28].

In the following, we will introduce linear operators that are built around the Perron-Frobenius operator. We will introduce some noise or diffusion to the purely deterministic action of T, perform some "change of measure space" transformations, and combine these constructions with their duals in order to solve specific dynamical transport problems. Before getting into details, we briefly describe the problems we are interested in solving and prototype operator constructions for solving them.

[1] A map T is non-singular if $\ell(T^{-1}A) = 0$ when $\ell(A) = 0$; thus volume cannot be created from nothing by pulling back with T, or alternatively cannot be completely destroyed by pushing forward with T.

An overarching goal is to detect and locate slow mixing dynamical structures. These structures should be macroscopic in size and by "slow mixing" we have in mind a geometric mixing rate that is slower than $1/\Lambda$, where Λ is the largest[2] positive Lyapunov multiplier. Thus, such "slow mixing" cannot be explained by *local* stretching, but is instead due to the way in which the dynamics acts *globally*.

2.1 Autonomous, time-independent, or periodically forced dynamics

In this setting we have a single map $T : X \to X$ and repeated iteration of T defines the dynamics. The map T may arise as a time-t map of an autonomous ODE $\dot{x} = F(x)$. Alternatively, if the dynamics is periodic with period p, then in discrete time one may study $T = T_p \circ \cdots \circ T_2 \circ T_1$ (or cyclic permutations) as a single map, or in continuous time, $\dot{x} = F(x,t)$, where $F(x,t) = F(x,t+p)$ for all $x \in X$, $t \in \mathbb{R}$, one may set T to the time-p map of the ODE.

We assume that there is some T-invariant[3] probability measure μ, and that one is interested in tracking the transport with respect to this measure. For example, if T is volume-preserving, then a natural choice for μ is Lebesgue measure.

As the dynamics is fixed in time, the global slow-mixing structures we seek are also fixed sets. If a set $A \subset X$ satisfies $A \approx T^{-1}A$, then the set of points that are *currently in A and will remain in A after one time step*, namely $A \cap T^{-1}A$, is large relative to A in the sense of μ-measure, i.e. $\frac{\mu(A \cap T^{-1}A)}{\mu(A)} \approx 1$. Thus, the probability to leave the set A in one time step is low; because of this property and the approximate invariance[4] property, these structures are known in the literature as *almost-invariant sets* or *metastable sets*. We will be interested in optimal almost-invariant sets; optimal in the sense that the conditional probability to leave the set in one time step is minimal.

In terms of operators, in the following sections we will work with an operator $\mathbf{L}$ that is dynamically similar to $\mathbf{P}$, but with the property that $\mathbf{L}\mathbf{1}_X = \mathbf{1}_X$. The invariance condition $A \approx T^{-1}A$ can be translated into the operator equation $\mathbf{L}\mathbf{1}_A \approx \mathbf{1}_A$. For technical reasons discussed later (in addition to a formal definition of $\mathbf{L}$), solving the problem of finding *optimal* almost-invariant sets is helped by forming a symmetrised operator $(\mathbf{L} + \mathbf{L}^*)/2$, where $\mathbf{L}^*$ is the dual of $\mathbf{L}$. While the action of $\mathbf{L}$ can be interpreted as a push-forward under the dynamics T, the action of $\mathbf{L}^*$ can be interpreted as a pull-back (applying T^{-1}). Thus, our symmetrised operator is effectively checking how mass in transported in forward and backward time. Because our measure μ is T-invariant, it does not matter in which time direction we check

[2] As we are considering structures of full dimension, these structures will be stretched and mixed according to the largest positive Lyapunov exponent.

[3] A probability measure μ is called T-invariant if $\mu \circ T^{-1} = \mu$.

[4] A set A is called *invariant* if $A = T^{-1}A$.

for mass loss from a set A; we do however benefit from optimality properties of self-adjoint operators, so the symmetrisation construction is important.

2.2 Nonautonomous, time-dependent, or aperiodically forced dynamics: single time direction

In contrast to the single-map setting above, we now have a sequence of maps $T_{t+\tau} \circ \cdots \circ T_{t+1} \circ T_t$ over a duration of τ time steps. Alternatively, in continuous time, one studies the flow of an ODE $\dot{x} = F(x,t)$ from some initial time t to some final time $t+\tau$. In both cases, we construct a single map $T : X \to T(X)$, which in the discrete time setting is simply $T = T_{t+\tau} \circ \cdots \circ T_{t+1} \circ T_t$, and in the continuous time setting is the flow map of the ODE from time t to time $t+\tau$.

Superficially, we may now appear to be back in the autonomous setting discussed in Section 2.1, however, there are two crucial differences. Firstly, in the autonomous setting, our symmetrisation construction implicitly assumed that the dynamics in backward time is T^{-1}. However, in a general nonautonomous setting, the backward time dynamics is $T_{t-\tau}^{-1} \circ \cdots \circ T_{t-2}^{-1} \circ T_{t-1}^{-1}$, which is different to $T^{-1} = T_t^{-1} \circ T_{t+\tau-1}^{-1} \circ T_{t+\tau}^{-1}$. Secondly, because of the general time-dependence, the slowly mixing structures we are interested in may be *time-varying*, rather than fixed; for example, consider oceanic eddies or atmospheric vortices, both of which move over time. Because of this general time-dependence, we do *not* insist on tracking the finite-time transport with respect to an *invariant* measure, but rather begin with a user-specified probability measure μ at time t. This measure describes a mass distribution that we are interesting in transporting; by the final time $t+\tau$, μ will have transformed into another probability measure ν via the transformation[5] $\nu = \mu \circ T^{-1}$.

As our slowly mixing structures are time-varying, we call them *coherent sets*, to distinguish them from the temporally fixed almost-invariant sets. If the sets $A_t, A_{t+\tau}$ satisfy $A_t \approx T^{-1}A_{t+\tau}$ then the set of points that are *currently in A_t and will be carried to $A_{t+\tau}$ after τ time steps*, namely $A_t \cap T^{-1}A_{t+\tau}$, is large relative to A_t, i.e. $\frac{\mu(A_t \cap T^{-1}A_{t+\tau})}{\mu(A_t)} \approx 1$. Thus, the conditional probability to *not* be carried from A_t to $A_{t+\tau}$ is low, and we say that such $A_t, A_{t+\tau}$ constitute a pair of coherent sets. We will be interested in optimal coherent sets; optimal in the sense that the conditional probability to not be carried from A_t to $A_{t+\tau}$ is minimal.

In terms of operators, we again make use of an operator $\mathbf{L}$ that is dynamically similar to $\mathbf{P}$, where $\mathbf{P}$ is the Perron-Frobenius operator for the map T which governs the dynamics from time t to $t+\tau$; roughly speaking, $\mathbf{L}$ will have the property that $\mathbf{L}\mathbf{1}_X = \mathbf{1}_{T(X)}$. The equivariance condition $A_t \approx T^{-1}A_{t+\tau}$ can be translated into the operator equation $\mathbf{L}\mathbf{1}_{A_t} \approx \mathbf{1}_{A_{t+\tau}}$. Clearly, this operator equation cannot be directly solved as an eigenequation as $\mathbf{1}_{A_t}$ and $\mathbf{1}_{A_{t+\tau}}$ may be very different functions. Instead, our approach will be to push forward $\mathbf{1}_{A_t}$ with $\mathbf{L}$, to obtain something close to $\mathbf{1}_{A_{t+\tau}}$,

[5] Note that $\nu = \mu \circ T^{-1}$ is the natural push-forward of μ under T. For example if μ is supported on a set $A \subset X$, then $\nu(T(A)) = \mu \circ T^{-1}T(A) \geq \mu(A)$ shows that ν is supported on $T(A)$.

and then pull back with $\mathbf{L}^*$ to return to something close to $\mathbf{1}_{A_t}$. The operator $\mathbf{L}$ and $\mathbf{L}^*$ will include a small amount of diffusive dynamics, and at the operator level, the way in which mass is not carried from A_t to $A_{t+\tau}$ is by mass diffusing from both A_t and $A_{t+\tau}$. Leaving technical details for later, we consider eigenfunctions[6] of $\mathbf{L}^*\mathbf{L}$ (push forward from t to $t+\tau$, then pull back from $t+\tau$ to t) to determine the set A_t and eigenfunctions of $\mathbf{L}\mathbf{L}^*$ (pull back from $t+\tau$ to t then push forward from t to $t+\tau$) to determine the set $A_{t+\tau}$. Both $\mathbf{L}^*\mathbf{L}$ and $\mathbf{L}^*\mathbf{L}$ are self-adjoint in appropriate Hilbert spaces and we make use of associated optimality properties to find the *optimal* pair of coherent sets.

2.3 Nonautonomous, time-dependent, or aperiodically forced dynamics: both time directions

The dynamical setting is almost the same as in Section 2.2, except that we consider two segments of time. Firstly, we consider dynamics from time t to $t-\tau_1$, governed by a map T_- and secondly from t to $t+\tau_2$, governed by a map T_+. The maps T_- and T_+ are generated in the same way as in Section 2.2. Our focus is on finding a coherent set at the intermediate time t; one that is coherent in *both*[7] forward and backward time. Thus, we wish to select a triple of sets $A_{t-\tau_1}, A_t, A_{t+\tau_2}$ with the property that $A_{t-\tau_1} \approx T_- A_t$ *and* $A_{t+\tau_2} \approx T_+ A_t$. As in Section 2.2, the user prescribes a probability measure μ at time t that represents a mass distribution we are interesting in transporting.

In terms of operators, we again make use of operators $\mathbf{L}_-$ and $\mathbf{L}_+$ that are dynamically similar to $\mathbf{P}_{T_-}$ and $\mathbf{P}_{T_+}$, where $\mathbf{P}_{T_\pm}$ is the Perron-Frobenius operator for the map $T_\pm$. The equivariance conditions $A_{t-\tau_1} \approx T_- A_t$ and $A_t \approx T_+^{-1} A_{t+\tau_2}$ can be translated into the operator equations $\mathbf{L}_-\mathbf{1}_{A_t} \approx \mathbf{1}_{A_{t-\tau_1}}$ and $\mathbf{L}_+\mathbf{1}_{A_t} \approx \mathbf{1}_{A_{t+\tau_2}}$, respectively. As in Section 2.2, these operator equations cannot be directly solved as an eigenequation. In analogy to the single-direction case, our approach will be to push forward $\mathbf{1}_{A_t}$ with $\mathbf{L}_+$, to obtain something close to $\mathbf{1}_{A_{t+\tau_2}}$, and then pull back with $\mathbf{L}_+^*$ to return to something close to $\mathbf{1}_{A_t}$; this utilises the dynamics on the interval $[t, t+\tau_2]$. Similarly, we also push forward $\mathbf{1}_{A_t}$ (under backward-time dynamics) with $\mathbf{L}_-$, and then pull back with $\mathbf{L}_-^*$ to again return to something close to $\mathbf{1}_{A_t}$, however, this time utilising the dynamics from t to $t-\tau_1$. Both $\mathbf{L}_+^*\mathbf{L}_+$ and $\mathbf{L}_-^*\mathbf{L}_-$ are self-adjoint in the same Hilbert space (anchored at time t). We may now average the effect of the dynamics over the intervals $[t-\tau_1]$ and $[t, t+\tau_2]$ by averaging the operators to form $(\mathbf{L}_-^*\mathbf{L}_- + \mathbf{L}_+^*\mathbf{L}_+)/2$, again obtaining a self-adjoint operator, which has the necessary optimality properties to find the *optimal* triple of coherent sets.

[6] In numerical computations it is cheaper to compute singular vectors of $\mathbf{L}$, than to form $\mathbf{L}^*\mathbf{L}$ and compute eigenvectors.

[7] A similar result could also be achieved by a more direct application of the construction in the previous section; we discuss this further in Section 5.2.

3 Two key tools

In this section we introduce two key objects for the analysis that follows. The first is the operator $\mathbf{L}$, which acts as a building block for purpose-built operators for each of the three dynamical settings we consider. The second is a class of self-adjoint operators, which we use heavily to obtain optimality results.

3.1 A building block operator

Our dynamical system is acting on a subset M of $\mathbb{R}^d$ and our transport analysis will be confined to a compact subset X of M and neighbourhoods of X and $T(X)$. We construct our basic building block operator $\mathbf{L}$ from the Perron-Frobenius operator for T. The constructions here are common to the three settings we consider and may be found in greater detail in [13]. Specialisations for the autonomous case will be pointed out after the main construction. We endow X with a probability measure μ, which we assume is absolutely continuous with positive density h_μ.

We begin building the operator $\mathbf{L}$ by pre- and post- applying diffusion to $\mathbf{P}$. We define a diffusion operator on X, $\mathbf{D}_{X,\varepsilon} : L^1(X,\ell) \to L^1(X_\varepsilon,\ell)$ by $\mathbf{D}_{X,\varepsilon} f(y) = \int_X \alpha_{X,\varepsilon}(y-x) f(x) dx$, where $\alpha_{X,\varepsilon} : X_\varepsilon \to \mathbb{R}_+$ is bounded and satisfies $\int_{X_\varepsilon} \alpha_{X,\varepsilon}(y-x)\ dy = 1$ for all $x \in X$. Similarly we define $\mathbf{D}_{Y'_\varepsilon,\varepsilon} : L^1(Y'_\varepsilon,\ell) \to L^1(Y_\varepsilon,\ell)$ by $\mathbf{D}_{Y'_\varepsilon,\varepsilon} f(y) = \int_{Y'_\varepsilon} \alpha_{Y,\varepsilon}(y-x) f(x) dx$, where $\alpha_{Y,\varepsilon}$ is bounded and satisfies $\int_{Y_\varepsilon} \alpha_{Y,\varepsilon}(y-x)\ dy = 1$ for $x \in Y'_\varepsilon$. We think of $X_\varepsilon = \mathrm{supp}(\mathbf{D}_{X,\varepsilon}\mathbf{1}_X)$, $Y'_\varepsilon = T(X_\varepsilon)$ and $Y_\varepsilon = \mathrm{supp}(\mathbf{D}_{Y'_\varepsilon,\varepsilon}\mathbf{1}_{Y'_\varepsilon})$. In terms of function spaces we have

$$L^1(X,\ell) \xrightarrow{\mathbf{D}_{X,\varepsilon}} L^1(X_\varepsilon,\ell) \xrightarrow{\mathbf{P}} L^1(Y'_\varepsilon,\ell) \xrightarrow{\mathbf{D}_{Y'_\varepsilon,\varepsilon}} L^1(Y_\varepsilon,\ell) \tag{1}$$

As an intermediate operator, we define $\mathbf{P}_\varepsilon = \mathbf{D}_{Y'_\varepsilon,\varepsilon}\mathbf{P}\mathbf{D}_{X,\varepsilon}$, our pre- and post- diffused Perron-Frobenius operator. In general $\mathbf{P}_\varepsilon \mathbf{1}_X \neq \mathbf{1}_{Y_\varepsilon}$; to obtain this property, we perform a "change of measure" transformation, and define

$$\mathbf{L}_\varepsilon f = \mathbf{P}_\varepsilon(f \cdot h_\mu)/\mathbf{P}_\varepsilon(h_\mu); \tag{2}$$

it is now obvious that $\mathbf{L}_\varepsilon \mathbf{1}_X = \mathbf{1}_{Y_\varepsilon}$ and in fact $\mathbf{L}_\varepsilon^* \mathbf{1}_{Y_\varepsilon} = \mathbf{1}_X$; see [13].

The setup described above covers both time-dependent situations under consideration; the map T can control the dynamics over the interval $[t, t+\tau]$ as in Section 2.2 or over one of the intervals $[t, t+\tau_2]$, $[t-\tau_1, t]$ as in Section 2.3. In the autonomous setting, there is one additional hypotheses, namely that h_μ is fixed by the advective and diffusive dynamics; that is[8], $\mathbf{P}_\varepsilon h_\mu = h_\mu$. It follows that $X = Y_\varepsilon$; examples of such a situation include:

1. X is boundaryless (eg. a solid torus represented by a solid cube in $\mathbb{R}^3$, identifying opposing faces),

[8] In general, $\mathbf{P}_\varepsilon h_\mu = h_{\nu_\varepsilon}$, where h_{ν_ε} is a density of the push-forward of h_μ by the dynamics.

2. X is an attractor,
3. X has boundaries, however, $\alpha_{X,\varepsilon}, \alpha_{Y,\varepsilon}$ are chosen in such a way that the diffusion does not perturb points across the boundary of X.

From chapter 4 onwards, we will consider the situations where one chooses $\alpha_{X,\varepsilon} = \alpha_{Y,\varepsilon} = (1/\ell(B_\varepsilon(0)))\mathbf{1}_{B_\varepsilon(0)}$. Dynamically, this means that one diffuses uniformly over an ε-ball, then applies the deterministic dynamics T, the diffuses uniformly over an ε-ball again. We have chosen such $\alpha_{X,\varepsilon}, \alpha_{Y,\varepsilon}$ as these are very natural choices of small random perturbations of the dynamics with bounded support. An important property of our building block operator is compactness.

Theorem 1 ([13]). *If Y_ε has finite Lebesgue measure and one chooses $\alpha_{X,\varepsilon} = \alpha_{Y,\varepsilon} = (1/\ell(B_\varepsilon(0)))\mathbf{1}_{B_\varepsilon(0)}$, then $\mathbf{L}_\varepsilon : L^2(X,\mu) \to L^2(Y_\varepsilon, \nu_\varepsilon)$ is compact operator.*

3.2 Optimality properties of compact self-adjoint operators on Hilbert space

We recall the minimax principle for compact self-adjoint operators, which forms a key part of our approach. Let $\mathbf{Q} : \mathscr{H} \to \mathscr{H}$ be a compact, self-adjoint operator on a Hilbert space $\mathscr{H}$. Then $\mathbf{Q}$ has only a countable number of distinct eigenvalues. We enumerate the positive eigenvalues of $\mathbf{Q}$, $\lambda_1^+ \geq \lambda_2^+ \geq \cdots$ and the negative eigenvalues, $\lambda_1^- \leq \lambda_2^- \leq \cdots$, where the number of occurrences equals the multiplicity of the eigenvalue. We choose orthonormal bases of eigenvectors $u_l^\pm$, enumerated so that $\mathbf{Q}u_l^\pm = \lambda_l^\pm u_l^\pm$. We may write

$$\mathbf{Q} = \sum_{l=1}^{M^-} \lambda_l^- \langle \cdot, u_l^- \rangle u_l^- + \sum_{l=1}^{M^+} \lambda_l^+ \langle \cdot, u_l^+ \rangle u_l^+, \tag{3}$$

where $M^\pm$ may be finite or infinite. (see eg. Theorem II.5.1 [6]).

If $\mathbf{Q}$ is also positive[9] then

$$\mathbf{Q} = \sum_{l=1}^{M^+} \lambda_l^+ \langle \cdot, u_l^+ \rangle u_l^+, \tag{4}$$

where M^+ may be finite or infinite. All eigenvalues of $\mathbf{Q}$ are nonnegative and are denoted $\lambda_1^+ \geq \lambda_2^+ \geq \cdots$; see eg. Proposition II.7.14 [6].

One has the following minimax principle (see eg. Theorem 9.2.4, p212 [3]):

Theorem 2. *Let $\mathbf{Q} : \mathscr{H} \to \mathscr{H}$ be compact and self-adjoint.*

$$\lambda_l^+ = \min_{V:\operatorname{codim} V \leq l-1 < M^+} \max_{0 \neq f \in V} \frac{\langle \mathbf{Q}f, f \rangle}{\langle f, f \rangle}, \ell = 1, \ldots, M^+ \tag{5}$$

[9] An operator $\mathbf{Q}$ on a Hilbert space $\mathscr{H}$ is called positive if $\langle \mathbf{Q}x, x \rangle \geq 0$ for all $x \in \mathscr{H}$.

and

$$\lambda_l^- = -\min_{V:\operatorname{codim} V \le l-1 < M^-} \max_{0 \neq f \in V} \frac{-\langle \mathbf{Q}f, f\rangle}{\langle f, f\rangle}, \ell = 1, \ldots, M^-. \tag{6}$$

Furthermore, the maximising fs are u_l^+ and u_l^-, respectively.

The eigenvalue we are particularly interested in is λ_2^+. Specialising to $\mathscr{H} = L^2(X,\mu)$, with $\langle f, g\rangle_\mu = \int f \cdot g \, d\mu$, one has

Corollary 1. *If λ_1^+ is simple, then the V in (5) is $\operatorname{sp}\{u_1^+\}^\perp$, one has*

$$\lambda_2^+ = \max_{f \in L^2(X,\mu)} \left\{ \frac{\langle \mathbf{Q}f, f\rangle_\mu}{\langle f, f\rangle_\mu} : \langle f, u_1^+\rangle_\mu = 0 \right\}. \tag{7}$$

and the maximising f is u_2^+.

4 Main constructions and results

Using our building block operator $\mathbf{L}$, in each of the three dynamical settings we construct a suitable operator-based optimisation problem whose solution yields optimal almost-invariant or coherent sets. We then form a relaxation of this problem and show that at its core is a self-adjoint operator. Finally, we utilise the optimality properties of this self-adjoint operator to exactly solve this relaxed problem via eigenvectors, and then use these eigenvectors to construct almost-invariant or coherent sets.

4.1 Autonomous dynamics

We consider the operator $\mathbf{L}_\varepsilon : L^2(X,\mu) \to L^2(X,\mu)$, defined as in (2) with $\mathbf{P}$ the Perron-Frobenius operator for T. We wish to measurably partition $X = A \cup A^c$ such that $\mathbf{L}_\varepsilon \mathbf{1}_A \approx \mathbf{1}_A$, $\mathbf{L}_\varepsilon \mathbf{1}_{A^c} \approx \mathbf{1}_{A^c}$ and $\mu(A) \approx \mu(A^c)$. This can be achieved by considering

$$\rho(A) = \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu}{\mu(A)} + \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu}{\mu(A^c)}. \tag{8}$$

The above expression can be interpreted as follows. Supposing for the moment that $\mathbf{L}_\varepsilon = \mathbf{L}_0$ defined as $\mathbf{L}_0 f = \mathbf{P}(f \cdot h_\mu)/h_\mu$; this is a "advection only" version of $\mathbf{L}_\varepsilon$. Then

$$\frac{\langle \mathbf{L}_0 \mathbf{1}_A, \mathbf{1}_A\rangle_\mu}{\mu(A)} = \frac{\langle \mathbf{P}(\mathbf{1}_A \cdot h_\mu)/h_\mu, \mathbf{1}_A\rangle_\mu}{\mu(A)} \tag{9}$$

$$= \frac{\langle \mathbf{P}(\mathbf{1}_A \cdot h_\mu), \mathbf{1}_A\rangle_\ell}{\mu(A)} = \frac{\langle \mathbf{1}_A \cdot h_\mu, \mathbf{1}_A \circ T\rangle_\ell}{\mu(A)} = \frac{\langle \mathbf{1}_A, \mathbf{1}_A \circ T\rangle_\mu}{\mu(A)} = \frac{\mu(A \cap T^{-1}A)}{\mu(A)}. \tag{10}$$

Similarly, $\langle \mathbf{L}_0 \mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu / \mu(A^c) = \mu(A^c \cap T^{-1}A^c)/\mu(A^c)$, thus, in the zero diffusion setting

$$\rho_0(A) = \frac{\mu(A \cap T^{-1}A)}{\mu(A)} + \frac{\mu(A^c \cap T^{-1}A^c)}{\mu(A^c)},$$

which is a natural expression to optimise over A to find the optimal almost-invariant set A. A difficulty with the purely deterministic point of view is that if one is prepared to make A very complicated, one can push the value of $\rho(A)$ arbitrarily close to 2. We use the addition of a small amount of diffusion via the operators $\mathbf{D}_{X,\varepsilon}$ to regularise this optimisation problem, while not changing greatly the expression for $\rho(A)$.

As $\langle \mathbf{L}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu = \overline{\langle \mathbf{L}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu} = \langle \mathbf{L}_\varepsilon^* \mathbf{1}_A, \mathbf{1}_A\rangle_\mu$ one has

$$\rho(A) = \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu}{\mu(A)} + \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu}{\mu(A^c)} = \frac{\langle \mathbf{Q}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu}{\mu(A)} + \frac{\langle \mathbf{Q}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu}{\mu(A^c)}, \qquad (11)$$

where $\mathbf{Q}_\varepsilon = (\mathbf{L}_\varepsilon + \mathbf{L}_\varepsilon^*)/2$. The advantage of $\mathbf{Q}_\varepsilon$ is that it is self-adjoint in $L^2(X,\mu)$. Now one has

$$\begin{aligned}
&\rho(A) - 1 \\
&= \frac{\langle \mathbf{Q}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu}{\mu(A)} + \frac{\langle \mathbf{Q}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu}{\mu(A^c)} - 1 \\
&= \left(\frac{\mu(A^c)}{\mu(A)} + 1\right) \langle \mathbf{Q}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu + \left(\frac{\mu(A)}{\mu(A^c)} + 1\right) \langle \mathbf{Q}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu - \left(\langle \mathbf{Q}_\varepsilon \mathbf{1}_A, \mathbf{1}_X\rangle_\mu + \langle \mathbf{Q}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_X\rangle_\mu\right) \\
&= \left(\frac{\mu(A^c)}{\mu(A)} \langle \mathbf{Q}_\varepsilon \mathbf{1}_A, \mathbf{1}_A\rangle_\mu + \frac{\mu(A)}{\mu(A^c)} \langle \mathbf{Q}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu\right) - \left(\langle \mathbf{Q}_\varepsilon \mathbf{1}_A, \mathbf{1}_X - \mathbf{1}_A\rangle_\mu + \langle \mathbf{Q}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_X - \mathbf{A^c}\rangle_\mu\right) \\
&= \left(\frac{\mu(A^c)}{\mu(A)} \int \mathbf{Q}_\varepsilon \mathbf{1}_A \cdot \mathbf{1}_A d\mu + \frac{\mu(A)}{\mu(A^c)} \int \mathbf{Q}_\varepsilon \mathbf{1}_{A^c} \cdot \mathbf{1}_{A^c} d\mu\right) - \left(\int \mathbf{Q}_\varepsilon \mathbf{1}_A \cdot \mathbf{1}_{A^c} d\mu + \int \mathbf{Q}_\varepsilon \mathbf{1}_{A^c} \cdot \mathbf{1}_A d\mu\right) \\
&= \left\langle \mathbf{Q}_\varepsilon \left(\sqrt{\frac{\mu(A^c)}{\mu(A)}} \mathbf{1}_A - \sqrt{\frac{\mu(A)}{\mu(A^c)}} \mathbf{1}_{A^c}\right), \left(\sqrt{\frac{\mu(A^c)}{\mu(A)}} \mathbf{1}_A - \sqrt{\frac{\mu(A)}{\mu(A^c)}} \mathbf{1}_{A^c}\right)\right\rangle_\mu
\end{aligned}$$

Thus, finding $A \subset X$ that maximises $\rho(A)$ is equivalent to the problem:

$$\max_{A \subset X} \left\langle \mathbf{Q}_\varepsilon \left(\sqrt{\frac{\mu(A^c)}{\mu(A)}} \mathbf{1}_A - \sqrt{\frac{\mu(A)}{\mu(A^c)}} \mathbf{1}_{A^c}\right), \left(\sqrt{\frac{\mu(A^c)}{\mu(A)}} \mathbf{1}_A - \sqrt{\frac{\mu(A)}{\mu(A^c)}} \mathbf{1}_{A^c}\right)\right\rangle_\mu . \qquad (12)$$

Note that $\left\langle \sqrt{\frac{\mu(A^c)}{\mu(A)}} \mathbf{1}_A - \sqrt{\frac{\mu(A)}{\mu(A^c)}} \mathbf{1}_{A^c}, \mathbf{1}_X \right\rangle_\mu = 0$ for any choice of A. Thus, a relaxed form of (12), where we remove the restriction that the argument of $\mathbf{Q}_\varepsilon$ is a difference of characteristic functions, but retain the orthogonality property of this ansatz is:

$$\max_{f \in L^2(X,\mu)} \left\{ \frac{\langle \mathbf{Q}_\varepsilon f, f\rangle_\mu}{\langle f, f, \rangle_\mu} : \langle f, \mathbf{1}_X\rangle_\mu = 0 \right\}. \qquad (13)$$

Recall that $\mathbf{Q}_\varepsilon \mathbf{1}_X = \mathbf{1}_X$, that $\mathbf{Q}_\varepsilon$ is self-adjoint and compact (by Theorem 1 setting $Y_\varepsilon = X$ and $\nu_\varepsilon = \mu$), but not positive, in general. Furthermore (see discussion in Section 5.3) $\lambda_1^+ = 1$ is simple. We may therefore apply Corollary 1 to see that the maximum value of (13) is λ_2, the second largest eigenvalue of $\mathbf{Q}_\varepsilon$, and that the maximising f is u_2, the corresponding eigenvector of $\mathbf{Q}_\varepsilon$. As (13) is a relaxation of (12), we immediately see that $\rho(A) \leq 1+\lambda_2$ for all measurable $A \subset X$. One can also obtain an *a priori* lower bound for $\rho(A)$:

Theorem 3.

$$2 - 2\sqrt{2(1-\lambda_2)} \leq \sup_{A \subset X} \rho(A) \leq 1 + \lambda_2.$$

Proof. See appendix A.

The above result is strongly related to classic conductance bounds in the reversible Markov chain literature [29, 25] and has been discussed in regard to almost-invariant sets in a matrix setting [11]. An *a posteriori* lower bound (relying on eigenvector computations) has been proposed by Huisinga and Schmidt [24].

Finally, to construct a partition A, A^c from the solution to (13), we set $A = A_\beta := \{x \in X : u_2(x) > \beta\}$ and $A^c = A_\beta^c := \{x \in X : u_2(x) \leq \beta\}$, where β is chosen to maximise $\rho(A_\beta)$. In numerical computations, β is obtained via a line search, as described in Section 6.3

4.2 Nonautonomous or time-dependent dynamics: single time direction

The constructions in this section may be found in greater detail in [13]; we recall only the main points here. We consider the operator $\mathbf{L}_\varepsilon : L^2(X,\mu) \to L^2(Y_\varepsilon, \nu_\varepsilon)$ defined as in (2) using $\mathbf{P}$ the Perron-Frobenius operator for T, where T represents the finite-time dynamics from some time t to $t+\tau$. We wish to measurably partition $X = A \cup A^c$ and $Y_\varepsilon = B \cup B^c$ such that

1. $\mathbf{L}_\varepsilon \mathbf{1}_A \approx \mathbf{1}_B$ and $\mathbf{L}_\varepsilon \mathbf{1}_{A^c} \approx \mathbf{1}_{B^c}$,
2. $\mu(A) = \nu_\varepsilon(B)$ and $\mu(A^c) = \nu_\varepsilon(B^c)$.

We would also like $\mu(A) \approx \mu(A^c)$ and $\nu_\varepsilon(B) \approx \nu_\varepsilon(B^c)$.

This can be achieved by considering

$$\rho(A,B) = \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_A, \mathbf{1}_B \rangle_{\nu_\varepsilon}}{\mu(A)} + \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{B^c} \rangle_{\nu_\varepsilon}}{\mu(A^c)}. \tag{14}$$

Using arguments identical to the autonomous case above, one can show that in the "advection only" setting, one has

$$\rho_0(A,B) = \frac{\mu(A \cap T^{-1}B)}{\mu(A)} + \frac{\mu(A^c \cap T^{-1}B^c)}{\mu(A^c)}.$$

Thus, it is now clear that to maximise $\rho(A,B)$ one should choose A so that it is close to the preimage of B (and likewise for the complements). Of course, one could choose $A = T^{-1}B$ to obtain $\rho(A,B) = 2$ and so there is a problem of nonuniqueness in solutions to an optimisation of $\rho(A,B)$ over A and B. As in the autonomous setting we require the small amount of diffusion via $\mathbf{D}_{X,\varepsilon}, \mathbf{D}_{Y'_\varepsilon,\varepsilon}$ to regularise the optimisation of $\rho(A,B)$ over A and B; this also (generically) produces a unique optimum.

Consider the optimisation problem

$$\sup_{A\subset X, B\subset Y_\varepsilon} \{\rho(A,B) - 1 : \mu(A) = \nu_\varepsilon(B)\}, \tag{15}$$

where we have subtracted a constant 1 for convenience. One has (see [13] for details)

$$\rho(A,B) - 1 \tag{16}$$

$$= \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_A, \mathbf{1}_B\rangle_{\nu_\varepsilon}}{\mu(A)} + \frac{\langle \mathbf{L}_\varepsilon \mathbf{1}_{A^c}, \mathbf{1}_{B^c}\rangle_{\nu_\varepsilon}}{\mu(A^c)} - 1 \tag{17}$$

$$= \left\langle \mathbf{L}_\varepsilon \left(\sqrt{\frac{\mu(A^c)}{\mu(A)}}\mathbf{1}_A - \sqrt{\frac{\mu(A)}{\mu(A^c)}}\mathbf{1}_{A^c} \right), \left(\sqrt{\frac{\nu_\varepsilon(B^c)}{\nu_\varepsilon(B)}}\mathbf{1}_B - \sqrt{\frac{\nu_\varepsilon(B)}{\nu_\varepsilon(B^c)}}\mathbf{1}_{B^c} \right) \right\rangle_{\nu_\varepsilon} \tag{18}$$

We note that $\left\langle \sqrt{\frac{\mu(A^c)}{\mu(A)}}\mathbf{1}_A - \sqrt{\frac{\mu(A)}{\mu(A^c)}}\mathbf{1}_{A^c}, \mathbf{1}_X \right\rangle_\mu = 0$ and $\left\langle \sqrt{\frac{\nu_\varepsilon(B^c)}{\nu_\varepsilon(B)}}\mathbf{1}_B - \sqrt{\frac{\nu_\varepsilon(B)}{\nu_\varepsilon(B^c)}}\mathbf{1}_{B^c}, \mathbf{1}_{Y_\varepsilon} \right\rangle_{\nu_\varepsilon} = 0$ for any choice of A,B. Thus, a relaxed form of (15) where we remove the restriction that the argument of $\mathbf{L}_\varepsilon$ is a difference of characteristic functions, but retain the orthogonality property of this ansatz is:

$$\max_{f\in L^2(X,\mu), g\in L^2(Y_\varepsilon,\nu_\varepsilon)} \left\{ \frac{\langle \mathbf{L}_\varepsilon f, g\rangle_{\nu_\varepsilon}}{\|f\|_\mu \|g\|_{\nu_\varepsilon}} : \langle f, \mathbf{1}_X\rangle_\mu = \langle g, \mathbf{1}_{Y_\varepsilon}\rangle_{\nu_\varepsilon} = 0, \right\}. \tag{19}$$

Proposition 1 [13] shows that the value of (19) is $\lambda_2^{1/2}$, the square root of the second largest eigenvalue of $\mathbf{Q}_\varepsilon := \mathbf{L}_\varepsilon^* \mathbf{L}_\varepsilon$ and the maximising f (resp. g) is u_2 (resp. v_2), the corresponding eigenvector of $\mathbf{Q}_\varepsilon$ (resp. $\mathbf{Q}_\varepsilon^*$). As (19) is a relaxation of (15), using (16)–(18) one obtains

Theorem 4. *[Froyland [13]]*

$$\max_{A\subset X, B\subset Y_\varepsilon} \{\rho(A,B) : \mu(A) = \nu_\varepsilon(B)\} \le 1 + \lambda_2^{1/2}.$$

To construct a partition A, A^c from the solution to (19), we set $A = A_\beta := \{x \in X : u_2(x) > \beta\}$, $A^c = A_\beta^c := \{x \in X : u_2(x) \le \beta\}$, $B = B_{\beta'} := \{y \in Y_\varepsilon : v_2(y) > \beta'\}$, $B^c = B_{\beta'}^c := \{y \in Y_\varepsilon : v_2(y) \le \beta'\}$, where β is chosen to maximise $\rho(A_\beta, B_{\beta'})$ and $\beta' = \beta'(\beta)$ is chosen so that $\mu(A_\beta) = \nu_\varepsilon(B'_\beta)$. In practice, the optimal value of β is found via a line search; see Section 6.4 for details.

4.3 Nonautonomous or time-dependent dynamics: both time directions

Recall that T_+ represents the finite-time dynamics from some time t to $t+\tau_2$ and that T_- represents the finite-time dynamics from some time t to $t-\tau_1$. We consider the operators $\mathbf{L}_{\varepsilon,+} : L^2(X,\mu) \to L^2(Y_\varepsilon, \nu_\varepsilon)$ and $\mathbf{L}_{\varepsilon,-} : L^2(X,\mu) \to L^2(Z_\varepsilon, \eta_\varepsilon)$. The operator $\mathbf{L}_{\varepsilon,+}$ is defined by substituting $\mathbf{P}_{T_+}$ (the Perron-Frobenius operator for T_+) for $\mathbf{P}$ in the expression (2) for $\mathbf{L}$ in Section 3.1. The operator $\mathbf{L}_{\varepsilon,-}$ is defined by substituting $\mathbf{P}_{T_-}$ (the Perron-Frobenius operator for T_-) for $\mathbf{P}$ in (2), and replacing $Y'_\varepsilon, Y_\varepsilon$ with $Z'_\varepsilon, Z_\varepsilon$. In summary,

$$L^1(Z_\varepsilon,\ell) \xleftarrow{\mathbf{D}_{Z'_\varepsilon,\varepsilon}} L^1(Z'_\varepsilon,\ell) \xleftarrow{\mathbf{P}_{T_-}} L^1(X_\varepsilon,\ell) \xleftarrow{\mathbf{D}_{X,\varepsilon}} L^1(X,\ell) \xrightarrow{\mathbf{D}_{X,\varepsilon}} L^1(X_\varepsilon,\ell) \xrightarrow{\mathbf{P}_{T_+}} L^1(Y'_\varepsilon,\ell) \xrightarrow{\mathbf{D}_{Y'_\varepsilon,\varepsilon}} L^1(Y_\varepsilon,\ell) \tag{20}$$

We wish to measurably partition $X = A \cup A^c$, $Y_\varepsilon = B \cup B^c$, and $Z_\varepsilon = C \cup C^c$ such that

1. $\mathbf{L}_{\varepsilon,+}\mathbf{1}_A \approx \mathbf{1}_B$ and $\mathbf{L}_{\varepsilon,+}\mathbf{1}_{A^c} \approx \mathbf{1}_{B^c}$,
2. $\mathbf{L}_{\varepsilon,-}\mathbf{1}_A \approx \mathbf{1}_C$ and $\mathbf{L}_{\varepsilon,-}\mathbf{1}_{A^c} \approx \mathbf{1}_{C^c}$,
3. $\mu(A) = \nu_\varepsilon(B) = \eta_\varepsilon(C)$ and $\mu(A^c) = \nu_\varepsilon(B^c) = \eta_\varepsilon(C^c)$.

We would also like $\mu(A) \approx \mu(A^c)$, $\nu_\varepsilon(B) \approx \nu_\varepsilon(B^c)$, and $\eta_\varepsilon(C) \approx \eta_\varepsilon(C^c)$.

We consider

$$\rho(A,B,C) \tag{21}$$
$$= \frac{1}{2}\left(\left(\frac{\langle \mathbf{L}_{\varepsilon,+}\mathbf{1}_A, \mathbf{1}_B\rangle_{\nu_\varepsilon}}{\mu(A)} + \frac{\langle \mathbf{L}_{\varepsilon,+}\mathbf{1}_{A^c}, \mathbf{1}_{B^c}\rangle_{\nu_\varepsilon}}{\mu(A^c)}\right) + \left(\frac{\langle \mathbf{L}_{\varepsilon,-}\mathbf{1}_A, \mathbf{1}_C\rangle_{\eta_\varepsilon}}{\mu(A)} + \frac{\langle \mathbf{L}_{\varepsilon,-}\mathbf{1}_{A^c}, \mathbf{1}_{C^c}\rangle_{\eta_\varepsilon}}{\mu(A^c)}\right)\right).$$

Using arguments identical to the autonomous case above, one can show that in the "advection only" setting, one has

$$\rho_0(A,B,C)$$
$$= \frac{1}{2}\left(\left(\frac{\mu(A \cap T_+^{-1}B)}{\mu(A)} + \frac{\mu(A^c \cap T_+^{-1}B^c)}{\mu(A^c)}\right) + \left(\frac{\mu(A \cap T_-^{-1}C)}{\mu(A)} + \frac{\mu(A^c \cap T_-^{-1}C^c)}{\mu(A^c)}\right)\right).$$

Consider the optimisation problem

$$\sup_{A\subset X, B\subset Y_\varepsilon, C\subset Z_\varepsilon} \{\rho(A,B,C) - 1 : \mu(A) = \nu_\varepsilon(B) = \eta_\varepsilon(C)\}, \tag{22}$$

where we have subtracted a constant 1 for convenience. Using the expression (21) for $\rho(A,B,C)$ we have (using (18) twice)

$$\rho(A,B,C)-1 \tag{23}$$
$$= \frac{1}{2}\left(\left(\frac{\langle \mathbf{L}_{\varepsilon,+}\mathbf{1}_A,\mathbf{1}_B\rangle_{\nu_\varepsilon}}{\mu(A)}+\frac{\langle \mathbf{L}_{\varepsilon,+}\mathbf{1}_{A^c},\mathbf{1}_{B^c}\rangle_{\nu_\varepsilon}}{\mu(A^c)}\right)+\left(\frac{\langle \mathbf{L}_{\varepsilon,-}\mathbf{1}_A,\mathbf{1}_C\rangle_{\eta_\varepsilon}}{\mu(A)}+\frac{\langle \mathbf{L}_{\varepsilon,-}\mathbf{1}_{A^c},\mathbf{1}_{C^c}\rangle_{\eta_\varepsilon}}{\mu(A^c)}\right)\right)-1$$
$$= \frac{1}{2}\left(\left\langle \mathbf{L}_{\varepsilon,+}\left(\sqrt{\frac{\mu(A^c)}{\mu(A)}}\mathbf{1}_A-\sqrt{\frac{\mu(A)}{\mu(A^c)}}\mathbf{1}_{A^c}\right),\sqrt{\frac{\nu_\varepsilon(B^c)}{\nu_\varepsilon(B)}}\mathbf{1}_B-\sqrt{\frac{\nu_\varepsilon(B)}{\nu_\varepsilon(B^c)}}\mathbf{1}_{B^c}\right\rangle_{\nu_\varepsilon}\right. \tag{24}$$
$$\left.+\left\langle \mathbf{L}_{\varepsilon,-}\left(\sqrt{\frac{\mu(A^c)}{\mu(A)}}\mathbf{1}_A-\sqrt{\frac{\mu(A)}{\mu(A^c)}}\mathbf{1}_{A^c}\right),\sqrt{\frac{\eta_\varepsilon(C^c)}{\eta_\varepsilon(C)}}\mathbf{1}_C-\sqrt{\frac{\eta_\varepsilon(C)}{\eta_\varepsilon(C^c)}}\mathbf{1}_{C^c}\right\rangle_{\eta_\varepsilon}\right). \tag{25}$$

Using the shorthand $\Psi_A=\sqrt{\frac{\mu(A^c)}{\mu(A)}}\mathbf{1}_A-\sqrt{\frac{\mu(A)}{\mu(A^c)}}\mathbf{1}_{A^c}$, $\Psi_B=\sqrt{\frac{\nu_\varepsilon(B^c)}{\nu_\varepsilon(B)}}\mathbf{1}_B-\sqrt{\frac{\nu_\varepsilon(B)}{\nu_\varepsilon(B^c)}}\mathbf{1}_{B^c}$ and $\Psi_C=\sqrt{\frac{\eta_\varepsilon(C^c)}{\eta_\varepsilon(C)}}\mathbf{1}_C-\sqrt{\frac{\eta_\varepsilon(C)}{\eta_\varepsilon(C^c)}}\mathbf{1}_{C^c}$, it is straightforward to verify that $||\Psi_A||_\mu = ||\Psi_B||_{\nu_\varepsilon}=||\Psi_C||_{\eta_\varepsilon}=1$ and $\langle\Psi_A,\mathbf{1}\rangle_\mu=\langle\Psi_B,\mathbf{1}\rangle_{\nu_\varepsilon}=\langle\Psi_C,\mathbf{1}\rangle_{\eta_\varepsilon}=0$. Thus, a relaxed form of (22), where we remove the restriction that the arguments of $\mathbf{L}_{\varepsilon,+}$ and $\mathbf{L}_{\varepsilon,-}$ are differences of characteristic functions, but retain the orthogonality properties of this ansatz is:

$$\max_{f\in L^2(X,\mu),g\in L^2(Y_\varepsilon,\nu_\varepsilon),h\in L^2(Z_\varepsilon,\eta_\varepsilon)}\left\{\frac{1}{2}\left(\frac{\langle \mathbf{L}_{\varepsilon,+}f,g\rangle_{\nu_\varepsilon}}{||f||_\mu||g||_{\nu_\varepsilon}}+\frac{\langle \mathbf{L}_{\varepsilon,-}f,h\rangle_{\eta_\varepsilon}}{||f||_\mu||h||_{\eta_\varepsilon}}\right)\right. \tag{26}$$
$$\left.:\langle f,\mathbf{1}\rangle_\mu=\langle g,\mathbf{1}\rangle_{\nu_\varepsilon}=\langle h,\mathbf{1}\rangle_{\eta_\varepsilon}=0\right\}$$

We now rewrite this relaxed optimisation problem in terms of the self-adjoint operator $\mathbf{Q}_\varepsilon=(\mathbf{L}^*_{\varepsilon,+}\mathbf{L}_{\varepsilon,+}+\mathbf{L}^*_{\varepsilon,-}\mathbf{L}_{\varepsilon,-})/2$.

$$(26)=\max_{f\in L^2(X,\mu)}\left\{\frac{1}{2}\left(\frac{\langle \mathbf{L}_{\varepsilon,+}f,\mathbf{L}_{\varepsilon,+}f\rangle_{\nu_\varepsilon}}{||f||_\mu\|\mathbf{L}_{\varepsilon,+}f\|_{\nu_\varepsilon}}+\frac{\langle \mathbf{L}_{\varepsilon,-}f,\mathbf{L}_{\varepsilon,-}f\rangle_{\eta_\varepsilon}}{||f||_\mu\|\mathbf{L}_{\varepsilon,-}f\|_{\eta_\varepsilon}}\right):\langle f,\mathbf{1}\rangle_\mu=0\right\}$$
$$=\max_{f\in L^2(X,\mu)}\left\{\frac{1}{2}\left(\frac{\langle \mathbf{L}^*_{\varepsilon,+}\mathbf{L}_{\varepsilon,+}f,f\rangle_\mu}{||f||_\mu\langle \mathbf{L}^*_{\varepsilon,+}\mathbf{L}_{\varepsilon,+}f,f\rangle_\mu^{1/2}}+\frac{\langle \mathbf{L}^*_{\varepsilon,-}\mathbf{L}_{\varepsilon,-}f,f\rangle_\mu}{||f||_\mu\langle \mathbf{L}^*_{\varepsilon,-}\mathbf{L}_{\varepsilon,-}f,f\rangle_\mu^{1/2}}\right):\langle f,\mathbf{1}\rangle_\mu=0\right\}$$
$$=\max_{f\in L^2(X,\mu)}\left\{\frac{1}{2}\left(\left(\frac{\langle \mathbf{L}^*_{\varepsilon,+}\mathbf{L}_{\varepsilon,+}f,f\rangle_\mu}{||f||_\mu^2}\right)^{1/2}+\left(\frac{\langle \mathbf{L}^*_{\varepsilon,-}\mathbf{L}_{\varepsilon,-}f,f\rangle_\mu}{||f||_\mu^2}\right)^{1/2}\right):\langle f,\mathbf{1}\rangle_\mu=0\right\}$$
$$\leq\max_{f\in L^2(X,\mu)}\left\{\left(\frac{1}{2}\left(\frac{\langle \mathbf{L}^*_{\varepsilon,+}\mathbf{L}_{\varepsilon,+}f,f\rangle_\mu}{||f||_\mu^2}+\frac{\langle \mathbf{L}^*_{\varepsilon,-}\mathbf{L}_{\varepsilon,-}f,f\rangle_\mu}{||f||_\mu^2}\right)\right)^{1/2}:\langle f,\mathbf{1}\rangle_\mu=0\right\}$$
$$=\max_{f\in L^2(X,\mu)}\left\{\left(\frac{\langle \mathbf{Q}_\varepsilon f,f\rangle_\mu}{||f||_\mu^2}\right)^{1/2}:\langle f,\mathbf{1}\rangle_\mu=0\right\}.$$

The operator $\mathbf{Q}_\varepsilon$ is self-adjoint, compact (by Theorem 1, noting that duals, compositions, and sums of compact operators are compact), and positive (since $\mathbf{L}^*_{\varepsilon,+}\mathbf{L}_{\varepsilon,+}$ and $\mathbf{L}^*_{\varepsilon,-}\mathbf{L}_{\varepsilon,-}$ are positive, and sums of positive operators are positive). Assum-

ing λ_1^+ is simple, by corollary 1 the value of (26) is $\lambda_2^{1/2}$, the square root of second largest eigenvalue of $\mathbf{Q}_\varepsilon := (\mathbf{L}_{\varepsilon,+}^* \mathbf{L}_{\varepsilon,+} + \mathbf{L}_{\varepsilon,-}^* \mathbf{L}_{\varepsilon,-})/2$ and the maximising f is u_2, the second eigenvector of $\mathbf{Q}_\varepsilon$. The corresponding maximising g and h are $\mathbf{L}_{\varepsilon,+}u_2/\|\mathbf{L}_{\varepsilon,+}u_2\|_{\nu_\varepsilon}$ and $\mathbf{L}_{\varepsilon,-}u_2/\|\mathbf{L}_{\varepsilon,-}u_2\|_{\eta_\varepsilon}$, respectively. As (26) is a relaxation of (24), the equalities (23)–(24) prove:

Theorem 5.

$$\max_{A\subset X, B\subset Y_\varepsilon, C\subset Z_\varepsilon} \{\rho(A,B,C) : \mu(A) = \nu_\varepsilon(B) = \eta_\varepsilon(C)\} \leq 1 + \lambda_2^{1/2}.$$

To construct a partition A, A^c from the solution to (26), we set $A = A_\beta := \{x \in X : u_2(x) > \beta\}$, $A^c = A_\beta^c := \{x \in X : u_2(x) \leq \beta\}$, $B = B_{\beta'} := \{y \in Y_\varepsilon : \mathbf{L}_{\varepsilon,+}u_2(y) > \beta'\}$, $B^c = B_{\beta'}^c := \{y \in Y_\varepsilon : \mathbf{L}_{\varepsilon,+}u_2(y) \leq \beta'\}$, $C = C_{\beta''} := \{z \in Z_\varepsilon : \mathbf{L}_{\varepsilon,-}u_2(z) > \beta''\}$, $C^c = C_{\beta''}^c := \{z \in Z_\varepsilon : \mathbf{L}_{\varepsilon,-}u_2(z) \leq \beta''\}$ where β is chosen to maximise $\rho(A_\beta, B_{\beta'}, C_{\beta''}))$ and $\beta' = \beta'(\beta)$, $\beta'' = \beta''(\beta)$ are chosen so that $\mu(A_\beta) = \nu_\varepsilon(B_{\beta'}) = \eta_\varepsilon(C_{\beta''})$. In practice, the optimal value of β is found via a line search; see Section 6.5 for details.

5 Further discussion

Having completed the description of the three dynamical setups, we now discuss some further properties, focussing mainly on the similarities and differences of Sections 4.2 and 4.3.

5.1 Single- vs. bi-directional coherence

One of the main features that we demonstrate in the numerical case studies is that coherent sets in the sense of Section 4.2 typically have boundaries that are approximately aligned along stable and unstable manifolds. In particular, the boundaries of the coherent sets at the initial time are aligned with stable directions and those at the final time are aligned with unstable directions. Why should this be the case? If the advective dynamics is invertible, the only way that mixing can occur is via diffusion. Therefore, an efficient way for a set to be rapidly mixed over a finite time interval would be for the set to be stretched into long filaments, thus greatly increasing the length of its boundary, allowing diffusion to have a much greater effect. Alternatively, choosing a set at the initial time that already has a very long boundary would also enhance mixing as the initial diffusion would have a large effect.

Optimally coherent sets resist this diffusive mixing by having short boundaries at both the initial and final times. Intuitively this is accomplished by the boundary of the initial set being mostly roughly aligned with stable directions; thus, under forward time evolution, these parts of the boundary are not stretched much, and when the final time is reached, they have evolved so as to roughly align with unstable di-

rections. Symmetrically, one can take the backward time viewpoint: coherent sets at the final time should have small boundary and evolve backward in time to sets with small boundary; this is achieved by the sets at the final time being roughly aligned with unstable directions. The length of the finite-time interval under consideration governs how well the boundaries are aligned with stable/unstable manifolds (the latter are time-asymptotic objects); the longer the interval, the stronger the alignment. In fact, we will show in Section 7 that the optimally coherent sets depend on the time interval considered, as they should.

We have specifically chosen case studies that have regions of strong hyperbolicity to illustrate this point. When the coherence is due to elliptic-type dynamics, as eg. in the polar vortex example in [19] or the Agulhas rings in [15], this phenomenon is not observed.

5.2 Creating a sequence of finite-time coherent sets

In applications, one may be interested in sets that are coherent over a finite-time duration of length τ, and in *tracking* such sets over a time horizon of *several multiples of* τ. If one were to compute coherent sets using the techniques of Section 4.2 from time $t-\tau$ to t, and then from t to $t+\tau$, the sets obtained at t would not match. This is because the discussion in the previous subsection indicates that the first experiment would yield a set at t with boundary roughly aligned with unstable directions, while the second experiment would yield a set at t with boundary roughly aligned with stable directions. The construction of Section 4.3, on the other hand, computing over the window $t-\tau$ to $t+\tau$ finds a "central" set at time t with the property that its boundary does not grow large in either backward or forward time. By applying the method of Section 4.3 to a series of windows of length 2τ, one obtains a sequence of such sets which should vary continuously with time. This is one of the main motivations for the constructions in Section 4.3.

A similar effect result could be achieved using the method of Section 4.2 in the following way:

1. Compute the transfer operator for the period $t-\tau$ to $t+\tau$, and establish optimally coherent sets at $t-\tau$ and $t+\tau$ via thresholding.
2. Push the left singular vector (at $t-\tau$) forward with the transfer operator from $t-\tau$ to t and threshold according to the conservation of mass principle.

By pushing forward to the intermediate time t, the sets obtained should also have the property that they have short boundaries when pushed forward to $t+\tau$ or pulled back to $t-\tau$. We will explore this possibility in future work.

5.3 Further mathematical properties of the coherent set framework

In [13] several properties of the operator $\mathbf{L}_\varepsilon$ from Section 4.2 are proved. We briefly mention some of these properties here in the situation where one uses ε-ball diffusion for $\alpha_{X,\varepsilon}$ and $\alpha_{Y,\varepsilon}$. These properties are dealt with in greater detail and generality in [13].

1. The analytical framework for identifying finite-time coherent sets based on the second singular vectors of $\mathbf{L}_\varepsilon$ is *frame-invariant* or *objective*. This means if the framework is applied in a general time-dependent rotating and translating frame, the coherent sets obtained will be the same (except rotated and translated) as those obtained in a static frame. The issue of frame-invariance does not apply to Section 4.1 as the dynamics should not be time-varying. The arguments in [13] could be adapted to the setting of Section 4.3.
2. If T is a diffeomorphism and $X = Y_\varepsilon = M$, the leading singular value $\sigma_1 = 1$ of $\mathbf{L}_\varepsilon$ is simple. Arguments similar to those in [13] could be applied to the material in Section 4.1 and 4.3 to demonstrate simplicity of λ_1.
3. If T is a diffeomorphism and $X = Y_\varepsilon = M$, a lower bound on the second singular value $\sigma_{2,\varepsilon} < 1$ is given in [13], depending on ε. The techniques in [13] could be adapted to the constructions in Section 4.1 and 4.3.
4. If T is a diffeomorphism and $X = Y_\varepsilon = M$, subdominant singular vectors of $\mathbf{L}_\varepsilon$ are 1/2-Hölder regular, with the Hölder constant having an explicit dependence on ε (larger ε, smaller constant). In particular, this places some limitations on the geometric shapes of the optimally coherent sets. One could apply the techniques in [13] to the constructions in both Section 4.1 and 4.3.

Regarding point 3. above, in the case of autonomous two-dimensional area-preserving maps T, Junge *et al.* [27] state that the probability to map out of a T-*invariant* set by an ε-perturbed systems is bounded to first order from above by ε (i.e. the amplitude of perturbation) multiplied by the ratio of set boundary length to set volume. They also state a lower bound on the second largest eigenvalue λ_2 of $\mathbf{R}_\varepsilon$; thus for a fixed ε if the invariant set of the unperturbed system has small boundary and its (normalized) Lebesgue measure is close to 1/2, then λ_2 is closer to 1. In [27], the self-adjoint operator $\mathbf{R}_\varepsilon$ is constructed similarly to ours, but using a one-sided diffusion only (i.e. akin to $(\mathbf{D}_{X,\varepsilon}\mathbf{P} + (\mathbf{D}_{X,\varepsilon}\mathbf{P})^*)/2$).

6 Numerical representations of transfer operators

In order to apply the theory developed in the previous sections to specific mathematical models, we require a computer representation of the operators $\mathbf{P}, \mathbf{L}, \mathbf{P}_\varepsilon$, and $\mathbf{L}_\varepsilon$. We recall here the standard approach of Ulam [45], adapted to our specific operator constructions. We represent these operators as a projected action on a finite-cardinality basis of characteristic functions. Let $\{B_1, \ldots, B_m\}$ denote a partition of X and define $\pi_{X,m} : L^1(X, \ell) \to \mathrm{sp}\{\mathbf{1}_{B_1}, \ldots, \mathbf{1}_{B_m}\}$ by

$$\pi_{X,m} f = \sum_{i=1}^{m} \left(\frac{1}{\ell(B_i)} \int_{B_i} f \, d\ell \right) \mathbf{1}_{B_i}.$$

Similarly, let $\{C_1, \ldots, C_n\}$ denote a partition of Y_ε and define $\pi_{Y_\varepsilon,n} : L^1(Y_\varepsilon, \ell) \to \mathrm{sp}\{\mathbf{1}_{C_1}, \ldots, \mathbf{1}_{C_n}\}$ by

$$\pi_{Y_\varepsilon,\ell} f = \sum_{j=1}^{n} \left(\frac{1}{\ell(C_j)} \int_{C_j} f \, d\ell \right) \mathbf{1}_{C_j},$$

where ℓ is Lebesgue measure. If μ, ν_ε are absolutely continuous with respect to ℓ, and the maximal diameter of the partition elements decrease to zero as $m,n \to \infty$, then $\pi_{X,m}, \pi_{Y_\varepsilon,n}$ converge strongly to the identity operator in $L^p(X,\mu), L^p(Y_\varepsilon, \nu_\varepsilon)$, $p = 1,2$ (see eg. Prop. 9, Chap. 6 [41]).

6.1 Numerically representing **P** *and* **L**

We consider the operator $\pi_{Y_\varepsilon,n} \mathbf{P} \pi_{X,m} : \mathrm{sp}\{\mathbf{1}_{B_1}, \ldots, \mathbf{1}_{B_m}\} \to \mathrm{sp}\{\mathbf{1}_{C_1}, \ldots, \mathbf{1}_{C_n}\}$, which has matrix representation[10]:

$$P_{ij} = \frac{\ell(B_i \cap T^{-1}(C_j))}{\ell(C_j)}. \tag{27}$$

In what follows, it will be useful to consider the related matrix

$$\bar{P}_{ij} = \frac{\ell(B_i \cap T^{-1}(C_j))}{\ell(B_i)}, \tag{28}$$

which may be considered as a discrete action of T on measures. $\bar{P}$ is row-stochastic, has leading eigenvalue 1, and has an interpretation as a transition matrix where the entry $\bar{P}_{ij}$ represents the conditional probability that a randomly chosen point in B_i lands in C_j after one application of T.

In what follows, we use the shorthand $\ell^B = [\ell(B_1), \ldots, \ell(B_m)]$ and $\ell^C = [\ell(C_1), \ldots, \ell(C_n)]$. Given μ, let $p_i = \mu(B_i)$ and we approximate h_μ as $\sum_{i=1}^m (p_i/\ell_i^B)\mathbf{1}_{B_i}$; from now on for brevity, we drop this functional representation for densities and measures and write them as vectors. The image density $h_{\nu_\varepsilon} = \mathbf{P}_\varepsilon h_\mu$ is estimated as $\sum_{i=1}^m (p_i/\ell_i^B) P_{ij}$ and the image measure ν_ε is estimated as $\sum_{i=1}^m (p_i/\ell_i^B) P_{ij} \ell_j^C = \sum_{i=1}^m p_i \bar{P}_{ij} =: q_j$. Thus, to construct an approximate matrix representation for $\mathbf{L}$, we use the definition (2) to obtain

[10] Li [30] contains the first statement of this result in the context of interval maps, but is straightforward to derive using the property that for each measurable $A \subset Y_\varepsilon$, $\int_A \mathbf{P} f \, d\ell = \int_{T^{-1}A} f \, d\ell$ for all $f \in L^1(X,\ell)$.

$$\begin{aligned} L_{ij} &= \frac{(p_i/\ell_i^B)P_{ij}}{\sum_{i=1}^m (p_i/\ell_i^B)P_{ij}} \\ &= \frac{(p_i/\ell_i^B)\bar{P}_{ij}(\ell_i^B/\ell_j^C)}{\sum_{i=1}^m (p_i/\ell_i^B)\bar{P}_{ij}(\ell_i^B/\ell_j^C)} \\ &= \frac{p_i\bar{P}_{ij}}{\sum_{i=1}^m p_i\bar{P}_{ij}}. \end{aligned} \tag{29}$$

This latter expression appeared in [19]. It is clear that $\mathbf{1}L = \mathbf{1}$. Denoting the inner products $\langle x,y\rangle_p$ and $\langle x,y\rangle_q$ by $\sum_{i=1}^m x_i y_i p_i$ and $\sum_{i=1}^n x_i y_i q_i$ respectively, it is straightforward to check that L^*, the matrix dual satisfying $\langle xL,y\rangle_q = \langle x,yL^*\rangle_p$ is $\bar{P}^\top$, and that $\mathbf{1}L^* = \mathbf{1}$.

In the autonomous setting, recall one has $X = Y_\varepsilon$ and $\mu = \nu_\varepsilon$ and μ should be T-invariant. To estimate the T-invariant μ, we use the leading eigenvector of $\bar{P}$, ie. choose p to be the (assumed unique) vector satisfying $p = p\bar{P}$. The expression for L now becomes

$$L_{ij} = \frac{p_i\bar{P}_{ij}}{p_j}, \tag{30}$$

and is a discrete approximation of $\mathbf{L} : L^2(X,\mu) \to L^2(X,\mu)$; in fact, in this setting L is nothing but a discrete approximation of the Perron-Frobenius operator $\mathbf{P}_\mu : L^1(X,\mu) \to L^1(X,\mu)$ defined with respect to μ, rather than ℓ. In the autonomous setting, there have been a number of papers that discuss choosing partitions in a way that Ulam's method produces the most accurate estimate of the physical invariant measure μ for T [20, 26, 36]. For the purposes of finding almost-invariant and coherent sets, as numerical diffusion plays an important role, we advise choosing partition sets that are approximately spherically symmetric (eg. squares or cubes) so that the implicit numerical diffusion that results is approximately isotropic.

We remark that the matrix representation (29) is not the same as the matrix representation of $\pi_{X,m}\mathbf{L}_0\pi_{Y,n}$; that is, L is not exactly a Galerkin projection of $\mathbf{L}_0$. We have chosen this alternative formulation for numerical convenience and do not believe that the numerical impact is great. To estimate the entry $\bar{P}_{ij}$ numerically, one may sample test points $x_{i,k}$, $k = 1,\ldots,K$ uniformly distributed over B_i (eg. on a uniform grid) and then compute $T(x_{i,k})$ and count how many fall in C_j; that is,

$$\bar{P}_{ij} \approx \frac{\#\{k : T(x_{i,k}) \in C_j\}}{K}. \tag{31}$$

The software package GAIO [7] is used to estimate the transition matrix entries. GAIO uses generalized rectangles (*boxes*) as partition elements and, using a multi-level tree-like data structure, can efficiently find which boxes contain image points.

6.2 Numerical representation of $\mathbf{P}_\varepsilon$ and $\mathbf{L}_\varepsilon$

Although in the previous section we have constructed a matrix representation of $\mathbf{L}_0$, not $\mathbf{L}_\varepsilon$ for some $\varepsilon > 0$, in fact, as a consequence of the discretisation we have already implicitly incorporated a low level of numerical diffusion of the order of the diameter of the partition elements. We now discuss two ways to construct a matrix representation of $\mathbf{L}_\varepsilon$, explicitly including diffusion of the type governed by $\mathbf{D}_{X,\varepsilon}$ and $\mathbf{D}_{Y'_\varepsilon,\varepsilon}$.

1. For each test point $x_{i,k} \in B_i$, $k = 1,\dots,K$, represent the diffusion over an ε-ball centred at $x_{i,k}$ by a second grid of points $y_{i,k,k'}, k' = 1,\dots,K'$ uniformly spread over $B_\varepsilon(x_{i,k})$. For example, one could select $y_{i,k,k'}$ on a grid centred at $x_{i,k}$; in terms of computer code, the most efficient approach is to create a "mask" of such points for an ε-ball centred at zero and merely add the vector $x_{i,k}$ to translate this fixed "mask". Now compute $T(y_{i,k,k'}), k = 1,\dots,K, k' = 1,\dots,K'$; in total this represents $K \cdot K'$ points for each box B_i. Finally, for each image point $T(y_{i,k,k'})$ we again create K'' points uniformly distributed in a ball of radius ε, centred at $T(y_{i,k,k'})$; call this final set of $K \cdot K' \cdot K''$ points $z_{i,k,k',k''}$. This final set of points can again be created via the "mask" procedure described above; now one adds the vector $T(y_{i,k,k'})$. Finally, we estimate

$$\bar{P}_{\varepsilon,i,j} = \frac{\#\{z_{i,k,k',k''} \in C_j\}}{K \cdot K' \cdot K''}.$$

   This was the approach taken in [13]. It has a high accuracy because it directly simulates the concatenation $\mathbf{D}_{Y'_\varepsilon,\varepsilon}\mathbf{P}\mathbf{D}_{X,\varepsilon}$ via test points and effectively applies a discretisation only once.
2. While the above approach is cheap from a memory point of view, there is an overhead to computing T-images of $K \cdot K'$ points. A faster (but somewhat less accurate) approach would be to discretise each of the three operators $\mathbf{D}_{Y'_\varepsilon,\varepsilon}, \mathbf{P}, \mathbf{D}_{X,\varepsilon}$ separately and then estimate their product by matrix multiplication. Such an approach is less accurate because the effects of the three discretisations are multiplied together, however if ε is much larger than the box diameters, the error should be comparatively small. An advantage to separately discretising is that one can try different diffusion amplitudes without having to recompute the discretisation $\mathbf{P}$.
   Let $\{B_1,\dots B_m\}$ denote a partition of X. We first set up the matrix $\bar{P}$ without diffusion as described in section 6.1. We construct a matrix $D_{X,\varepsilon}$ representing a discretised version of $\mathbf{D}_{X,\varepsilon}$ as follows. In each box B_i we choose K test points $x_{i,k} \in B_i$, $k = 1,\dots,K$. For each $x_{i,k} \in B_i$ we represent the diffusion over an ε-ball centred at $x_{i,k}$ by a second grid of points $y_{i,k,k'}$, $k' = 1,\dots,K'$, uniformly spread over $B_\varepsilon(x_{i,k})$ and estimate

$$D_{X,\varepsilon,i,j} = \frac{\#\{y_{i,k,k'} \in B_j\}}{K \cdot K'}.$$

We similarly construct a matrix $D_{Y'_\varepsilon,\varepsilon}$ based on sets $\{C_1,\ldots,C_n\}$. The matrix P_ε is then obtained as $P_\varepsilon = D_{Y'_\varepsilon,\varepsilon} P D_{X,\varepsilon}$.

In this paper we take an even faster and coarser approach whereby a ball of radius ε is approximated as a square or cube of side-length 2ε. This approach is faster as it makes use of the internal data structure of GAIO [7].

In the following sections we briefly describe how to appropriately put together the matrices L. From now on, we drop the ε subscript for the matrices.

6.3 Autonomous setting

In the autonomous setting, we construct $L_{ij} = p_i \bar{P}_{ij}/p_j$, where the map T used to construct $\bar{P}$ represents the dynamics over the fixed time duration we are interested in. As we are in the autonomous setting, only the duration matters, not the initial time. We consider

$$Q := (L + L^*)/2; \tag{32}$$

note that Q is self-adjoint in $\langle\cdot,\cdot\rangle_p$ (but is not a symmetric matrix in general), and that $\mathbf{1}Q = \mathbf{1}Q^* = \mathbf{1}$. By Corollary 1 the solution to (13) is given by the second left eigenvector u_2^+ of $\mathbf{Q}$; we will numerically estimate u_2^+ as the second left eigenvector of Q, normalised so that $\langle u_2^+, u_2^+\rangle_p = 1$. If there is strong almost-invariance present, we expect the values of u_2^+ to be around ± 1. Note that Froyland [11] proposed the use of the second *right* eigenvector of $(\bar{P}+\hat{P})/2$ to obtain almost-invariant sets (where $\hat{P}$ denotes the transition matrix for the time-reversed Markov chain governed by $\bar{P}$) As $(\bar{P}+\hat{P})/2 = Q^\top$, our use of the second *left* eigenvector of Q yields an identical result. Here we have incorporated the autonomous and nonautonomous constructions under a single unified set of constructions, and also have demonstrated how to naturally incorporate diffusive aspects of dynamics.

The aim is to find an optimal partition of X. We restate the algorithm as used in [18] for finding almost invariant sets using R.

Algorithm 1 (Almost invariant sets)

1. *Partition the state space X into a collection of connected sets $\{B_1,\ldots,B_m\}$ of small diameter.*
2. *Construct the Ulam matrix $\bar{P}$ using (31) and compute the (assumed unique) fixed left eigenvector p of $\bar{P}$. If explicit diffusion is to be added, use one of the methods in Section 6.2.*
3. *Construct the matrix L using (30) and Q using (32).*
4. *Compute the second largest eigenvalue $\lambda_2 < 1$ of Q and corresponding left eigenvector u_2, normalised so that $\langle u_2, u_2\rangle_p = 1$.*
5. *Denote $I(b) = \{i : u_{2,i} > b\}, I^c(b) = \{i : u_{2,i} \le b\}$. Perform a line search on b to maximise*[11]

[11] The expression (33) is a discrete form of (8).

$$\frac{\sum_{i\in I(b),j\in I(b)} p_i \bar{P}_{ij}}{\sum_{i\in I(b)} p_i} + \frac{\sum_{i\in I^c(b),i\in I^c(b)} p_i \bar{P}_{ij}}{\sum_{i\in I^c(b)} p_i}. \tag{33}$$

6. *Denote by $\hat{b}$ the optimal b and set $\hat{A} := \cup_{i\in I(\hat{b})} B_i, \hat{A}^c := \cup_{i\in I(\hat{b})^c} B_i$.*

6.4 Nonautonomous setting: single direction

In this setting we wish to study transport over the time interval $[t_0, t_1]$ and build the matrix $\bar{P}$ using a map T that describes the dynamics over this interval. We then construct L using (29). Proposition 1 [13] shows that the value of (19) is $\lambda_2^{1/2}$, the square root of the second largest eigenvalue of $\mathbf{Q}_\varepsilon := \mathbf{L}_\varepsilon^* \mathbf{L}_\varepsilon$ and the maximising f (resp. g) is u_2^+ (resp. v_2), the corresponding left eigenvector of $\mathbf{Q}_\varepsilon$ (resp. $\mathbf{Q}_\varepsilon^*$). One could define a matrix approximation of $\mathbf{Q}$ as $Q := LL^*$, and a matrix approximation of $\mathbf{Q}^*$ as $Q^* := L^*L$, however, it is more efficient numerically to find left and right singular vectors of L directly. The reason for this is that L is reasonably sparse, especially for small ε, and LL^* may be sigificantly more dense than the matrices involved in a calculation of singular value of L. The following algorithm was put forward in [19].

Algorithm 2 (Finite-time coherent sets, single time direction)

1. *Partition the domains X and Y_ε into a collection of connected sets $\{B_1,\dots,B_m\}$ and $\{C_1,\dots,C_n\}$ respectively, of small diameter. If Y_ε is not known precisely, then set Y_ε to be a neighbourhood of $T(X)$ that contains all possible perturbed image points.*
2. *Select the reference measure μ as the mass distribution to be tracked, and set $p_i = \mu(B_i)$.*
3. *Construct the Ulam matrix $\bar{P}$ as in (31), and compute $q = p\bar{P}$. If additional explicit diffusion is used, then use one of the approaches in Section 6.2.*
4. *Define diagonal matrices $(\Sigma_p)_{ii} = p_i$, $i = 1,\dots,m$, and $(\Sigma_q)_{jj} = q_j$, $j = 1,\dots,n$, compute the second largest singular value $\sigma_2 < 1$ of $\Sigma_p^{1/2}\bar{P}\Sigma_q^{-1/2}$ and corresponding left and right singular vectors $\tilde{u}_2, \tilde{v}_2$, and set $u_2 := \tilde{u}_2\Pi_p^{-1/2}$, $v_2 := \tilde{v}_2\Pi_q^{-1/2}$.*
5. *Denote $I(b) = \{i : u_{2,i} > b\}, I^c(b) = \{i : u_{2,i} \le b\}, J(b') = \{j : v_{2,j} > b'\}, J^c(b') = \{j : v_{2,j} \le b'\}$. Perform a line search on b to maximise*[12]

$$\frac{\sum_{i\in I(b),j\in J(b')} p_i \bar{P}_{ij}}{\sum_{i\in I(b)} p_i} + \frac{\sum_{i\in I^c(b),j\in J^c(b')} p_i \bar{P}_{ij}}{\sum_{i\in I^c(b)} p_i}, \tag{34}$$

   selecting $b' = b'(b)$ so that $|\sum_{i\in I(b)} p_i - \sum_{j\in J(b')} q_j|$ is minimised[13] *for each choice of b.*

[12] The expression (34) is a discrete form of (14).

[13] This is the discrete version of insisting that $\mu(A(b)) = \nu_\varepsilon(B(b'))$.

6. *Denote by* $\hat{b},\hat{b}'$ *the optimal* b,b' *and set* $\hat{A} := \cup_{i\in I(\hat{b})} B_i, \hat{A}^c := \cup_{i\in I(\hat{b})^c} B_i, \hat{B} := \cup_{j\in J(\hat{b}')} C_j, \hat{B}^c := \cup_{j\in J(\hat{b}')^c} C_j$.

In the numerical case studies section we will frequently plot the output of item 4 above, namely the vectors u_2 and v_2, and also the output of item 6 above, namely the sets $\hat{A},\hat{A}^c,\hat{B}$, and $\hat{B}^c$. If the vectors $\tilde{u}_2,\tilde{v}_2$ have ℓ^2-norm[14] 1, then the vectors u_2,v_2 will be normalised so that $\langle u_2,u_2\rangle_p = 1$ and $\langle v_2,v_2\rangle_q = 1$. If there is strong almost-invariance present, we expect the values of u_2 and v_2 to be around ± 1.

6.5 Nonautonomous setting: both directions

In this setting we are interested in sets at time point t that remain coherent both in forward and backward times. We build matrices $\bar{P}_+$ and $\bar{P}_-$ using maps T_+ and T_- that describe the dynamics from t to $t+\tau_2$ and from t to $t-\tau_1$, respectively. We then construct L_+, L_- using (29) and consider the matrix approximation to $\mathbf{Q}_\varepsilon$: $Q := (L_+L_+^* + L_-L_-^*)/2$ (recall we always use left multiplication). We propose the following algorithm:

Algorithm 3 (Finite-time coherent sets, both time directions)

1. *Partition the domains* X, Y_ε *and* Z_ε *into a collection of connected sets* $\{B_1,\ldots,B_m\}$, $\{C_1,\ldots,C_n\}$ *and* $\{E_1,\ldots,E_o\}$ *respectively, of small diameter. If* Y_ε *and* Z_ε *are not known precisely, then set* Y_ε *to be a neighbourhood of* $T_+(X)$ *that contains all possible perturbed image points, likewise* Z_ε *a neighbourhood of* $T_-(X)$.
2. *Select the reference measure* μ *at time* t *as the mass distribution to be tracked, and set* $p_i = \mu(B_i)$.
3. *Construct the Ulam matrix* $\bar{P}_+$ *as in (31) using* $T = T_+$ *and compute* $q^+ = p\bar{P}_+$. *To construct* $\bar{P}_-$, *use* $T = T_-$ *and replace* C_j *with* E_j *in (31). Set* $q^- = p\bar{P}_-$. *If additional explicit diffusion is used, then use one of the approaches in Section 6.2.*
4. *Construct* L_+ *as in (29) using* $\bar{P}_+$ *and* p, *and* L_- *using* $\bar{P}_-$ *and* p. *Form* $Q = (L_+L_+^* + L_-L_-^*)/2$ *by matrix multiplication where* $L_+^* = \bar{P}_+^\top$ *and* $L_-^* = \bar{P}_-^\top$.
5. *Compute the second largest eigenvalue* $\lambda_2 < 1$ *of* Q *and corresponding left eigenvector* u_2. *Set* $v_2^+ = u_2L_+$ *and* $v_2^- = u_2L_-$. *Normalise* u_2 *so that* $\langle u_2,u_2\rangle_p = 1$ *and normalise* $v_2^\pm$ *so that* $\langle v_2^\pm, v_2^\pm\rangle_{q^\pm} = 1$.
6. *Denote* $I(b) = \{i : u_{2,i} > b\}, I^c(b) = \{i : u_{2,i} \le b\}, J(b') = \{j : v_{2,j}^+ > b'\}, J^c(b') = \{j : v_{2,j}^+ \le b'\}$ *and* $K(b'') = \{l : v_{2,l}^- > b''\}, K^c(b'') = \{l : v_{2,l}^- \le b''\}$. *Perform a line search on* b *to maximise*[15]

$$\frac{\sum_{i\in I(b),j\in J(b')} p_i\bar{P}_{ij}^+}{\sum_{i\in I(b)} p_i} + \frac{\sum_{i\in I^c(b),j\in J^c(b')} p_i\bar{P}_{ij}^+}{\sum_{i\in I^c(b)} p_i} + \frac{\sum_{i\in I(b),k\in K(b'')} p_i\bar{P}_{ik}^-}{\sum_{i\in I(b)} p_i} + \frac{\sum_{i\in I^c(b)k\in K^c(b'')} p_i\bar{P}_{ik}^-}{\sum_{i\in I^c(b)} p_i}, \tag{35}$$

[14] This is the default output normalisation for MATLAB, for example.

[15] The expression (35) is a discrete form of (24).

selecting $b' = b'(b)$ *and* $b'' = b''(b)$ *so that* $|\sum_{i\in I(b)} p_i - \sum_{j\in J(b')} q_j^+|$ *and* $|\sum_{i\in I(b)} p_i - \sum_{k\in K(b'')} q_k^-|$ *are minimised*[16] *for each choice of* b.

7. *Denote by* $\hat{b}, \hat{b}', \hat{b}''$ *the optimal* b, b', b'' *and set* $\hat{A} := \cup_{i\in I(\hat{b})} B_i, \hat{A}^c := \cup_{i\in I^c(\hat{b})} B_i, \hat{B} := \cup_{j\in J(\hat{b}')} C_j, \hat{B}^c := \cup_{j\in J^c(\hat{b}')} C_j$ *and* $\hat{C} := \cup_{k\in K(\hat{b}'')} E_k, \hat{C}^c := \cup_{k\in K^c(\hat{b}'')} E_k$.

In the numerical case studies section we will frequently plot the output of item 5 above, namely the vectors u_2 , v_2^+ and v_2^-, and also the output of item 7 above, namely the sets $\hat{A}, \hat{A}^c, \hat{B}, \hat{B}^c, \hat{C}$ and $\hat{C}^c$. If there is strong coherence present, we again expect the entries of the vectors u_2, v_2^+, and v_2^- to be around ± 1.

7 Numerical examples

In this section we will apply the different constructions to two well-known example systems. First we consider a periodically driven double gyre flow [43], which has frequently been used as a testbed for different tools for the numerical analysis of transport. Due to the system's periodicity we will analyse the system both with respect to almost-invariant sets as well as finite-time coherent sets and point out the differences of the constructions. We discuss the effects of diffusion for almost-invariant sets, and for coherent sets, additionally time direction and flow duration.

As a second example system we consider the transitory double gyre flow as introduced in [35]. Here the dynamics is only nonautonomous on a finite time interval, but autonomous outside. Therefore the system is well-suited to analysing finite-time coherent structures. Special emphasis will be placed on how the structure of finite-time coherent sets depends on the time direction and flow duration.

While in these examples the domain at the initial time and final time remains the same, the coherent set framework also easily handles situations where the initial domain and final domain do not intersect at all, as in eg. [15].

7.1 Case study 1: Periodically driven double gyre flow

We consider the time-dependent system of differential equations [43]

$$\begin{aligned} \dot{x} &= -\pi A \sin(\pi f(x,t)) \cos(\pi y) \\ \dot{y} &= \pi A \cos(\pi f(x,t)) \sin(\pi y) \frac{df}{dx}(x,t), \end{aligned} \tag{36}$$

where $f(x,t) = \delta \sin(\omega t) x^2 + (1 - 2\delta \sin(\omega t)) x$.

For detailed discussions of the system we refer to [43, 18]. As in [18] in we fix parameter values $A = 0.25$, $\delta = 0.25$ and $\omega = 2\pi$ and obtain a flow of period

[16] This is the discrete version of insisting that $\mu(A(b)) = \nu_\varepsilon(B(b')) = \eta_\varepsilon(C(b''))$.

$p = 1$. The system preserves Lebesgue measure on $[0,2] \times [0,1]$. We partition the domain[17] $M = [0,2] \times [0,1]$ in $n = 32768 = 2^{15}$ square boxes. Here we will identify $M = X = Y_\varepsilon = Z_\varepsilon$. We form matrices $\bar{P}$ by integrating with a Runge Kutta scheme with constant stepsize $h = 0.01$ from $t = 0$ over different time spans, i.e. over one period ($\tau = 1$) and $\tau = \pm 2.5$, using $K = 400$ uniformly distributed test points per box (inner grid[18] points). In this set-up, with box diameter 0.0078, we have a numerically induced diffusion of about $\varepsilon \approx 0.0039$. In addition, we consider explicit diffusion via left and right multiplication of $\bar{P}$ with diffusion matrices $D_{X,\varepsilon}$, $D_{Y'_\varepsilon,\varepsilon}$. Here we choose 100 inner grid points per box, for each test point we approximate a ball of radius $\varepsilon = 0.02$ by 25 inner grid points of a box of diameter 2ε centred in the respective point.

In the following, we will demonstrate the usage of the different constructions and discuss effects of diffusion, flow duration, and time-direction.

Almost-invariant sets

As the system is periodic with period $p = 1$ by construction, the time-1 flow map T describes an autonomous dynamical system. Thus, we want to determine fixed regions in phase space that are almost-invariant under the dynamics of T. To this end, we compute[19] $p = p\bar{P}$ and form matrices L and Q. Following Algorithm 1 we consider the second left eigenvector u_2 to eigenvalue $\lambda_2 = 0.9998$, normalized such that $\langle u_2, u_2 \rangle_p = 1$ (Figure 2 (a)). Optimal almost invariant sets are obtained for $\hat{b} = \pm 0.3801$, giving $\rho(\hat{A}) = 1.9907$, see Figure 1. As indicated, regular phase space structures, i.e. truly invariant sets are picked up.

If these invariant structures are removed from the domain, one finds $\hat{b} = 0$ is now the global optimum; see Figure 2 (b). A similar experiment was carried out in [18]. Note that for the optimal partition, $\rho(\hat{A})$ is within the upper $(1 + \lambda_2 = 1.9998)$ and lower bounds $(2 - 2\sqrt{2(1-\lambda_2)} = 1.96)$.

In Figure 2 (c) and (d) we show the results of explicit diffusion, incorporated by left and right multiplication of diffusion matrices $D_{X,\varepsilon}, D_{Y'_\varepsilon,\varepsilon}$ ($\varepsilon = 0.02$) with $\bar{P}$ as described above. Applying Algorithm 1 to the diffused matrices results in a visibly smoother eigenvector u_2 (Figure 2 (c)) compared to the case without explicit diffusion (Figure 2 (a)). Moreover, the optimal partitions now are given by $\hat{b} = 0$ and the boundaries are smoother and shorter. (Figure 2 (d) compared to Figure 2 (b)). As expected, diffusion regularizes eigenvectors of Q; this in turn usually shortens and smoothens the boundaries of almost invariant sets.

Another important phenomenon that has been illustrated here is an apparent bifurcation of the almost-invariant sets with variation in noise amplitude. When only numerical diffusion was present, the invariant sets shown in Figure 1 were selected

[17] Note that the boundaries of M are invariant.

[18] See [7] for a description of inner grid points.

[19] As Lebesgue measure is preserved, p should give equal weight $1/m$ to each of the m partition sets. However, in order to account for possible sampling-induced numerical inaccuracies when setting up $\bar{P}$, we will use the numerically obtained p.

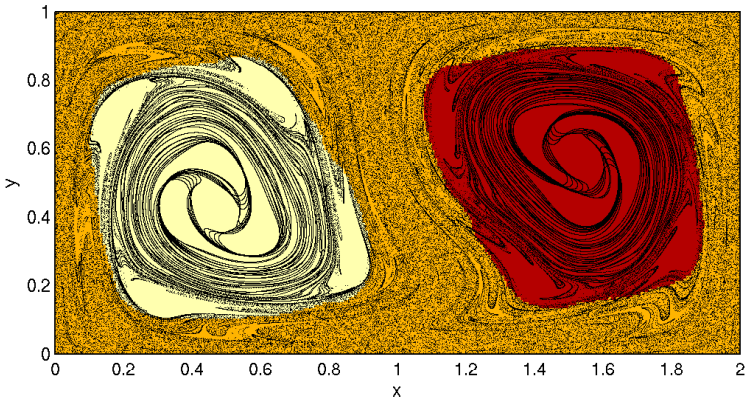

Fig. 1 Almost invariant sets in the double gyre flow on the time interval $[0,1]$ - no explicit diffusion. Due to symmetry in the system, both $\hat{b} = -0.3801$ and $\hat{b} = 0.3801$ are optimal thresholds in Algorithm 1. We obtain $\rho(\hat{A}) = 1.9907$, where $\hat{A}$ is the light ($\hat{b} = -0.3801$) or dark set ($\hat{b} = 0.3801$) and $\hat{A}^c = M \setminus \hat{A}$. In this setting, the almost invariant sets correspond to regular (invariant) structures in phase space.

by Algorithm 1. The addition of explicit diffusion resulted in Algorithm 1 selecting the sets shown in Figure 2 (d). The reason for this switch is that with very low diffusion, advective flux dominates diffusive flux and the invariant sets in Figure 1 minimise transport across their boundaries. As the diffusion amplitude is raised, the optimal almost-invariant structures change to those shown in Figure 2 (d), which have a shorter, non-invariant boundary, leading to a reduction in diffusive flux, but an increase in advective flux; the net effect, is however, lower flux than the invariant sets in Figure 1 under the $\varepsilon = 0.02$ diffusion regime. We will further investigate bifurcation aspects in future work.

In Figure 2 (b) and (d) we have also plotted approximations[20] to stable and unstable manifolds of hyperbolic periodic orbits (oscillating around $x = 1$ on the $y = 0$ and $y = 1$ axes). Apparently, as already discussed in [18], the transfer operator approach finds a decomposition into almost invariant sets, which is influenced by the underlying manifold structure but tries to find a more optimal decomposition than a geometrical approach such as lobe dynamics [39, 40] would suggest. While the stable and unstable manifolds concern time-asymptotic dynamics, the optimal almost-invariant sets identified are tuned to a finite flow time of the dynamics.

Finite-time coherent sets: single time direction

In this section we are no longer interested in spatially fixed sets, but in possibly dynamical regions in phase space that remain coherent during some finite time span.

[20] We refer the reader to [18] for more details.

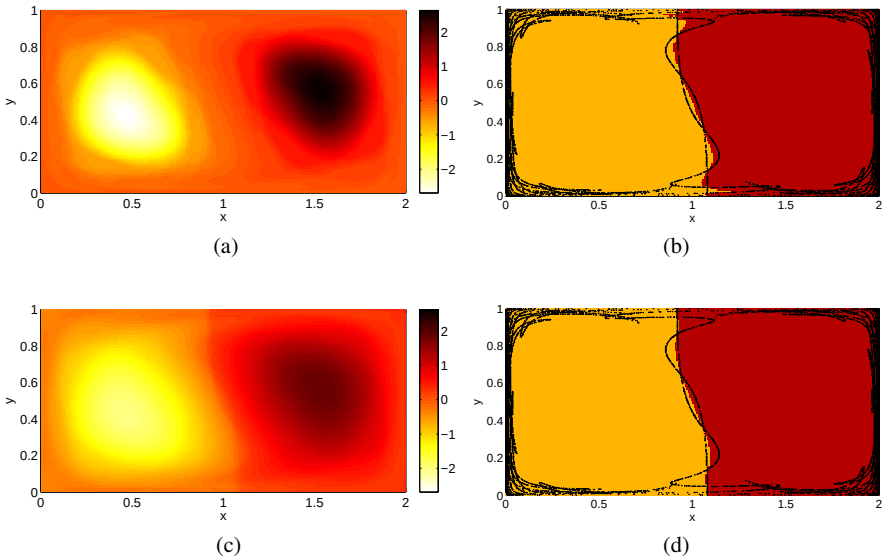

Fig. 2 Second eigenvectors and corresponding almost invariant sets in the double gyre flow on the finite time interval $[0,1]$, with and without explicit diffusion. (a) Second eigenvector u_2 of Q ($\lambda_2 = 0.9998$), $\varepsilon = 0$. (b) Partition by sign results in $\rho(A_0) = 1.9803$. Such a partition would also be obtained when ignoring the invariant, regular structures as visible in Figure 1, see [18] for a related case study. (c,d): Effects of explicit diffusion on second eigenvectors and almost invariant sets in the double gyre flow: (c) Second eigenvector u_2 to eigenvalue $\lambda_2 = 0.9974$ of $Q_\varepsilon = (L_\varepsilon + L_\varepsilon^*)/2$, where $\varepsilon = 0.02$. (d) Optimal partition into almost invariant sets based on u_2 in (c). Here $\hat{b} = 0$ is the optimal threshold, resulting in $\rho(\hat{A}) = 1.9758$.

We will first study finite-time coherent sets obtained from considering the dynamics in a single time direction. We restrict ourselves to the case without explicit diffusion.

As a first case study, we will consider the dynamics on $[0,1]$ such as in the previous paragraph, but apply Algorithm 2 to the respective matrix $\bar{P}$. The resulting (normalized) singular vectors u_2 and v_2 with respect to the singular value$\sigma_2 = 1 - 1.6 \times 10^{-5}$ are shown in Figure 3 (a) and (c). The thresholds $\hat{b} = 0$ and $\hat{b}' = 0$ turn out to define optimal coherent pairs shown in Figure 3 (b) and (d). We obtain $\rho(\hat{A}, \hat{B}) = 1.9976$, which is bounded from above by $1 + \sigma_2$. In addition, we have overlaid the optimal partition in Figure 3 (b) with an approximation to the stable manifold of the periodic orbit on the $y = 0$ axis and the optimal partition in Figure 3 (d) with an approximation to the unstable manifold of the periodic orbit on the $y = 1$ axis. The (asymptotic) geometric structures influence the shape of the coherent sets but due to the short finite-time horizon on which their computation is based on a we do not get a nearly exact correspondence.

We repeat this study by considering the dynamics on $[0,2.5]$ as well as on $[-2.5,0]$, i.e. longer time intervals which are not integer multiples of the period

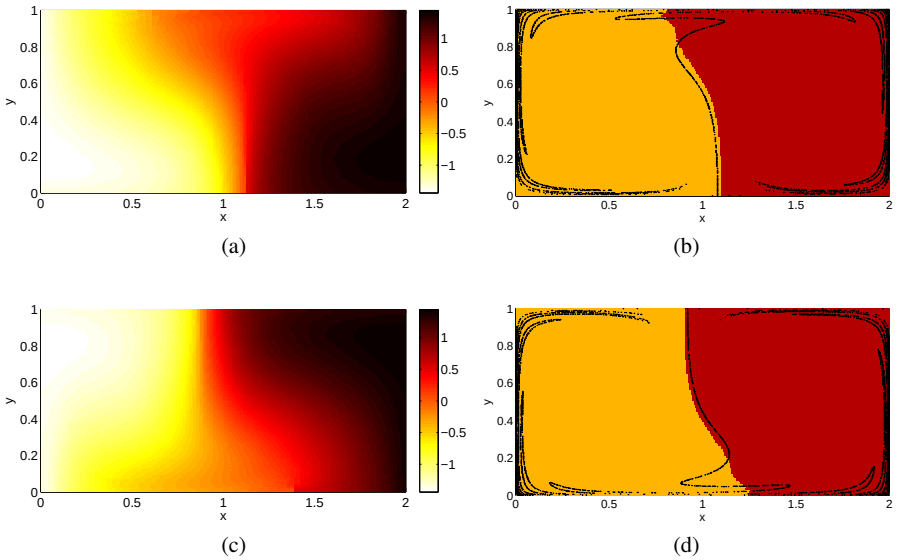

Fig. 3 Singular vectors and finite-time coherent pairs ($\hat{b} = \hat{b}' = 0$ in Algorithm 2) in the double gyre flow on the time interval $[0,1]$. (a) Left (normalized) singular vector u_2 to singular value $\sigma_2 = 1 - 1.6 \times 10^{-5}$ as obtained from Algorithm 2. (b) Finite-time coherent sets $(\hat{A}, \hat{A}^c)$ at $t = 0$ from u_2. (c) Right (normalized) singular vector v_2. (d) Finite-time coherent sets $(\hat{B}, \hat{B}^c)$ at $t = 1$ from v_2. We obtain $\rho(\hat{A}, \hat{B}) = 1.9976$, which is bounded from above by $1 + \sigma_2$.

of the flow (though the latter property is not so important here). Figure 4 (a) and (c) show the outcome of an application of Algorithm 2 on $[0, 2.5]$ and Figure 4 (b) and (d) the results for the interval $[-2.5, 0]$.

Again $\hat{b} = \hat{b}' = 0$ turn out to be the optimal thresholds for a decomposition into finite-time coherent sets[21]. The decompositions $(\hat{A}, \hat{A}^c)$ at $t = 0$ and $(\hat{B}, \hat{B}^c)$ at $t = 2.5$ are shown in Figure 4 (a) and (c), whereas the respective decompositions $(\hat{A}, \hat{A}^c)$ at $t = -2.5$ and $(\hat{B}, \hat{B}^c)$ at $t = 0$ can be seen in Figure 4 (b) and (d). For both settings, we obtain $\sigma_2 = 0.9999$ and $\rho(\hat{A}, \hat{B}) = 1.9951$, bounded from above by $1 + \sigma_2$.

We compare Figure 4 (a) with Figure 3 (b) to explain the effect of flow duration. Both of these images describe coherent sets at $t = 0$; the only difference is the flow duration. One see that the boundary in Figure 4 (a) is longer than the boundary in Figure 3 (b) and is also closer to the stable manifold of the hyperbolic point on the $x = 0$ axis. The reason for this is that with increasing flow time, advective flux effects increase, relative to diffusive flux. Thus, a short boundary is less important for the longer flow duration; instead a boundary that grows in length at a slower rate over the longer time interval is more important, and so the boundary has grown longer

[21] Varying b and b' around $b = b' = 0$, one obtains partitions that are very close to optimal. Because of the symmetry in the system we concentrate on the partition by sign.

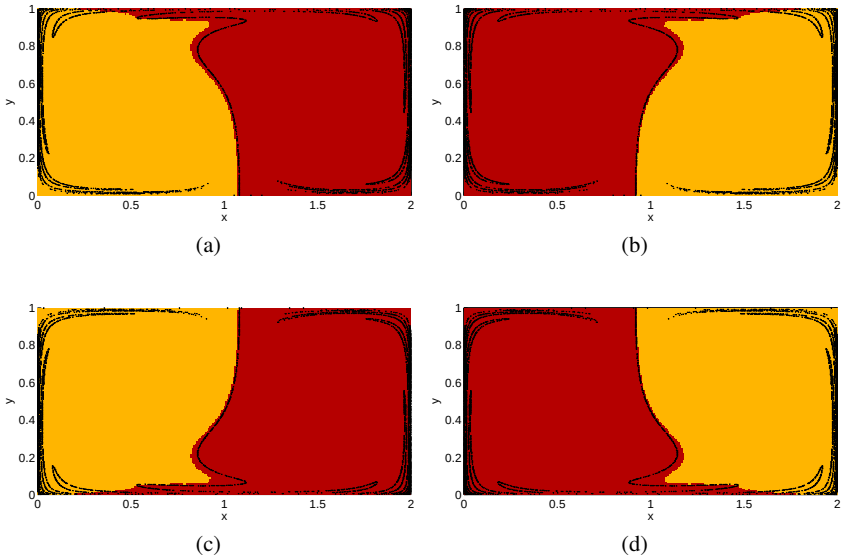

Fig. 4 Finite-time coherent pairs in the double gyre flow on the time interval $[0,2.5]$ and $[-2.5,0]$. (a) Finite-time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0$ from left (normalized) singular vector u_2 as obtained from Algorithm 2. Here the dynamics on $[0,2.5]$ is considered. (b) Finite-time coherent sets $(\hat{A},\hat{A}^c)$ at $t=-2.5$ from left (normalized) singular vector u_2 as obtained from Algorithm 2. Here the dynamics on $[-2.5,0]$ is considered. (c) Finite-time coherent sets $(\hat{B},\hat{B}^c)$ at $t=2.5$ from right (normalized) singular vector v_2 on $[0,2.5]$. (d) Finite-time coherent sets $(\hat{B},\hat{B}^c)$ at $t=0$ from right (normalized) singular vector v_2 from dynamics on $[-2.5,0]$. For both settings we obtain $\rho(\hat{A},\hat{B}) = 1.9951$, which is bounded from above by $1+\sigma_2$, with singular value $\sigma_2 = 0.9999$. The optimal thresholds in Algorithm 2 are $\hat{b}=\hat{b}'=0$. The partitions are overlaid with approximations of the stable and unstable manifolds of hyperbolic periodic orbits.

and moved closer to the stable manifold relative to Figure 3. This ensures that the boundary in Figure 4 (c) is shorter than the push-forward of the boundary in Figure 3 (b) to time $t=2.5$ would be. In each case (flows times 1 and 2.5), the coherent sets are tuned to their particular finite-time duration to resist mixing over that period.

Secondly, we notice that although the finite-time coherent sets in Figure 4 (a) and 4 (d) are both defined at $t=0$, the different time spans under consideration ($[0,2.5]$ and $[-2.5,0]$) produce markedly different results. This can be explained as follows. For the interval $[0,2.5]$, the sets at $t=0$ should have boundaries that are small (to reduce the diffusive effect at $t=0$) and the images of the sets at $t=2.5$ should have the same property (to reduce the diffusive effect at $t=2.5$). The optimal coherent pair at $t=0$ therefore have a boundary that is approximately the stable manifold of the hyperbolic point at the bottom of the rectangle at $t=0$; then under forward iteration, this boundary will not grow very much, and indeed will tend to align with the unstable manifold of the hyperbolic point on top of the rectangle. The deviation from the true stable and unstable manifolds is a function of the diffusion

level and the finite time duration. As the diffusion goes to zero and the time duration goes to infinity, we expect the boundaries to approach the true stable and unstable manifolds. However, for a given diffusion level and finite time duration, the sets shown here have less flux transfer than the true manifolds (which are complicated objects that would create a very long boundary); the coherent sets are optimised to resist diffusion-assisted mixing over the time interval $[0,2.5]$. Turning now to the interval $[-2.5,0]$, we can apply the same argument as above, except that now $t=0$ is the final time in the interval and by the above argument one expects the coherent set boundaries to be roughly aligned along unstable directions, rather than stable directions. This is indeed what we see in Figure 4.

Finite-time coherent sets: both time directions

Finally, we consider the coherent set framework that takes into account both time directions. For this we study again both the dynamics on $[0,2.5]$ and the backwards time dynamics on $[-2.5,0]$ as in the previous paragraph and apply Algorithm 3. We obtain the second left eigenvector u_2 (after normalization) at $t=0$, corresponding to the eigenvalue $\lambda_2 = 0.9998$ as well as corresponding optimal vectors v_2^+ at $t=2.5$ and v_2^- at $t=-2.5$. The optimal thresholds in Algorithm 3 are given by $\hat{b}=\hat{b}'=\hat{b}''=0$, defining finite time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0$, as shown in Figure 5 (a), $(\hat{B},\hat{B}^c)$ at $t=2.5$ (Figure 5 (c)) and $(\hat{C},\hat{C}^c)$ at $t=-2.5$ (Figure 5 (e)). One obtains $\rho(\hat{A},\hat{B},\hat{C}) = 1.9878$, which is bounded from above by $1+(\lambda_2)^{1/2} = 1.9999$. Here $\hat{b}=\hat{b}'=\hat{b}''=0$.

While the boundaries of the coherent sets at $t=0$ obtained when considering only one time direction are roughly aligned along stable or unstable manifolds (see Figure 4 (a) and (d)), the result of the triple construction using both time directions simultaneously is a dynamical compromise, influenced by both stable and unstable directions at $t=0$ (see Figure 5 (a)).

We also remark that the longer flow time (5 time units in Figure 5 compared to 2.5 time units in Figure 4) result in an even stronger correspondence of Figure 5 (c)–(f) with stable/unstable manifolds (a zoom is shown in Figure 5 (g)).

We also consider explicit diffusion with $\varepsilon = 0.02$ and obtain coherent triples $\rho(\hat{A},\hat{B},\hat{C}) = 1.9628$ and $(\lambda_2)^{1/2} = 0.9972$, see Figure 5 (b), (d) and (f). The boundaries between the respective finite-time coherent sets are shorter and smoother than in the non-diffusive case, and the small disconnected pieces (lower left of Figure 5 (c) (dark), upper right of Figure 5 (e) (light)) no longer appear in Figure 5 (d) and (f), respectively, as the extra diffusive flux makes the small disconnected pieces non-optimal.

7.2 Case study 2: Transitory double gyre flow

We consider the transitory dynamical system introduced in [35]

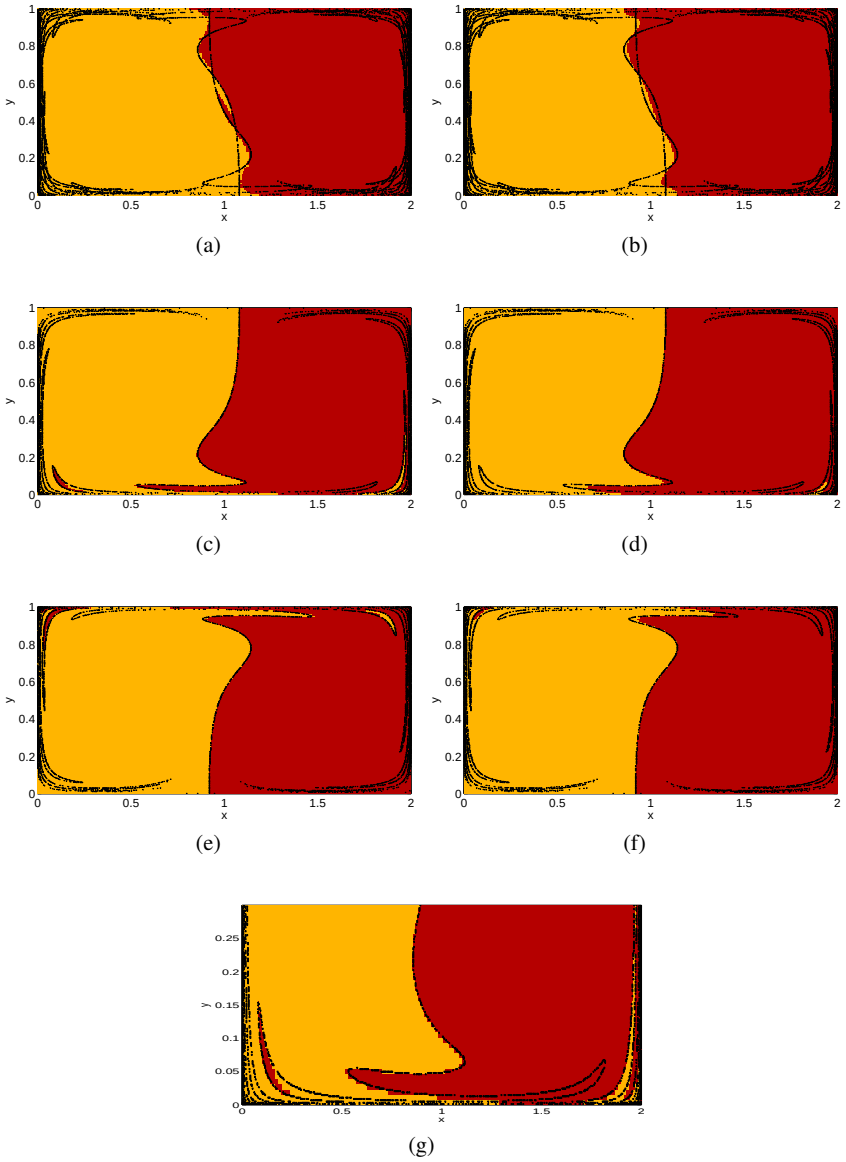

Fig. 5 Finite-time coherent sets using forwards and backwards time dynamics on $[0,2.5]$ and $[-2.5,0]$ based on non-diffusive setting (left) and for the case with diffusion (right). (a) Finite time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0$ obtained from Algorithm 3, no explicit diffusion. (b) Same as (a) but for $\varepsilon = 0.02$. (c) Corresponding finite time coherent sets $(\hat{B},\hat{B}^c)$ at $t=2.5$ ($\varepsilon = 0$); (g) shows a close-up. (d) Same as (c) but for $\varepsilon = 0.02$. (e) Corresponding finite time coherent sets $(\hat{C},\hat{C}^c)$ at $t=-2.5$ ($\varepsilon = 0$). (f) Same as (e) but for $\varepsilon = 0.02$. The coherence of the triples for the case $\varepsilon = 0$ can be estimated as $\rho(\hat{A},\hat{B},\hat{C}) = 1.9878$, which is bounded from above by $1+(\lambda_2)^{1/2} = 1.9999$. For $\varepsilon = 0.02$ one obtains $\rho(\hat{A},\hat{B},\hat{C}) = 1.9628$ and $(\lambda_2)^{1/2} = 0.9972$.

$$\dot{x} = -\frac{\partial}{\partial y}\Psi, \quad \dot{y} = \frac{\partial}{\partial x}\Psi,$$

with stream function

$$\begin{aligned}\Psi(x,y,t) &= (1-s(t))\Psi_P + s(t)\Psi_F \\ \Psi_P(x,y) &= \sin(2\pi x)\sin(\pi y) \\ \Psi_F(x,y) &= \sin(\pi x)\sin(2\pi y)\end{aligned}$$

and transition function

$$s(t) = \begin{cases} 0, & t < 0, \\ t^2(3-2t), & 0 \le t \le 1, \\ 1, & t > 1. \end{cases}$$

A horizontal double gyre pattern (described by the "past" system with stream function Ψ_P) on the unit square is rotated anti-clockwise during times $0 < t < 1$ into a vertical double gyre pattern (described by the "future" system with stream function Ψ_F). For $t \le 0$ and $t \ge 1$ the system is autonomous, thus the interesting finite-time behaviour is restricted to $[0,1]$. Figure 6 illustrates the complex mixing processes on $[0,1]$. Of particular importance are the separatrices for the "past" and "future" autonomous systems, i.e. the vertical line at $x = 0.5$ for $t \le 0$ (Figure 6 (a)) and the horizontal line at $y = 0.5$ for $t \ge 1$ (Figure 6 (f)). Their images and preimages in the transitory time interval $[0,1]$ completely describe the transport mechanism as can be seen in Figure 6. We refer to [35] for a detailed discussion and analysis of this transitory dynamical system.

For our numerical study we partition the domain $M = [0,1] \times [0,1]$ in $n = 16384 = 2^{14}$ square boxes and form matrices $\bar{P}$ by integrating with a Runge Kutta scheme with constant stepsize $h = 0.01$ for different time spans on $[0,1]$, using $K = 400$ uniformly distributed test points per box (inner grid points). In this set-up, again with box diameter 0.0078, we have a numerically induced diffusion of about $\varepsilon \approx 0.0039$, but we do not take into account any explicit diffusion.

First we consider the flow on the entire transition interval $[0,1]$. The results of an application of Algorithm 2 are shown in Figure 7. The (normalized) singular vectors u_2 and v_2 with respect to the singular value $\sigma_2 = 0.9997$ are shown in Figure 7 (a) and (c). The thresholds $\hat{b} = -0.0525$, $\hat{b}' = -0.0530$ turn out to define optimal coherent sets $(\hat{A}, \hat{A}^c)$ at $t = 0$ and $(\hat{B}, \hat{B}^c)$ at $t = 1$, shown in Figure 7 (b) and (d). The coherence can be estimated as $\rho(\hat{A}, \hat{B}) = 1.9885$, which is bounded from above by $1 + \sigma_2 = 1.9997$. The finite-time coherent sets appear to pick up the dominant light structures in Figure 6 (c) and (f), whose boundaries match those of the sets considerably.

However, one may wonder why the left hand "blob" is picked up as the optimal coherent set in Figures 7(b), and not the right hand "blob". An inspection of Figure 6(c) reveals that a small horizontal "cut" across the thin light filament at approximately $(x,y) = (1, 0.5)$ will separate the light image blob in Figure 6(c). This also holds for the light preimage blob in Figure 6(d) (here cutting at approximately

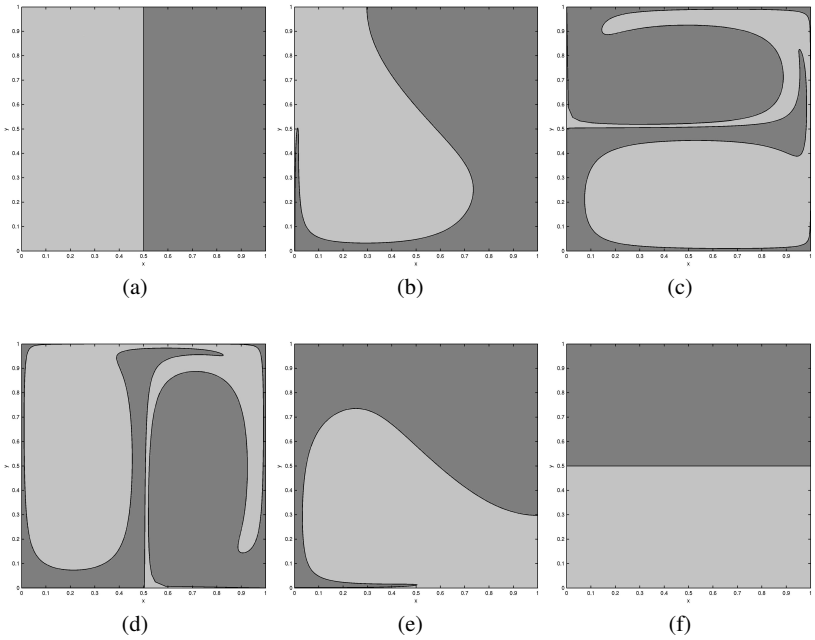

Fig. 6 Complex mixing dynamics of the transitory double gyre system. (a) Two vertical sets, corresponding to invariant sets for the autonomous dynamics for $t \leq 0$, are initialized at time $t = 0$ and evolved forward under the dynamics. The sets are bounded by a heteroclinic orbit connecting two saddle point equilibria at $(0.5, 0)$ and $(0.5, 1)$ in the "past" autonomous system. (b) Image sets of (a) at $t = 0.5$ and (c) at $t = 1$. (d) Preimage at $t = 0$ of the two horizontal sets in (f) that are bounded by the separatrix of the "future" autonomous system. (e) same as (d) but at $t = 0.5$. (e) Two sets corresponding to invariant sets for the autonomous dynamics for $t \geq 1$, are initialized at time $t = 1$ and evolved backwards under the dynamics. The sets are bounded by a heteroclinic orbit connecting two saddle point equilibria at $(0, 0.5)$ and $(1, 0.5)$ in the "future" autonomous system.

$(x, y) = (0.5, 1)$). Such a cut will lead to only a small advective flux across the cut and both the light blob and its image have relatively small boundary, to reduce diffusive flux. On the other hand, to similarly separate the dark image blob in Figure 6(c) would require either a longer horizontal cut around $(x, y) = (0.1, 0.9)$, leading to a larger advective flux, or a small vertical cut near eg. $(x, y) = (0.2, 1)$, leading to small advective flux, but slightly more diffusive flux as the boundary of the dark image blob has been lengthened. Thus, there is some slight asymmetry in the system, which is clearly picked up by the coherent set calculations.

In a second set of numerical experiments we consider transport and mixing on the subintervals $[0, 0.5]$ and $[0.5, 1]$. The results of applying Algorithm 2 are shown in Figure 8. The decompositions $(\hat{A}, \hat{A}^c)$ at $t = 0$ and $(\hat{B}, \hat{B}^c)$ at $t = 0.5$ w.r.t. the dynamics on $[0, 0.5]$ can be seen in Figure 8 (a) and (c), whereas $(\hat{A}, \hat{A}^c)$ at $t = 0.5$ and $(\hat{B}, \hat{B}^c)$ at $t = 1$ w.r.t. the dynamics on $[0.5, 1]$ are shown in Figure 8 (b) and (d).

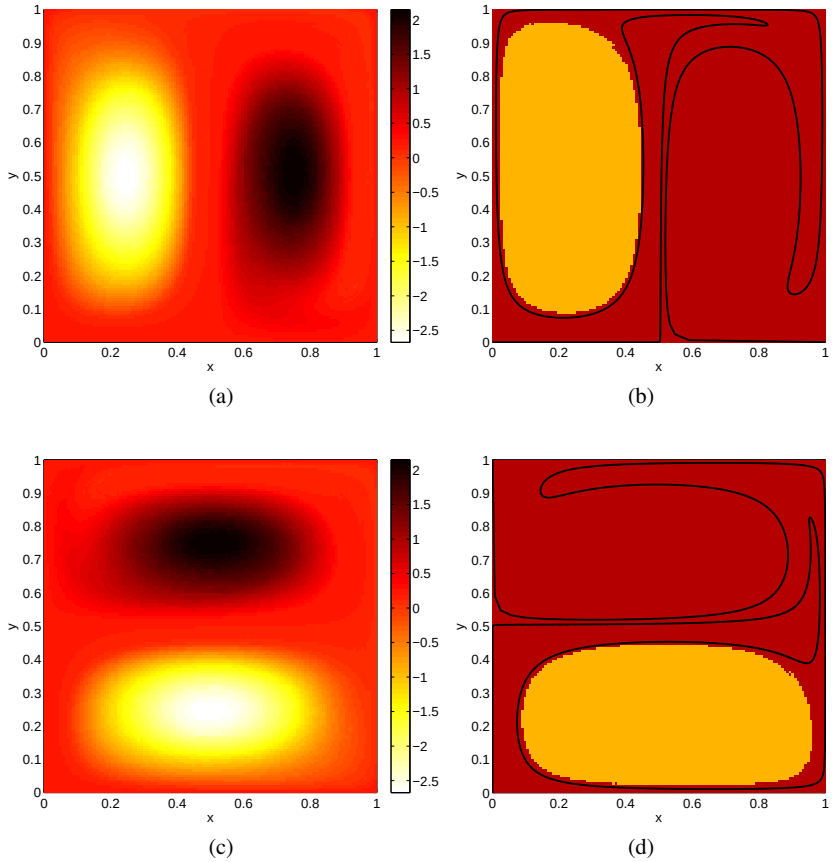

Fig. 7 Singular vectors and finite-time coherent pairs in the transitory double gyre flow on $[0,1]$. (a) Left (normalized) singular vector u_2 to singular value $\sigma_2 \approx 0.9997$ obtained from Algorithm 2. (b) Finite-time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0$ from u_2. (c) Right (normalized) singular vector v_2. (d) Finite-time coherent sets $(\hat{B},\hat{B}^c)$ at $t=1$ from v_2. We obtain $\rho(\hat{A},\hat{B}) = 1.9885$, which is well bounded from above by $1+\sigma_2$. The transport barriers from Figure 6 (c) and (f) are overlaid, delineating considerable parts of the boundaries of the finite-time coherent sets.

On both time intervals we get the same second singular value $\sigma_2 = 0.9998$ and $\hat{b} = 0.1488$, $\hat{b}' = 0.1511$ turn out to be the optimal thresholds for a decomposition into finite-time coherent sets.The coherence values $\rho(\hat{A},\hat{B}) = 1.9919$ for both settings are bounded by the theoretical upper bound of $1+\sigma_2$. While the shapes of the sets are visibly influenced by the respective geometric structures, the boundaries of the sets do not fully match the images of $\{x = 1/2\}$ and $\{y = 1/2\}$. The reason is again the small level of diffusion. The black curves in Figure 8 (a)–(d) are all much

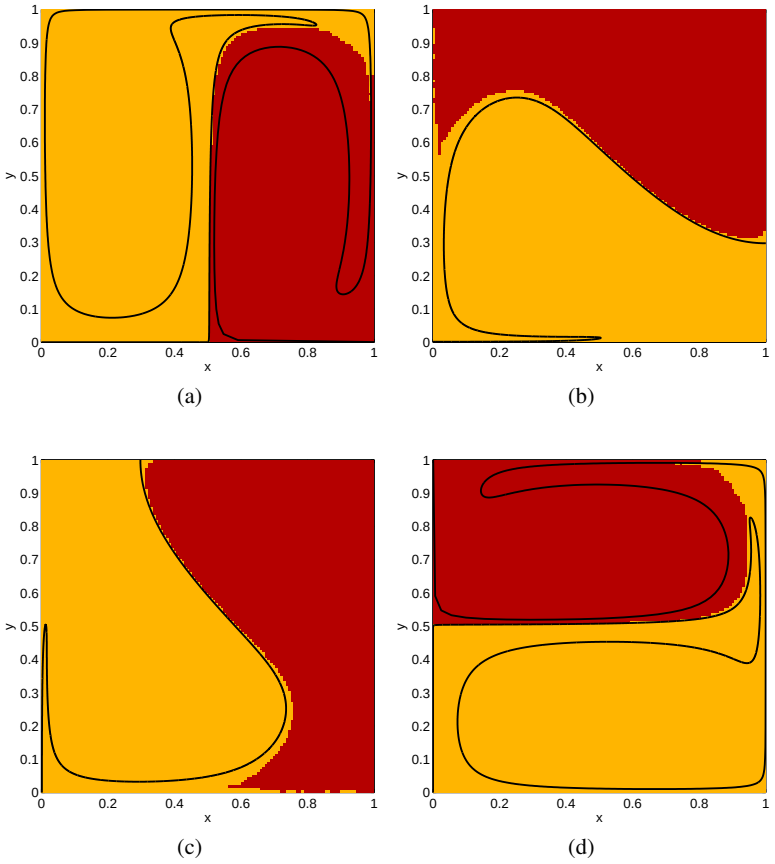

Fig. 8 Finite-time coherent sets in the transitory double gyre flow on $[0,0.5]$ (left) as well as on $[0.5,1]$ (right) using the one-sided construction in Algorithm 2. (a) Finite-time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0$ from dynamics on $[0,0.5]$. (b) Finite-time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0.5$ from dynamics on $[0.5,1]$. (c) Finite-time coherent sets $(\hat{B},\hat{B}^c)$ at $t=0.5$ from dynamics on $[0,0.5]$. (d) Finite-time coherent sets $(\hat{B},\hat{B}^c)$ at $t=1$ from dynamics on $[0.5,1]$. For both settings $\rho(\hat{A},\hat{B}) = 1.9919 < 1+\sigma_2 = 1.9998$. The optimal thresholds in Algorithm 2 are $\hat{b}=0.1488$, $\hat{b}'=0.1511$. The corresponding transport barriers from Figure 6 are overlaid.

longer than the boundaries of the coherent sets. Thus, again, the coherent sets are tuned to a particular flow duration and small diffusion level.

Finally, we consider again the triple construction using the matrices on $[0,0.5]$ and $[0.5,1]$ from the previous paragraph. The results of an application of Algorithm 3 are shown in Figure 9. In particular, Figure 9 (a) shows the optimal decomposition into finite-time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0.5$, (b) the sets $(\hat{B},\hat{B}^c)$ at $t=1$ and (c) the corresponding partition $(\hat{C},\hat{C}^c)$ at $t=0$. With thresholds $\hat{b}=-0.2532$, $\hat{b}'=\hat{b}''=-0.2464$ we obtain $\rho(\hat{A},\hat{B},\hat{C})=1.9883$ bounded by $1+(\lambda_2)^{1/2}$, where $\lambda_2=0.9994$.

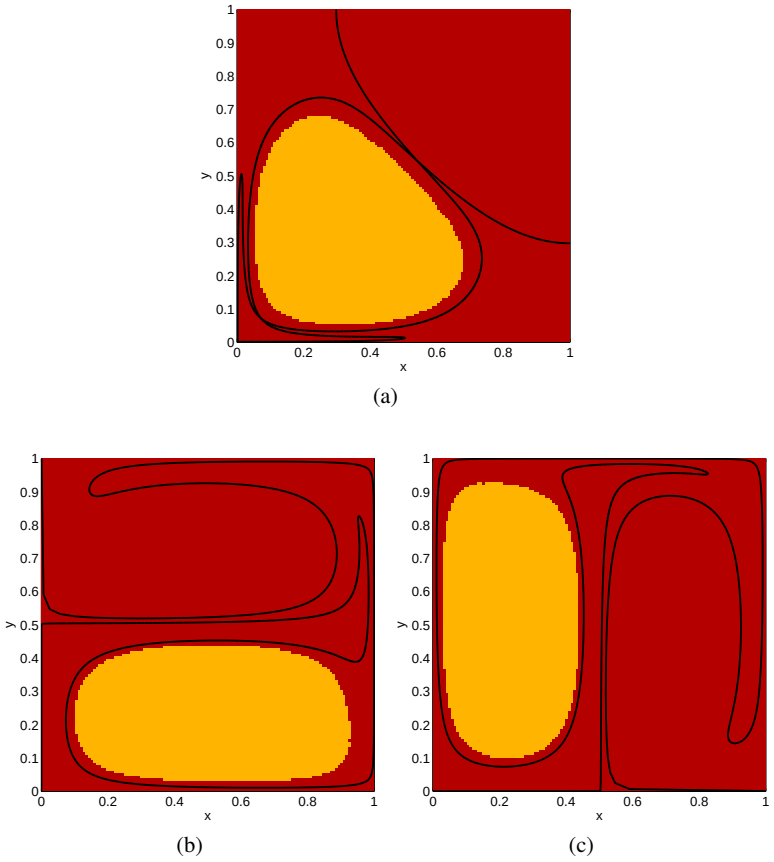

Fig. 9 Finite-time coherent sets for transitory double gyre flow on $[0,0.5]$ and $[0.51]$ using the two-sided construction (Algorithm 3. (a) Optimal finite time coherent sets $(\hat{A},\hat{A}^c)$ at $t=0.5$. (b) $(\hat{B},\hat{B}^c)$ at $t=1$. (c) $(\hat{C},\hat{C}^c)$ at $t=0$. Here $\hat{b}=-0.2532$, $\hat{b}'=\hat{b}''=-0.2464$ and $\rho(\hat{A},\hat{B},\hat{C})=1.9883$.

While the optimal sets do not exactly match the boundaries of the geometric structures in Figure 6, the partitions at $t=0$ and $t=1$ are very similar to the ones in Figure 7. Moreover, the central sets at $t=0.5$ (Figure 9 (a)) are influenced both by stable and unstable directions and thus form, as expected, a dynamical compromise which accounts for both the forward and backward time dynamics. As we have observed before, geometrical structures such as invariant manifolds may not necessarily bound sets of minimal leakage. This is again due to the fact that deterministic time-asymptotic objects do not account for the finite-time and diffusive effects that are central to our transfer operator construction.

8 Summary

In this chapter we proposed a unified setting for finite-time almost-invariant and coherent set constructions. The constructions were based around a building block operator that combined advective dynamics via a Perron-Frobenius operator with small amplitude diffusive dynamics, developed in [13]. This building block operator was then manipulated to set up suitable optimisation problems for the autonomous and nonautonomous settings. The unified setting clarified the similarities and differences of the dynamical problems being solved in the autonomous and non-autonomous cases. These optimisation problems made use of the fact that the underlying transfer operator was compact and self-adjoint, leading to a simple analytic solution given by the eigenfunction corresponding to the second largest eigenvalue.

Via two detailed case studies we investigated the dependence of almost-invariant and coherent sets on three aspects of the dynamics: the level of diffusion, the flow duration, and the time-direction of the dynamics. We also compared the boundaries of the coherent sets with the time-dependent stable and unstable manifolds of organising hyperbolic points in flows.

As proved formally in [13], for fixed flow times, we showed that increased levels of diffusion produce more regular eigenfunctions, and coherent sets with shorter boundaries. Intuitively, this is because diffusive flux is proportional to the boundary lengths of coherent sets, and so the boundary would like to be shortened to minimise the effects of diffusive flux. On the other hand, shortening the boundary usually means that the boundary is less equivariant (or in the autonomous case, invariant) in an advective sense, increasing the advective flux. The result is that an optimal compromise is reached, tuned to the specific diffusion level and flow duration, with a somewhat shorter boundary. These remarks apply equally to the autonomous and non-autonomous settings. We also demonstrated an instance where the optimal almost-invariant sets appear to undergo a bifurcations as the diffusion level is increased.

Increasing flow duration, but fixing the diffusion level, has the opposite effect. Now, advective flux becomes a greater component of the overall flux, and the boundary tries to move in a way that more aligned with "stable" directions. This is so that at the final flow time, the boundary has not grown very much (which would lead to high diffusive flux at the final time). Again, an optimal balance is reached; the boundary of the coherent set moves close to the stable direction at the initial time, likely growing somewhat in length, the result tuned to the the specific flow time and diffusion level.

In hyperbolic settings, we demonstrated numerically that the boundary of the coherent set is approximately aligned with stable directions at the initial time and *unstable* directions at the *final* time. This can be intuitively explained by an argument identical to those above; the coherent set at the final time should have relatively short boundary, and under backward advection, the boundary should not grow very long, otherwise there will be high diffusive flux at the initial time. In order for this to occur, the coherent set is approximately aligned along unstable directions at the final time so that under backward advection the boundary does not grow very long;

the exact positioning of the coherent sets are tuned to the particular flow time and diffusion level to minimise flux out of the sets.

Finally, in order to produce a sequence of coherent sets over a sliding window of fixed finite-time duration, we proposed a new construction where the focus is on coherent sets in the middle of the finite-time window; these sets remain coherent in both forward and backward time and can be used to create a natural time-dependent sequence of coherent sets over several translated finite-time windows.

Future work will include further investigation of bifurcation phenomena and computational improvements to the bi-directional coherent set calculations.

Acknowledgements GF and KPG thank the Banff International Research Station for hospitality during the workshop on "Open Dynamical Systems: Ergodic Theory, Probabilistic Methods and Applications", which provided the catalyst for this work. GF's research was partially supported by an ARC Disovery Project (DP110100068) and an ARC Future Fellowship.

Appendix

8.1 Proof of Theorem 3

Our building block operator $\mathbf{L}$ may be written as $\mathbf{L}f(y) = \int_X k(x,y)f(x)\ d\mu(x)$, where $k \in L^2(X, Y_\varepsilon)$, satisfies $k \geq 0$. Thus $\mathbf{Q} = (\mathbf{L}+\mathbf{L}^*)/2$ is a self-adjoint operator $\mathbf{Q} : L^2(X,\mu) \to L^2(X,\mu)$ defined by $\mathbf{Q}f(y) = \int \kappa(x,y)f(x)\ d\mu(x)$, with $\mathbf{Q1} = \mathbf{1}$ (see [13] for the specific forms of k and κ). From self-adjointness it follows that $\kappa(x,y) = \kappa(y,x)$ and thus

$$\int \kappa(x,y)\ d\mu(y) = 1 = \int \kappa(x,y)\ d\mu(x). \tag{37}$$

Lower bound

The proof of Lemma 1 draws heavily on the proof of Theorem 2.1 [29] (which is effectively a continuous time version of Lemma 1) and includes techniques from Theorem 4.3, Chapter 6 [5].

Lemma 1. *Let λ_2 denote the second largest eigenvalue of $\mathbf{Q}$. Then $1-\lambda_2 \geq c^2/8$ where $c = \inf_A(\int \mathbf{Q}\mathbf{1}_A \cdot \mathbf{1}_{A^c}\ d\mu)/(\mu(A)\mu(A^c))$.*

Proof. We know that $\sup\{\langle \mathbf{Q}f, f\rangle_\mu / \langle f, f\rangle_\mu : \langle f, \mathbf{1}\rangle_\mu = 0\} = \lambda_2$, where λ_2 is the second largest eigenvalue of $\mathbf{Q}$. Thus $\inf\{\langle (I-\mathbf{Q})f, f\rangle_\mu / \langle f, f\rangle_\mu : \langle f, \mathbf{1}\rangle_\mu = 0\} = 1-\lambda_2$.

Now,

$$
\begin{aligned}
&\langle (I-\mathbf{Q})f, f\rangle_\mu \\
&= \int f^2 \, d\mu - \int \kappa(x,y) f(y) \, d\mu(y) \cdot f(x) \, d\mu(x) \\
&= \int f^2 \, d\mu - \int \kappa(x,y) f(y) \, d\mu(y) \cdot f(x) \, d\mu(x) \\
&= \int f(x) \left(\int \kappa(x,y)(f(x)-f(y)) \, d\mu(y) \right) d\mu(x) \quad \text{using (37)} \\
&= \int f(y) \left(\int \kappa(x,y)(f(y)-f(x)) \, d\mu(y) \right) d\mu(x) \quad \text{interchanging } x \text{ and } y \\
&= (1/2) \int \left(\int \kappa(x,y)(f(x)-f(y))^2 \, d\mu(y) \right) d\mu(x) \quad \text{combining previous 2 lines.}
\end{aligned}
$$

Now

$$
\begin{aligned}
&\int \kappa(x,y)(f(x)+f(y))^2 \, d\mu(x) d\mu(y) \\
&\leq 2 \int \kappa(x,y)(f(x)^2+f(y)^2) \, d\mu(x) d\mu(y) \quad \text{using } (a+b)^2 \leq 2(a^2+b^2) \\
&= \int 4\kappa(x,y) f(x)^2 \, d\mu(x) d\mu(y) \quad \text{by symmetry of } \kappa \\
&= \int 4 f(x)^2 \, d\mu(x) \quad \text{by (37).}
\end{aligned}
$$

So

$$
\begin{aligned}
&(1/2) \int \left(\int \kappa(x,y)(f(x)-f(y))^2 \, d\mu(y) \right) d\mu(x) \\
&\geq (1/2) \int \left(\int \kappa(x,y)(f(x)-f(y))^2 \, d\mu(y) \right) d\mu(x) \cdot \frac{\int \kappa(x,y)(f(x)+f(y))^2 \, d\mu(x) d\mu(y)}{\int 4 f(x)^2 \, d\mu(x)} \\
&\geq (1/8) \frac{\left(\int \kappa(x,y) |(f(x)^2 - f(y)^2)| \, d\mu(x) d\mu(y) \right)^2}{\int f^2 \, d\mu} \quad \text{by Hölder.}
\end{aligned} \tag{38}
$$

Now

$$
\begin{aligned}
&\int \kappa(x,y)|(f(x)^2-f(y)^2)|\,d\mu(x)d\mu(y)\\
&= 2\int \kappa(x,y)\mathbf{1}_{\{f(x)^2-f(y)^2>0\}}(f(x)^2-f(y)^2)\,d\mu(x)d\mu(y) \quad \text{by symmetry of } \kappa\\
&= 2\int_0^\infty d\alpha \int \kappa(x,y)\mathbf{1}_{\{f(x)^2>\alpha\geq f(y)^2\}}\,d\mu(x)d\mu(y)\\
&= 2\int_0^\infty d\alpha \int \kappa(x,y)\mathbf{1}_{A_\alpha}(x)\cdot\mathbf{1}_{A_\alpha^c}(y)\,d\mu(x)d\mu(y) \quad \text{where } A_\alpha=\{f^2>\alpha\}\\
&= 2\int_0^\infty d\alpha \int \mathbf{Q}\mathbf{1}_{A_\alpha}(y)\cdot\mathbf{1}_{A_\alpha^c}(y)\,d\mu(y)\\
&\geq 2c\int_0^\infty d\alpha\,\mu(A_\alpha)\mu(A_\alpha^c), \quad \text{where } c=\inf_A(\textstyle\int \mathbf{Q}\mathbf{1}_A\cdot\mathbf{1}_{A^c}\,d\mu)/(\mu(A)\mu(A^c))\\
&= 2c\int_0^\infty d\alpha \int \mathbf{1}_{\{f(x)^2>\alpha\geq f(y)^2\}}d\mu(x)d\mu(y)\\
&= 2c\int_0^\infty d\alpha \int \mathbf{1}_{\{f(x)^2-f(y)^2>0\}}(f(x)^2-f(y)^2)d\mu(x)d\mu(y)\\
&= c\int |f(x)^2-f(y)^2|\,d\mu(x)d\mu(y).
\end{aligned}
$$

By (38) we have

$$
\begin{aligned}
\langle (I-\mathbf{Q})f,f\rangle_\mu &\geq (1/8)\frac{c^2\left(\int |f(x)^2-f(y)^2|\,d\mu(x)d\mu(y)\right)^2}{(\int f^2\,d\mu)}\\
&\geq (1/8)c^2\kappa'\int f^2\,d\mu\\
&\geq (1/8)c^2\int f^2\,d\mu \quad \text{by [29], Prop. 2.2 (see for a definition of } \kappa').
\end{aligned}
$$

Thus $1-\lambda_2\geq c^2/8$.

Define

$$
\rho(A)=\frac{\int \mathbf{Q}\mathbf{1}_A\cdot\mathbf{1}_A\,d\mu}{\mu(A)}+\frac{\int \mathbf{Q}\mathbf{1}_{A^c}\cdot\mathbf{1}_{A^c}\,d\mu}{\mu(A^c)} \quad \text{and} \quad c(A)=\frac{\int \mathbf{Q}\mathbf{1}_A\cdot\mathbf{1}_{A^c}\,d\mu}{\mu(A)\mu(A^c)}.
$$

Lemma 2. $\rho(A)=2-c(A)$.

Proof.

$$
\begin{aligned}
&\frac{\int \mathbf{Q}\mathbf{1}_A\cdot\mathbf{1}_A\,d\mu}{\mu(A)}+\frac{\int \mathbf{Q}\mathbf{1}_{A^c}\cdot\mathbf{1}_{A^c}\,d\mu}{\mu(A^c)}\\
&=\frac{\mu(A^c)(\int \mathbf{Q}\mathbf{1}_A\cdot\mathbf{1}\,d\mu-\int \mathbf{Q}\mathbf{1}_A\cdot\mathbf{1}_{A^c}\,d\mu)}{\mu(A)\mu(A^c)}+\frac{\mu(A)(\int \mathbf{Q}\mathbf{1}_{A^c}\cdot\mathbf{1}\,d\mu-\int \mathbf{Q}\mathbf{1}_{A^c}\cdot\mathbf{1}_A\,d\mu)}{\mu(A)\mu(A^c)}\\
&=\frac{\mu(A^c)\mu(A)-\mu(A^c)\int \mathbf{Q}\mathbf{1}_A\cdot\mathbf{1}_{A^c}\,d\mu)}{\mu(A)\mu(A^c)}+\frac{\mu(A)\mu(A^c)-\mu(A)\int \mathbf{Q}\mathbf{1}_{A^c}\cdot\mathbf{1}_A\,d\mu)}{\mu(A)\mu(A^c)}\\
&=2-\mu(A^c)c(A)-\mu(A)c(A^c)=2-c(A) \quad \text{as } c(A)=c(A^c).
\end{aligned}
$$

Corollary 2. $\rho := \sup_A \rho(A) \geq 2 - \sqrt{8(1-\lambda_2)}$.

Proof. $\sup_A \rho(A) = 2 - \inf_A c(A) = 2 - c = 2 - \sqrt{8(1-\lambda_2)}$.

Upper bound

Lemma 3. $\rho \leq \lambda_2 + 1$.

Proof. We know that $\langle \mathbf{Q}f, f\rangle_\mu \leq \lambda_2$ for all $\langle f, \mathbf{1}\rangle_\mu = 0$. Consider the test function $f = \sqrt{\mu(A^c)/\mu(A)}\mathbf{1}_A - \sqrt{\mu(A)/\mu(A^c)}\mathbf{1}_{A^c}$. We have

$$
\begin{aligned}
&\langle \mathbf{Q}f, f\rangle_\mu \\
&= \langle \sqrt{\mu(A^c)/\mu(A)}\mathbf{Q}\mathbf{1}_A, \sqrt{\mu(A^c)/\mu(A)}\mathbf{1}_A\rangle_\mu - \langle \sqrt{\mu(A^c)/\mu(A)}\mathbf{Q}\mathbf{1}_A, \sqrt{\mu(A^c)/\mu(A)}\mathbf{1}_{A^c}\rangle_\mu \\
&\qquad - \langle \sqrt{\mu(A^c)/\mu(A)}\mathbf{Q}\mathbf{1}_{A^c}, \sqrt{\mu(A^c)/\mu(A)}\mathbf{1}_A\rangle_\mu + \langle \sqrt{\mu(A^c)/\mu(A)}\mathbf{Q}\mathbf{1}_A^c, \sqrt{\mu(A^c)/\mu(A)}\mathbf{1}_{A^c}\rangle_\mu \\
&= \frac{\mu(A^c)}{\mu(A)}\langle \mathbf{Q}\mathbf{1}_A, \mathbf{1}_A\rangle_\mu - 2\langle \mathbf{Q}\mathbf{1}_A, \mathbf{1}_{A^c}\rangle_\mu + \frac{\mu(A)}{\mu(A^c)}\langle \mathbf{Q}\mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu \\
&= \frac{\mu(A^c)}{\mu(A)}\langle \mathbf{Q}\mathbf{1}_A, \mathbf{1}_A\rangle_\mu + \frac{\mu(A)}{\mu(A^c)}\langle \mathbf{Q}\mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu \\
&\qquad - (\langle \mathbf{Q}\mathbf{1}_A, \mathbf{1}\rangle_\mu - \langle \mathbf{Q}\mathbf{1}_A, \mathbf{1}_A\rangle_\mu + \langle \mathbf{Q}\mathbf{1}_{A^c}, \mathbf{1}\rangle_\mu - \langle \mathbf{Q}\mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu) \\
&= \frac{\langle \mathbf{Q}\mathbf{1}_A, \mathbf{1}_A\rangle_\mu}{\mu(A)} + \frac{\langle \mathbf{Q}\mathbf{1}_{A^c}, \mathbf{1}_{A^c}\rangle_\mu}{\mu(A^c)} - \mu(A) - \mu(A^c) \\
&\leq \lambda_2.
\end{aligned}
$$

As A is arbitrary, the result follows.

References

1. H. Aref. The development of chaotic advection. *Physics of Fluids*, 14(4):1315–1325, 2002.
2. L. Billings and I. B. Schwartz. Identifying almost invariant sets in stochastic dynamical systems. *Chaos*, 18:023122, 2008.
3. M. S. Birman and M. Z. Solomjak. *Spectral theory of self-adjoint operators in Hilbert space*. D. Reidel Publishing Co., Inc., 1986.
4. E. M. Bollt, L. Billings, and I. B. Schwartz. A manifold independent approach to understanding transport in stochastic dynamical systems. *Physica D*, 173:153–177, 2002.
5. P. Brémaud. *Markov chains. Gibbs fields, Monte Carlo simulation, and queues*, volume 31 of *Texts in Applied Mathematics*. Springer, New York, 1999.
6. J.B. Conway. *A course in functional analysis*, volume 96 of *Graduate texts in mathematics*. Springer, 2nd edition, 1990.
7. M. Dellnitz, G. Froyland, and O. Junge. *The Algorithms behind GAIO – Set Oriented Numerical Methods for Dynamical Systems*, pages 145–174. Springer, 2001.
8. M. Dellnitz and O. Junge. Almost invariant sets in Chua's circuit. *Int. J. Bif. and Chaos*, 7(11):2475–2485, 1997.
9. M. Dellnitz and O. Junge. On the approximation of complicated dynamical behaviour. *SIAM Journal for Numerical Analysis*, 36(2):491–515, 1999.

10. P. Deuflhard, W. Huisinga, A. Fischer, and C. Schütte. Identification of almost invariant aggregates in nearly uncoupled Markov chains. *Linear Algebra and its Applications*, 315:39–59, 2000.
11. G. Froyland. Statistically optimal almost-invariant sets. *Physica D*, 200:205–219, 2005.
12. G. Froyland. Unwrapping eigenfunctions to discover the geometry of almost-invariant sets in hyperbolic maps. *Physica D*, 237:840–853, 2008.
13. G. Froyland. An analytic framework for identifying finite-time coherent sets in time-dependent dynamical systems. preprint, http://arxiv.org/pdf/1210.7418, 2012.
14. G. Froyland and M. Dellnitz. Detecting and locating near-optimal almost-invariant sets and cycles. *SIAM J. Sci. Comput.*, 24(6):1839–1863, 2003.
15. G. Froyland, C. Horenkamp, V. Rossi, N. Santitissadeekorn, and A. Sen Gupta. Three-dimensional characterization and tracking of an Agulhas ring. *Ocean Modelling*, 52–53:69 –75, 2012.
16. G. Froyland, S. Lloyd, and A. Quas. Coherent structures and isolated spectrum for Perron-Frobenius cocycles. *Ergodic Theory and Dynamical Systems*, 30:729–756, 2010.
17. G. Froyland, S. Lloyd, and N. Santitissadeekorn. Coherent sets for nonautonomous dynamical systems. *Physica D*, 239:1527–1541, 2010.
18. G. Froyland and K. Padberg. Almost-invariant sets and invariant manifolds – connecting probabilistic and geometric descriptions of coherent structures in flows. *Physica D*, 238:1507–1523, 2009.
19. G. Froyland, N. Santitissadeekorn, and A. Monahan. Transport in time-dependent dynamical systems: Finite-time coherent sets. *Chaos*, 20:043116, 2010.
20. R. Guder, M. Dellnitz, and E. Kreuzer. An adaptive method for the approximation of the generalized cell mapping. *Chaos, Solitons & Fractals*, 8(4):525 – 534, 1997.
21. G. Haller. Finding finite-time invariant manifolds in two-dimensional velocity fields. *Chaos*, 10:99–108, 2000.
22. G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D*, 149:248–277, 2001.
23. G. Haller. A variational theory of hyperbolic Lagrangian Coherent Structures. *Physica D*, 240:574–598, 2011.
24. W. Huisinga and B. Schmidt. Metastability and dominant eigenvalues of transfer operators. In B. Leimkuhler, C. Chipot, R. Elber, A. Laaksonen, A. Mark, T. Schlick, C. Schütte, R. Skeel, T. J. Barth, M. Griebel, D. E. Keyes, R. M. Nieminen, D. Roose, and T. Schlick, editors, *New Algorithms for Macromolecular Simulation*, volume 49 of *Lecture Notes in Computational Science and Engineering*, pages 167–182. Springer Berlin Heidelberg, 2006.
25. M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989.
26. O. Junge. An adaptive subdivision technique for the approximation of attractors and invariant measures: proof of convergence. *Dynamical Systems*, 16(3):213–222, 2001.
27. O. Junge, J. E. Marsden, and I. Mezic. Uncertainty in the dynamics of conservative maps. *Proc. 43rd IEEE Conf. Dec. Control*, pages 2225–2230, 2004.
28. A. Lasota and M. C. Mackey. *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*, volume 97 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2nd edition, 1994.
29. G. F. Lawler and A. D. Sokal. Bounds on the L^2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality. *Trans. Amer. Math. Soc.*, 309:557–580, 1988.
30. T.-Y. Li. Finite approximation for the Frobenius-Perron operator. a solution to Ulam's conjecture. *Journal of Approximation Theory*, 17:177–186, 1976.
31. W. Liu and G. Haller. Strange eigenmodes and decay of variance in the mixing of diffusive tracers. *Physica D*, 188:1–39, 2004.
32. J. D. Meiss. Symplectic maps, variational principles, and transport. *Rev. Mod. Phys.*, 64(3):795–848, 1992.
33. I. Mezic. *On the geometrical and statistical properties of dynamical systems: theory and applications*. PhD thesis, California Institute of Technology, 1994.

34. I. Mezić and S. Wiggins. A method for visualization of invariant sets of dynamical systems based on the ergodic partition. *Chaos*, 9(1):213–218, 1999.
35. B. Mosovsky and J.D. Meiss. Transport in transitory dynamical systems. *SIAM J. Dyn. Syst.*, 10:35–65, 2011.
36. R. Murray. Optimal partition choice for invariant measure approximation for one-dimensional maps. *Nonlinearity*, 17(5):1623, 2004.
37. A. Pikovsky and O. Popovych. Persistent patterns in deterministic mixing flows. *Europhys. Lett.*, 61(5):625–631, 2003.
38. O.V. Popovych, A. Pikovsky, and B. Eckhardt. Abnormal mixing of passive scalars in chaotic flows. *Physical Review E*, 75:036308, 2007.
39. V. Rom-Kedar, A. Leonard, and S. Wiggins. An analytical study of transport, mixing and chaos in an unsteady vortical flow. *Journal of Fluid Mechanics*, 214:347–394, 1990.
40. V. Rom-Kedar and S. Wiggins. Transport in two-dimensional maps. *Archive for Rational Mechanics and Analysis*, 109:239–298, 1990.
41. H.L. Royden. *Real analysis*. Pearson, third edition, 1988.
42. C. Schütte, W. Huisinga, and P. Deuflhard. Transfer operator approach to conformational dynamics in biomolecular systems. In B. Fiedler, editor, *Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems*, pages 191–223. Springer, Berlin, 2001.
43. S. C. Shadden, F. Lekien, and J. E. Marsden. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D*, 212:271–304, 2005.
44. M.K. Singh, M.F.M. Speetjens, and P.D. Anderson. Eigenmode analysis of scalar transport in distributive mixing. *Physics of Fluids*, 21:093601–093601, 2009.
45. S. Ulam. *Problems in Modern Mathematics*. Interscience, 1964.
46. S. Wiggins. *Chaotic Transport in Dynamical Systems*. Springer-Verlag, New York, NY, 1992.
47. S. Wiggins. The dynamical systems approach to Lagrangian transport in oceanic flows. *Annu. Rev. Fluid Mech.*, 37:295–328, 2005.