A Tutorial on the Dynamic Laplacian

Gary Froyland

Abstract Spectral techniques are popular and robust approaches to data analysis. A prominent example is the use of eigenvectors of a Laplacian, constructed from data affinities, to identify natural data groupings or clusters, or to produce a simplified representation of data lying on a manifold. This tutorial concerns the dynamic Laplacian, which is a natural generalisation of the Laplacian to handle data that has a time component and lies on a time-evolving manifold. In this dynamic setting, clusters correspond to long-lived "coherent" collections. We begin with a gentle recap of spectral geometry before describing the dynamic generalisations. We also discuss computational methods and the automatic separation of many distinct features through the SEBA algorithm. The purpose of this tutorial is to bring together many results from the dynamic Laplacian literature into a single short document, written in an accessible style.

Keywords Dynamic Laplacian · Dynamic cheeger constant · Finite-time coherent set · Inflated dynamic Laplacian

1 Introduction

Spectral geometry is concerned with relationships between the geometry of manifolds on the one hand, and spectral information from differential operators on manifolds on the other, most commonly the Laplace operator. While so-called direct problems seek spectral information from manifold geometry, inverse problems attempt to learn manifold geometry from the operator spectrum. Inverse problems are of great interest in practical areas such as spectral approaches to dimension reduction of high-dimensional manifolds and geometric clustering of data.

This brief tutorial focusses on the inverse problem, and its generalisation to answer questions concerning long-lived structurs in dynamical systems. In this dynamic set-

G. Froyland (✉)
School of Mathematics and Statistics, UNSW Sydney, Sydney, NSW 2052, Australia
e-mail: g.froyland@unsw.edu.au

© The Author(s) 2026

S. Oishi et al. (eds.), *Recent Developments in Industrial and Applied Mathematics*,
ICIAM2023 Springer Series 1, https://doi.org/10.1007/978-981-95-1446-5_5

ting, the manifold is subjected to nonlinear transformations representing the evolution of the manifold over time. The inverse problem in this setting concerns learning the average geometry of the evolved manifold. Corresponding to this dynamic inverse problem is a natural dynamic generalisation of the Laplace operator to a *dynamic Laplace operator*. Using both the spectrum and eigenfunctions of the dynamic Laplacian we create a dynamic spectral geometry, and illustrate its use in identifying *finite-time coherent sets* in dynamical systems. These sets are the natural generalisations of Cheeger sets to the dynamic setting.

2 Recap on Properties of the Laplace Operator

We begin by recapping some basic spectral geometry for manifolds in the absence of dynamics; see for example Lablée (2015). Consider a full-dimensional manifold[1] $M \subset \mathbb{R}^d$. Expressed in local coordinates, the Laplace operator Δ acts on functions $f : M \to \mathbb{R}$ by $\Delta f = \sum_{i=1}^{d} \frac{\partial^2}{\partial x_i^2} f$. The eigenproblem is

$$\Delta f = \lambda f, \tag{1}$$

with $\langle \nabla f(x), n(x) \rangle = 0$ for all $x \in \partial M$, where $n(x)$ is a vector normal to the boundary of M at $x \in \partial M$, and $\langle \cdot, \cdot \rangle$ is the usual Euclidean inner product. With these Neumann boundary conditions, the eigenvalues consist of a countable sequence $0 = \lambda_1 \geq \lambda_2 \geq \cdots$, of negative numbers, diverging to $-\infty$. The leading eigenvalue $\lambda_1 = 0$ and the corresponding eigenfunction is constant on M. The eigenfunctions $f_1, f_2, \ldots$ form a complete orthonormal set for $L^2(M)$.

2.1 *Optimality Properties of Laplace Eigenvalues and Eigenfunctions*

There are variational expressions for the Laplace eigenvalues, for example, for each $n \geq 1$,

$$\sum_{k=1}^{n} \lambda_k = \min \left\{ \sum_{k=1}^{n} \int_M |\nabla f_k(x)|^2 \, dx : f_1, \ldots, f_n \text{ are orthonormal} \right\}. \tag{2}$$

Moreover, for each $n \geq 1$, the minimizing collection $\{f_1, \ldots, f_n\}$ is precisely the n leading Laplace eigenfunctions. Thus, denoting $L_n = \text{span}\{f_1, \ldots, f_n\}$, we see that

[1] All of the properties discussed here hold on general Riemannian manifolds, but for simplicity of presentation we will consider flat manifolds M endowed with the Euclidean metric.

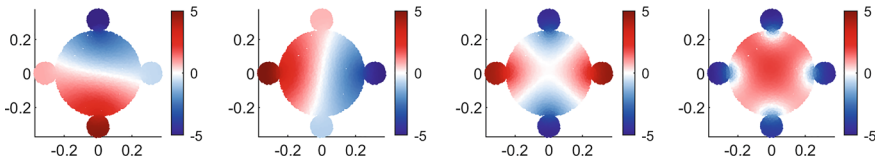

Fig. 1 Laplace eigenfunctions $f_2, \ldots, f_5$ of a two-dimensional manifold. The corresponding eigenvalues are $\lambda_2 \approx -33.42$, $\lambda_3 \approx -33.80$, $\lambda_4 \approx -49.99$, $\lambda_5 \approx -82.42$

the Laplace eigenfunctions are the most regular (with respect to the L^2 norms of gradients) orthonormal basis for L_n. Equivalently to (2) we have for each $n \geq 1$

$$\lambda_n = \min \left\{ \int_M |\nabla f_n(x)|^2 \, dx : \|f_n\|_2 = 1, \, f_n \perp L_{n-1} \right\}, \tag{3}$$

where L_0 is the trivial subspace. Figure 1 shows the first four nontrivial eigenfunctions of the Laplace operator on a two-dimensional manifold consisting of a large disk with four smaller disks attached near the boundary of the large disk. The decreasing regularity–left to right in Fig. 1–as indicated in (2) and (3) is apparent.

2.2 The Heat Equation and Geometry

The reader may also have noticed relationships between the eigenfunctions in Fig. 1 and the geometry of the manifold. Indeed there are well-known links between the Laplace eigendata of a manifold and the geometric shape of the manifold. Firstly, consider the heat equation on a manifold M, where the scalar value $u(t, x)$ denotes the temperature at position x and time t:

$$\frac{\partial u}{\partial t} = \Delta u. \tag{4}$$

We impose the same Neumann boundary conditions as above; this means that there is no heat lost at the boundary due to a zero derivative in the temperature field normal to the boundary. If we initialise a heat flow with $u(0, x) = f_k(x)$ then we obtain $u(t, x) = e^{\lambda_k t} f_k(x), k = 1, 2, \ldots$. The minimality properties (2) of the eigenvalues λ_k corresponds to a minimal decay of an orthonormal set of initial conditions $f_1, \ldots, f_n$ toward the heat equilibrium of $f_1 \equiv 1/\sqrt{\text{vol}(M)}$. In particular, in the two left panels of Fig. 1 we see that positive "heat" is concentrated to one side of the manifold and negative heat is concentrated on the other side, with the greatest concentration in the opposing small disks. Such an initialisation will lead to a lower heat relaxation toward equilibrium. Similarly, in the right two panels of Fig. 1 the heat is initialised predominantly in the small disks around the periphery of M. Because of the narrow

necks joining the small disks to the main disk, the heat flow through these necks will contribute to slow relaxation to equilibrium.

2.3 Cheeger Problems and Laplace Eigenfunctions

Similar notions can be constructed in a purely geometric fashion through *Cheeger problems* Cheeger (1970) and Buser (1982). Let $\mathring{M}$ denote the interior of M and $|A|$ denote the full-dimensional Lebesgue volume of A. For a set $A \subset M$ with smooth boundary, we call

$$H(A) := \frac{|\partial A \cap \mathring{M}|}{|A|} \tag{5}$$

the *Neumann Cheeger ratio* of A, which quantifies the boundary size of A relative to its interior size. We call a collection $A_1, \ldots, A_n$ of n disjoint subsets of M an *n-packing* and define $H(\{A_1, \ldots, A_n\}) = \max_{1 \le k \le n} H(A_k)$. Cheeger (or isoperimetric) problems are concerned with decomposing manifolds into subsets with small boundaries relative to their interiors. To this end, we define the n^{th} *Neumann Cheeger constant* of M as

$$h_n := \inf\{H(\{A_1, \ldots, A_n\}) : A_1, \ldots, A_n \subset M \text{ are disjoint}\}; \tag{6}$$

see Froyland and Rock (2025). It is not hard to show that h_n is increasing in n. Cheeger inequalities for n-packings of boundaryless manifolds have been studied in e.g. Miclo (2015), with analogous definitions.

In general, it is difficult to find the minimising collection $A_1, \ldots, A_n$, however the Laplace eigenfunctions can provide good candidate collections. A *nodal domain* of f_k is a maximal connected component of M where the sign of f_k is constant. We have the following result

Theorem 1.1 (Theorem 3.7 Froyland and Rock (2025)) *Consider an eigenpair* (λ, f).

1. *If f has at least n nodal domains then $h_n \le \sqrt{-2\lambda}$.*
2. *For each $k = 1, \ldots, n$, let $N_k \subset M$ denote the k^{th} nodal domain and let $A_k(c_k) := \{x \in N_k : f(x)^2 \ge c_k\} \subset M$ denote the superlevel set of f^2 restricted to N_k. Then for a positive measure set of $(c_1, \ldots, c_n) \in \mathbb{R}^n$, one has*

$$H(A_1(c_1), \ldots, A_n(c_n)) \le \sqrt{-2\lambda}. \tag{7}$$

Part 1 says that the closer λ is to zero, the more easily M can be decomposed into n pieces, each with a small Cheeger ratio. This is consistent with our discussion of the heat equation because the closer λ is to zero, the more slowly the initial condition $u(0, x) = f(x)$ will relax to equilibrium. When f has n nodal domains, there are

n subsets of M containing obstacles that slow down this equilibriation. In Fig. 1, the left two panels show that f_2 and f_3 have two nodal domains (separated by the white strip), each roughly half of M, f_4 has four nodal domains, and f_5 has five nodal domains. Thus, by part 1 of Theorem 1.1, using the eigenvalues quoted in the caption of Fig. 1 we have $h_2 \leq 8.18$, $h_3 \leq h_4 \leq 10$, and $h_5 \leq 12.84$.

Part 2 of Theorem 1.1 provides a recipe for constructing n-packings of M with low Cheeger ratios, namely level sets of f^2 individually restricted to each nodal domain. In this regard, note that the nodal lines of f_2 and f_3 in Fig. 1 are slightly offset from the horizontal and vertical, respectively; this is not an accident. If the nodal lines were horizontal and vertical, they would split the small disks at the periphery of M, leading to larger Cheeger ratios for the corresponding blue and red nodal domains. Cheeger inequalities without nodal domain counts are available for boundaryless manifolds Miclo (2015), but they are much weaker.

2.4 *Sparse Eigenbasis Approximation (SEBA)*

Finding good candidate n-packings of a manifold with superlevel sets of individual nodal domains of f^2 can be practically challenging. Computing *individual* nodal domains can be nontrivial because of the connectivity requirement in the definition. Moreover, because the irregularity of the $\{f_k\}$ grows with k, we expect superior n-packings to be obtained from as few f_k as possible. The n^{th} eigenfunction f_n contains *at most* n nodal domains (by the Courant Nodal Domain Theorem, Courant and Hilbert (1953)), and may contain fewer.

By taking linear combinations of $\{f_1, \ldots, f_n\}$ in combination with soft thresholding we can create a sparse collection of functions $\{s_1, \ldots, s_n\}$ that approximately span $\text{span}\{f_1, \ldots, f_n\}$. Each function s_k will support an individual element of an n packing; see Froyland and Rock (2025) for a discussion of SEBA in the manifold setting, including a direct analogue to Part 2 of Theorem 1.1 for sets $A_k(c_k) := \{x \in M : s_k(x) \geq c_k\}$. Implementational details, further discussion of constructing n-packings, and the short code may be found in Froyland et al. (2019) and at https://github.com/gfroyland/SEBA. Figure 2 illustrates this approach using eigenfunctions $f_1, \ldots, f_4$ (i.e. the constant function f_1 along with the left 3 panels of Fig. 1). One can be sure that SEBA has worked well when the sparse basis

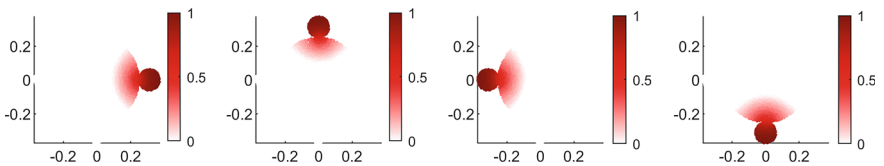

Fig. 2 Sparse functions $s_1, \ldots, s_4$ output by the SEBA algorithm applied to $f_1, \ldots, f_4$. One has $\text{span}\{s_1, \ldots, s_4\} \approx \text{span}\{f_1, \ldots, f_4\}$ and may define A_k as a superlevel set of s_k, $k = 1, \ldots, 4$

functions $s_1, \ldots, s_n$ are all nonnegative. For example, applying SEBA to the collections of eigenfunctions $\{f_1, \ldots, f_n\}$, $2 \le n \le 5$ in Fig. 1, we obtain nonnegative basis functions for $n = 2, 4, 5$, but not $n = 3$ because there is no natural encoding of three sets in the leading three eigenfunctions.

Finally, we remark that the k-means algorithm is often used to post-process spectral clustering output. Our Laplacian eigenanalysis is precisely spectral clustering on the manifold M and SEBA provides a more robust post-processing step for discrete data than k-means. This is because SEBA will zero out data points that do not naturally belong to a cluster and SEBA provides clear guidelines for rejecting spurious sparse vectors (s_k with large negative values should be rejected); see Froyland et al. (2019) for details.

3 Dynamical Systems and Finite-Time Coherent Sets

We now introduce dynamics on manifolds and explain how dynamic versions of the spectral geometric properties of the Laplacian discussed in the previous section can identify coherent sets in dynamical systems. We consider a domain[2] M and one-parameter family[3] of smooth transformations $\Phi^t : M \to \Phi^t(M)$ denoting the evolution of points in M for a duration $t \in [0, \tau]$. For example, Φ^t could be the flow generated by a time-dependent ODE $\dot{x} = F(t, x)$, i.e. $\frac{d}{ds}\Phi^s|_{s=t}(\cdot) = F(t, \cdot)$. For simplicity of presentation we will continue to assume that M is d-dimensional compact submanifold of $\mathbb{R}^d$ endowed with the Euclidean metric, and that Φ^t preserves Euclidean volume. All of the constructions below naturally generalise to evolving families of Riemannian manifolds endowed with metrics independent of the dynamics, and to non-Riemannian-volume-preserving dynamics Φ^t Froyland and Kwok (2020).

3.1 Chaotic Advection, Mixing, and Coherence

The above finite-time evolution is often highly heterogeneous in space, and an important characterisation of the dynamics is the decomposition of the phase space into coherent regions that are relatively non-dispersive and mixing regions that rapidly disperse. We will say that a subset A of M is *coherent over the interval* $[0, \tau]$ if the boundary of A remains small relative to its volume. In the presence of small diffusion–for example considering an SDE with drift F as above and small diffusion, $dX_t = F(t, X_t) + (\epsilon^2/2)dW_t$–the points in the interior of A that escape from A are those that are near to the boundary, and any escape is by diffusion alone; see Fig. 3. This approach is consistent with methods of quantification of advective mix-

[2] In dynamical systems, the domain is often called phase space or state space.

[3] All of the constructions developed here have obvious analogues if time is discrete.

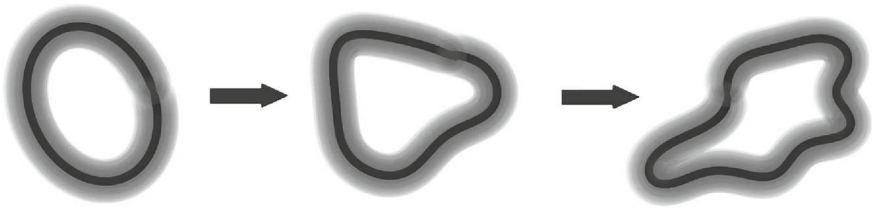

Fig. 3 Advective mixing in the presence of small diffusion. Left: a two-dimensional boundary shown as a black curve with ϵ-scale band of diffusion over the boundary. Centre and Right: under advective evolution with Φ^t the boundary evolves in a deterministic way, with the amount of mixing proportional to the size of the boundary

ing Ottino (1989), Pierrehumbert (1994), Mathew et al. (2005), Thiffeault (2012) and Haller and Beron-Vera (2013).

We briefly illustrate this idea in a simple example, based on the rotating double gyre Mosovsky and Meiss (2011). Let $M = [0, 1]^2$ and Φ^t, $t \in [0, 1]$ be generated by the time-dependent ODE given by Eq. (3.1) Mosovsky and Meiss (2011). Figure 4(left) shows two sets $A_1, A_2 \subset M$ with small Neumann Cheeger ratio (see (5)): $H(A_1) = H(A_2) = 1/(1/2) = 2$. In fact, these A_1, A_2 form an optimal 2-packing of M, with $h_2 = 2$. But when these sets are evolved under the dynamics Φ^t, their boundaries grow significantly: $H(\Phi^1(A_1)) = H(\Phi^1(A_2)) \approx 8.3057/(1/2) = 16.6114$. Thus, in a short period of time, exchange of fluid between the two shaded regions in Fig. 4 rapidly increases.

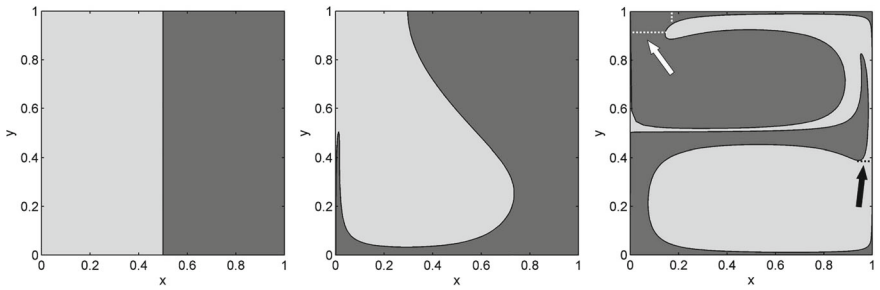

Fig. 4 Illustration of the dynamics of the rotating double gyre flow. Left: at time $t = 0$ we initialise with two sets A_1, A_2 given by the left and right halves of the domain. Centre: at time $t = 1/2$ the boundaries of A_1, A_2 begins to grow. Right: at the final evolution time $t = 1$ the boundaries have grown significantly, leading to advective mixing in the neighbourhoods of the boundaries. The arrows indicate asymmetries in the evolved sets. Image from Froyland and Padberg-Gehle (2014)

3.2 Dynamic Cheeger Ratios and Coherent Sets

In view of Fig. 3 and the above discussion on advective mixing, coherent subsets A of M are those for which the total exchange of trajectories from A to its complement, relative to the volume of A, is minimal. This total exchange is proportional to the average Cheeger ratio of $\Phi^t(A)$ over the time interval $[0, \tau]$. We therefore introduce a *dynamic Cheeger ratio* Froyland (2015) and Froyland and Rock (2025) in analogue to the Neumann Cheeger ratio (5):

$$H^D(A) := \frac{1}{\tau}\int_0^\tau \frac{|\partial(\Phi^t(A)) \cap \Phi^t(\mathring{M})|}{|\Phi^t(A)|}\, dt. \tag{8}$$

For $A = A_1$ in Fig. 4, the numerator of (8) is simply the average length of the common boundary of A_1 and A_2; in particular, we do not include the boundaries[4] of $\Phi^t(M)$. We therefore seek finite-time coherent sets as minimisers of $H^D(\cdot)$. Our approach will mirror our search for minimisers of $H(\cdot)$ in Sect. 2. Directly minimising the geometric expression (8) is difficult, so we instead develop a theory of *dynamic spectral geometry*, and find good candidates for minimisers with eigenfunctions of a *dynamic Laplace operator*.

3.3 Encoding Evolved Geometries with Pullback Metrics

If our dynamic Laplace operator is going to help us find finite-time coherent sets, it must incorporate the deformed geometry of M under the action of Φ^t. For example, in the right panel of Fig. 4 we computed the length of the evolved common boundary (let us denote this family of curves $\Phi^t(\Gamma)$)–with respect to the Euclidean metric–to be about 8.3057 at time $t = 1$. To encode this deformation using a single operator acting on functions on M, we need to create a geometry on M that reflects the Euclidean distances between points in $\Phi^t(M)$. Under such a geometry, the length of $\Phi^t(\Gamma)$–computed with the Euclidean metric e on $\Phi^t(M)$–would be the length of Γ on M. This geometry is described by the pullback metric $(\Phi^t)^* e =: g_t$. Because we have assumed that Φ^t preserves volume on M, the metric g_t has a particularly simple pointwise description in standard coordinates, namely it can be represented by the $d \times d$ positive-definite matrix $\mathbf{g}_t(x) := (D\Phi^t(x))^\top D\Phi^t(x)$ for $x \in M$, which is known as the Cauchy–Green tensor in fluid dynamics.

[4] In Sect. 3.6 we will discuss the Dirichlet dynamic Cheeger ratio and the circumstances in which to use each of the two types of dynamic Cheeger ratios.

Continuing with the example in Fig. 4, let $\gamma(s) = (1/2, s), 0 \le s \le 1$ be a parameterisation of Γ in the square domain M. The length of Γ according to the Euclidean metric on the square is

$$\int_0^1 \|\gamma'(s)\|_2 \, ds = \int_0^1 1 \, ds = 1. \tag{9}$$

We may similarly compute the length of $\Phi^t(\Gamma)$. A little notation before beginning: on the tangent space of M at $x \in M$, denoted $T_x M$, we define an inner product $\langle \cdot, \cdot \rangle_{g_t}$ by $\langle u, v \rangle_{g_t} = u^\top \mathbf{g}_t(x) v$ for $u, v \in T_x M$. We now may calculate the length of $\Phi^t(\Gamma)$ as:

$$\begin{aligned}
\int_0^1 \|D\Phi^t(\gamma(s)) \cdot \gamma'(s)\|_2 \, ds &= \int_0^1 \left(\gamma'(s)^\top (D\Phi^t(\gamma(s)))^\top D\Phi^t(\gamma(s)))\gamma'(s)\right)^{1/2} \, ds \\
&= \int_0^1 \left(\gamma'(s)^\top \mathbf{g}_t(\gamma(s))\gamma'(s)\right)^{1/2} \, ds \\
&= \int_0^1 \langle \gamma'(s), \gamma'(s) \rangle_{g_t}^{1/2} \, ds \\
&= \int_0^1 \|\gamma'(s)\|_{g_t} \, ds. \qquad (10)
\end{aligned}$$

Note that (10) is almost identical in form to the left-hand-side of (9); we have merely *changed the norm* in which we compute the length of the derivative vector γ'.

As discussed above, the metric g_t is the *pullback metric* we obtain by pulling back the Euclidean metric e with the diffeomorphism Φ^t. As Riemannian manifolds, Φ^t is an isometry between (M, g_t) and $(\Phi^t(M), e)$, because lengths computed in the tangent space $T_x M$ with g_t are identical to lengths computed in the tangent space $T_{\Phi^t x}\Phi^t(M)$ with the Eulidean metric. The concept of pullback metrics allows us to *transfer geometry on* $\Phi^t(M)$ *back onto* M *by applying a new metric on* M.

If M is d-dimensional, $d \ge 2$, we need to compute areas of $d-1$-dimensional hypersurfaces. We proceed in an analogous way to the two-dimensional case above, again using the same pullback metric g_t. A bounded $d-1$-dimensional hypersurface S can be parameterised by a function $\gamma : [0, 1]^{d-1} \to \mathbb{R}^d$. To compute the $d-1$-dimensional area of the hypersurface $\Phi^t(S)$ one integrates

$$\int_{[0,1]^{d-1}} \left(\det(\mathrm{Gram}_{g_t}(\mathbf{s}))\right)^{1/2} \, d\mathbf{s},$$

where $d\mathbf{s} = ds_1 \, ds_2 \cdots ds_{d-1}$ and $\mathrm{Gram}_{g_t}(\mathbf{s})_{ij} = \langle \partial\gamma(\mathbf{s})/\partial s_i, \partial\gamma(\mathbf{s})/\partial s_j \rangle_{g_t}$ is the $(d-1) \times (d-1)$ Gram matrix formed by partial derivatives of the parameterisation γ with respect to the $d-1$ parameters. In the case $d = 2$, the Gram matrix is scalar-valued and we obtain (10).

3.4 The Dynamic Laplacian

Consider a Riemannian manifold (M, g) with volume form ω_g. The divergence of a vector field X on M is the unique scalar-valued function $\mathrm{div}(X) : M \to \mathbb{R}$ satisfying $\mathrm{div}(X)\omega_g = X(\omega_g)$, where $X(\omega_g)$ is the Lie derivative of ω_g with respect to X. The gradient $\nabla_g f$ is the unique vector field on M satisfying $\langle \nabla_g f, X\rangle_g = X(f)$ for all vector fields X on M, where $X(f)$ is the Lie derivative of f with respect to X. The Laplace–Beltrami operator with respect to a metric g is defined as $\Delta_g f = \mathrm{div}(\nabla_g f)$. The dynamic Laplacian Froyland (2015) for Φ^t over the interval $[0, \tau]$ combines Laplace–Beltrami operators for the metrics g_t, $0 \le t \le \tau$: Let

$$\Delta^D f := \frac{1}{\tau}\int_0^\tau \Delta_{g_t} f \, dt, \tag{11}$$

where Δ_{g_t} is the Laplace–Beltrami operator on M with respect to the metric g_t. Recall that each of the pullback metrics is a metric on M and so (11) is an average of Laplace–Beltrami operators, each of which is defined on M.

We now consider the (Neumann) eigenproblem for the dynamic Laplacian

$$\Delta^D f = \lambda f, \tag{12}$$

with $\int_0^\tau \langle \nabla_{g_t} f(x), n(x)\rangle_e \, dt = 0$ for all $x \in \partial M$. Because Δ^D is a linear combination of Laplace–Beltrami operators, it inherits the symmetry and ellipticity properties of these operators and therefore also many of their spectral properties. The eigenspectrum of Δ^D is a countable sequence $0 = \lambda_1^D \ge \lambda_2^D \ge \cdots$ of negative numbers diverging to $-\infty$, and the leading eigenfunction is constant. The eigenfunctions $f_1^D, f_2^D, \ldots$ form a complete orthonormal set for $L^2(M)$ Froyland (2015).

3.5 Dynamic Cheeger Problems and Dynamic Laplace Eigenfunctions

We now discuss the links between dynamic Cheeger problems and eigenfunctions of the dynamic Laplacian, in analogy to Sect. 2.3. Recall the dynamic Cheeger ratio $H^D(A)$ for a set $A \subset M$ in Eq. (8). When $H^D(A)$ is low, this means that the relative evolved boundary size of $\Phi^t(A)$ is low on average over the time interval $[0, \tau]$. Given a dynamical evolution of a manifold M under Φ^t, $0 \le t \le \tau$, to find n finite-time coherent sets we search for dynamic n-packings $\{A_1, \ldots, A_n\}$ of M with low $H^D(\{A_1, \ldots, A_n\}) := \max_{1\le k\le n} H^D(A_k)$. To quantify how well a choice can be

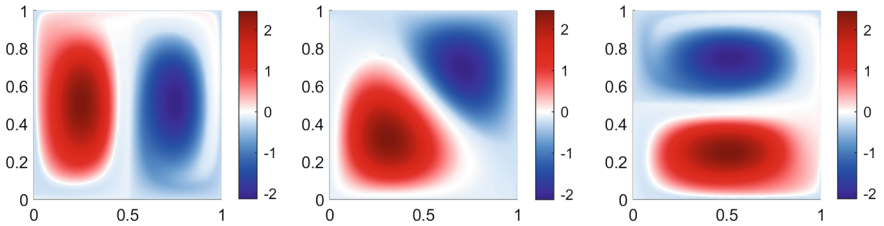

Fig. 5 Left: Leading nontrivial eigenfunction f_2^D for the rotating double gyre shown in Fig. 4. Centre: Pushforward of f_2^D by 1/2 a time unit to $t = 1/2$, namely $f_2^D \circ (\Phi^{1/2})^{-1}$. Right: Pushforward of f_2^D by 1 time unit to $t = 1$, namely $f_2^D \circ (\Phi^1)^{-1}$

made for a given finite-time dynamics we define the *dynamic Cheeger constant* of (Φ^t, M) for $t \in [0, \tau]$ as[5]

$$h_n^D := \inf\{H^D(\{A_1, \ldots, A_n\}) : A_1, \ldots, A_n \subset M \text{ are disjoint}\}. \tag{13}$$

We will use eigenfunctions of the dynamic Laplace operator to find good candidate dynamic n-packings.

Theorem 1.2 (Theorem 3.19 Froyland and Rock (2025)) *Consider an eigenpair (λ^D, f^D) for the dynamic Laplacian Δ^D.*

1. *If f^D has at least n nodal domains then $h_n^D \le \sqrt{-2\lambda^D}$.*
2. *For each $k = 1, \ldots, n$, let $N_k \subset M$ denote the k^{th} nodal domain and let $A_k(c_k) := \{x \in N_k : f^D(x)^2 \ge c_k\} \subset M$ denote the superlevel set of $(f^D)^2$ restricted to N_k. Then for a positive measure set of $(c_1, \ldots, c_n) \in \mathbb{R}^n$, one has*

$$H^D(A_1(c_1), \ldots, A_n(c_n)) \le \sqrt{-2\lambda^D}. \tag{14}$$

The interpretation of this theorem follows the discussion of Theorem 1.1. By (8), part 1 says that the closer the eigenvalue λ^D is to zero, the more easily we can decompose M into n disjoint sets $A_1, \ldots, A_n$ such that the average boundary size of each evolved family $\{\Phi^t(A_k)\}_{0 \le t \le \tau}$ is small. Part 2 provides an ansatz for identifying such families using the dynamic Laplacian eigenfunctions.

Figure 5(left) shows the leading nontrivial eigenfunction f_2^D of the dynamic Laplacian for the rotating double gyre dynamics discussed in Sect. 3.1. If we use superlevel sets of $(f_2^D)^2$ as per Part 2 of Theorem 1.2 we will obtain two sets A_1, A_2, which we can visualise as the deeper red and blue areas in Fig. 5(left). We may push these red and blue regions forward under the dynamics Φ^t in the usual way, to obtain a function $f_2^D \circ (\Phi^t)^{-1}$ on $\Phi^t(M)$; this is shown at $t = 1/2$ and $t = 1$ in Fig. 5(centre & right). These are snapshots of the evolving coherent families of sets $\Phi^t(A_1)$, $\Phi^t(A_2)$ for $t \in [0, \tau]$. The red and blue structures that we see follow the

[5] See Froyland (2015) for $n = 2$, and Froyland and Rock (2025) for general n.

dynamics visible in Fig. 4, but the evolved boundaries of our sets A_1, A_2 are much smaller than the gray regions in Fig. 4, and therefore much more coherent.

3.6 Dirichlet Boundary Conditions

We have discussed the eigenproblem for the dynamic Laplace operator with Neumann boundary conditions. These boundary conditions are particularly appropriate when the domains $\Phi^t(M)$ have a physically meaningful boundary and the finite-time coherent sets one seeks may touch the boundary. In situations where one explicitly wishes to exclude coherent sets that touch the boundary, or where the boundary of $\Phi^t(M)$ may be arbitrary because the manifolds $\Phi^t(M), 0 \le t \le \tau$ live inside a larger ambient space, then Dirichlet boundary conditions–i.e. the eigenfunctions are zero on ∂M–may be applied. The main changes are as follows Froyland and Junge (2018) and Froyland and Rock (2025):

1. λ_1^D is now strictly negative and f_1^D may be chosen to be strictly positive, except on ∂M where it vanishes.
2. Equation (8) becomes

$$h^D(A) := \frac{1}{\tau}\int_0^\tau \frac{|\partial(\Phi^t(A))|}{|\Phi^t(A)|}\,dt, \tag{15}$$

   so that the entire boundaries of $\Phi^t(A)$ contribute, even if part of these boundaries coincide with the boundaries of $\Phi^t(M)$. With this adjustment Theorem 1.2 holds for eigenpairs (λ^D, f^D) satisfying Dirichlet boundary conditions.
3. Dirichlet boundary conditions force the superlevel sets $A_k(s_k)$ in Theorem 1.2 to not touch the boundary of M.

4 Computation

Our starting point will be several trajectories of a time-dependent dynamical system $x_{i,t_l}, 0 \le i \le N, 1 \le l \le T$, where $x_{i,t_l} \in \Phi^{t_l}(M)$ is the location of the i^{th} trajectory at time t_l. These trajectories may be numerically generated from a known dynamical system Φ^t or they may be state observations made from some unknown dynamics. The basic building block we need to approximate numerically is the Laplace–Beltrami operator for the pullback metric g_t. Estimating $g_t = (\Phi^t)^* e$ numerically requires high-resolution data in order to accurately estimate the matrix $(D\Phi^t)^\top D\Phi^t$ (see Sect. 3.3). To avoid this high computational cost, we instead use the important relationship

$$\Delta_{g_t} f = \Delta_e(f \circ \Phi^{-t}) \circ \Phi^t = (\Phi^t)^*(\Delta_e(\Phi^t_* f)), \tag{16}$$

where $(\Phi^t)_* f = f \circ \Phi^{-t}$ is the pushforward of f under Φ^t, and $(\Phi^t)^* f = f \circ \Phi^t$ is the pullback of f under Φ^t.

Because we wish to estimate Laplace–Beltrami operators, a natural established numerical framework is the finite-element method. This has been adapted to the dynamic Laplacian in Froyland and Junge (2018). We wish to solve the eigenproblem $\Delta^D f = \lambda f$ with the dynamic Neumann boundary condition. We begin by integrating both sides of this equality against a test function $\tilde{f}$ using the volume measure V_e for the metric e, i.e. Lebesgue measure. This yields the weak form of the dynamic Laplacian eigenproblem, namely

$$\int_M (\Delta^D f)\tilde{f}\, dV_e = \lambda \int_M f\tilde{f}\, dV_e \text{ for all test functions } \tilde{f}. \tag{17}$$

The weak form has the immediate advantage of eliminating the Neumann[6] boundary condition Froyland and Junge (2018). By (11) we see that the LHS of (17) is a time integral of $\int_M (\Delta_{g_t} f)\tilde{f}\, dV_e$ over the time interval $[0, \tau]$. Taking one of these terms and applying (16) we obtain

$$\int_M (\Phi^t)^* (\Delta_e (\Phi^t_* f))\tilde{f}\, dV_e.$$

Changing variables with $\Phi^t(M)$ and performing integration by parts yields

$$-\int_{\Phi^t(M)} \nabla(\Phi^t_* f) \cdot \nabla(\Phi^t_* \tilde{f})\, dV_e. \tag{18}$$

We now introduce the FEM discretisation, which is a basis of nodal hat functions $\{\varphi_i^{t_l}\}_{j=1}^N$ defined on a mesh of $\Phi^t(M)$ with N nodes $\{\Phi^{t_l} x_i\}_{i=1}^N$, namely the trajectories we were initially supplied with. As a real-world example, Fig. 6 shows a mesh of around 3000 deep-ocean ARGO floats at December 2013. Each $\varphi_i^{t_l}$ is locally supported on mesh elements containing the node $\Phi^{t_l} x_i$ as a vertex, and $\varphi_i^{t_l}(\Phi^{t_l} x_i) = 1$, $i = 1, \ldots, N$. We form a (sparse) matrix $D_{ij}^{t_l}$ by evaluating (18) with $\Phi^t_* f = \varphi_i^{t_l}$ and $\Phi^t_* \tilde{f} = \varphi_j^{t_l}$, $1 \le i, j, \le N$:

$$D_{ij}^{t_l} := -\int_{\Phi^{t_l}(M)} \nabla\varphi_i^{t_l} \cdot \nabla\varphi_j^{t_l}\, dV_e, \tag{19}$$

which is a standard FEM integral. We have thus handled the left hand side of (17), obtaining the matrices $D_{ij}^{t_l}$ for t in a discrete set of points $\{t_1, \ldots, t_T\}$ sampled from the time interval $[0, \tau]$.

[6] Dirichlet boundary conditions are also eliminated by simply restricting the functions f to those that vanish at the boundary of M.

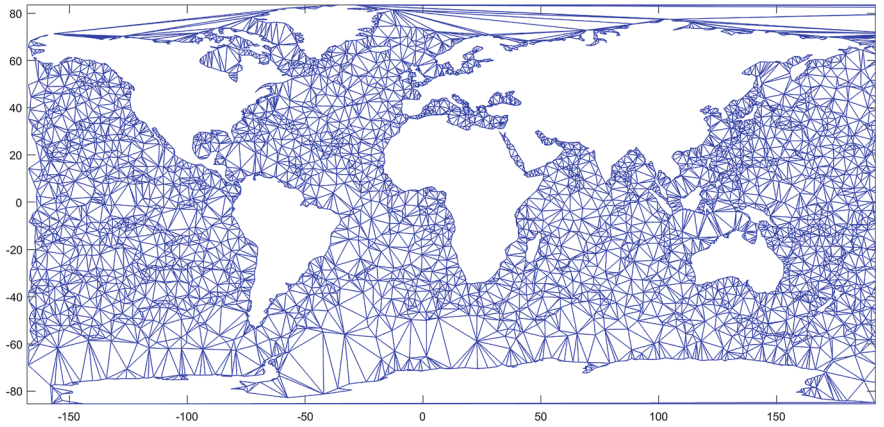

Fig. 6 Delaunay mesh of approximately 3000 Argo float positions (nodes) at December 2013. Abernathey et al. (2022) identified the major long-term coherent sets at 1000 m depth by meshing Argo float locations $x_i^{t_l}$ at each month t_l over six years

For the right hand side of (17) we substitute $f = \varphi_i^0$, $\tilde{f} = \varphi_j^0$ to obtain the $N \times N$ sparse matrix $M_{ij} := \int_M \nabla\varphi_i^0 \cdot \nabla\varphi_j^0 \, dV_e$. One now solves the discretised problem

$$\left(\frac{1}{T}\sum_{l=1}^{T} D^{t_l}\right) v = \lambda M v \tag{20}$$

to find eigenvalues λ and corresponding basis coefficient vectors v, and reconstructs eigenfunctions on M as $f = \sum_{j=1}^{N} v_j \varphi_j^0$.

This FEM approach is very robust and accurate, and is numerically efficient for dimensions 2 and 3; it has been used for all eigenfunction images in this chapter. It has been tested in many situations on sparse trajectory data Froyland and Junge (2018), Froyland et al. (2019) and Abernathey et al. (2022) and is insensitive to highly varying densities of trajectory data sampling (see e.g. Fig. 10 Froyland and Junge (2018)). The method also copes very well with situations where observed trajectories are incomplete and many observations are missing. Figure 12 Froyland and Junge (2018) illustrates this for the rotating double gyre eigenfunction in Fig. 5.

The identification of coherent regions in the global ocean at 1000 m depth from Argo drifter observations is an example of a real-world computation requiring a method than can handle sparsity, varying data density, and missing data. Figure 6 illustrates the sparsity and the varying data density. Figure 7 displays eigenfunctions of the dynamic Laplacian built from around 8000 individual deep-sea Argo float trajectories, where only 3000 floats are providing location signals in any given month, and only 10% of floats continue to transmit over the six years from 2011–2017. In Fig. 7 the dynamic Laplacian was estimated in a single computation constructed from the many partial and non-overlapping-in-time trajectories of Argo floats.

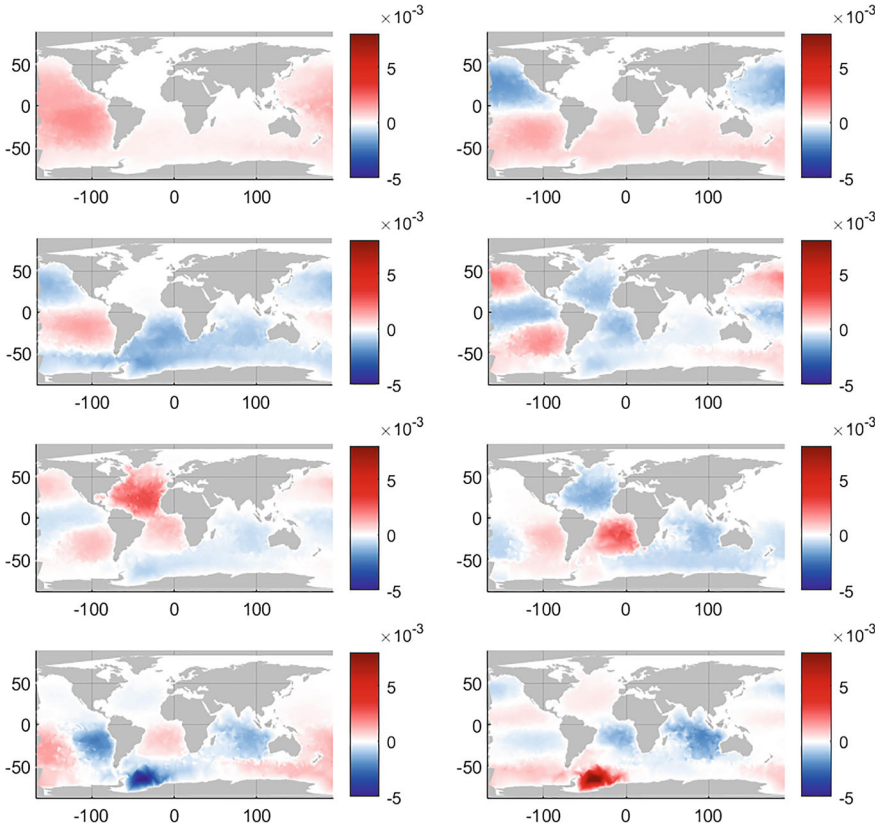

Fig. 7 Leading eight eigenfunctions $f_1^D, \ldots, f_8^D$ of the dynamic Laplacian computed from Argo trajectory data over the six years 2011–2017. The eigenfunctions are represented as scalar fields at January 2011. Image from Abernathey et al. (2022)

Finally, we note that there are no free parameters to choose and the method preserves the symmetric structure of the original eigenproblem. MATLAB code and demos are available at https://github.com/gaioguy/FEMDL. Alternate FEM-based strategies and generalisations to weighted manifolds are discussed in detail in Froyland and Junge (2018).

4.1 Alternate Approaches to Estimate Δ^D and Related Objects

Ulam and finite-difference

In Froyland (2015) the pushforwards and pullbacks Φ^t_* and $(\Phi^t)^*$ in (16) were estimated with Ulam-discretisations and the Laplacian Δ_e was discretised by finite-difference. This approach was generalised to the measure-preserving setting on weighted manifolds in Froyland and Kwok (2020), and we refer the reader to Froyland (2015) and Froyland and Kwok (2020) for a full explanation of this approach. A large number of points were used in Froyland (2015) and Froyland and Kwok (2020) to estimate the Ulam matrices representing Φ^t_* and $(\Phi^t)^*$, but similar results can be obtained with many fewer points. Because this approach is grid-based, a reasonably uniform sampling of initial conditions is needed.

Radial basis function collocation

The paper Froyland and Junge (2015) replaces the indicator function bases of the Ulam approach with smooth radial basis functions and proceeds to approximate Φ^t_*, $(\Phi^t)^*$ and Δ^D by collocation. The shape parameter–which determines the width of the basis functions–needs to be selected and the method requires a suitable positioning of basis function centres.

Related objects

We noted earlier that the dynamic Laplacian is a type of spectral clustering on manifolds. A lower-order approach, compared to FEM, is to construct a graph Laplacian based on initial conditions of trajectories Hadjighasem et al. (2016), where edges are weighted according to a dynamic distance introduced in Froyland and Padberg-Gehle (2015). In a similar direction, Padberg-Gehle and Schneide (2017) apply constant edge weights that record if one trajectory comes within a threshold of another. In these approaches, cutoffs for weights or distances need to be selected. Varying point density can influence the results, whereas this is automatically adjusted for in the FEM approach, which can itself be interpreted as a graph Laplacian, see the discussion in Sect. 3.4 Froyland and Junge (2018).

5 Finding Multiple Coherent Sets: SEBA

In Sect. 2.4 we discussed how the SEBA algorithm could create sparse linear combinations of Laplace eigenfunctions to automatically separate individual nodal domains and create additional candidate sets for n-packings. The same methodology can be applied to the eigenfunctions $f^D_1, \ldots, f^D_n$ of the dynamic Laplace operator to cre-

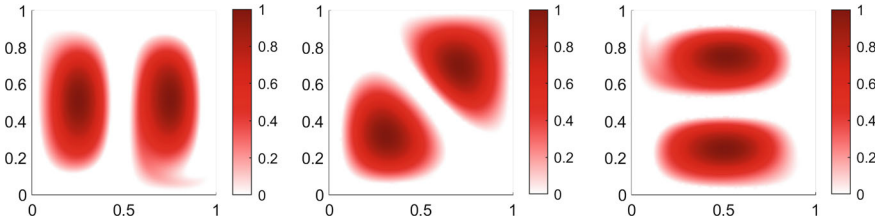

Fig. 8 Left: Sum of two SEBA vectors arising from SEBA applied to the dynamic Laplacian eigenfunctions $f_1^D \equiv \mathbf{1}$ and f_2^D from Fig. 5(left) for the rotating double gyre. Centre: Pushforward of the left panel by 1/2 a time unit with $\Phi_*^{1/2}$ to $t = 1/2$. Right: Pushforward of the left panel by 1 time unit with Φ_*^1 to $t = 1$

ate a sparse basis $\{s_1^D, \ldots, s_n^D\}$ whose span approximates span$\{f_1^D, \ldots, f_n^D\}$. Each sparse function s_k^D supports a single coherent set and in this way we disentangle individual coherent sets from multiple eigenfunctions. Performance guarantees for candidate coherent sets $A_k(c_k) := \{x \in M : s_k^D(x) \geq c_k\}$ are available in direct analogue to Part 2 of Theorem 1.2; see Froyland and Rock (2025). By applying SEBA to the leading eigenfunctions $f_1^D \equiv \mathbf{1}$ and f_2^D from Fig. 5(left), we obtain two (sparse) SEBA vectors s_1^D, s_2^D taking values in [0, 1]. The sum of these two vectors is shown in red in Fig. 8(left) and their pushforwards under Φ^t shown in the centre and right panels.

Similarly, we can apply SEBA to the eight eigenfunctions shown in Fig. 7 to separate the eight dominant coherent regions in the global ocean at 1000 m depth; see Fig. 9. With the strongest coherent regions clearly separated, it has been noted in Abernathey et al. (2022) that they are largely aligned with the major ocean boundaries, but with a bias toward the east of the basin due to the stronger-mixing currents on their western boundaries. The Weddell gyre off the Antarctic coast is highlighted. The North and South Atlantic is separated by the equatorial countercurrent. The subpolar and subtropical gyres in the North Pacific are separated and the South Pacific Gyre splits into eastern and western components. The reader is referred to Abernathey et al. (2022) for more detailed analysis. These results are broadly consistent with coarser results Froyland and Padberg-Gehle (2015) and Banisch and Koltai (2017) using floating drifters on the ocean surface.

Sometimes the trajectory data is such that it does not support exactly n coherent sets in the leading n eigenfunctions. We have already seen this issue in Sect. 2.4 where the eigenfunctions in Fig. 1 supported decompositions into 2, 4, and 5 regions, but not 3 regions. The SEBA vectors themselves contain a simple "reliability" test, namely the magnitude of the most negative value; see Sects. 4.1 and 4.2 in Froyland et al. (2019)). A discussion of assessing spectral gaps is in Sect. 4.2 Froyland et al. (2019).

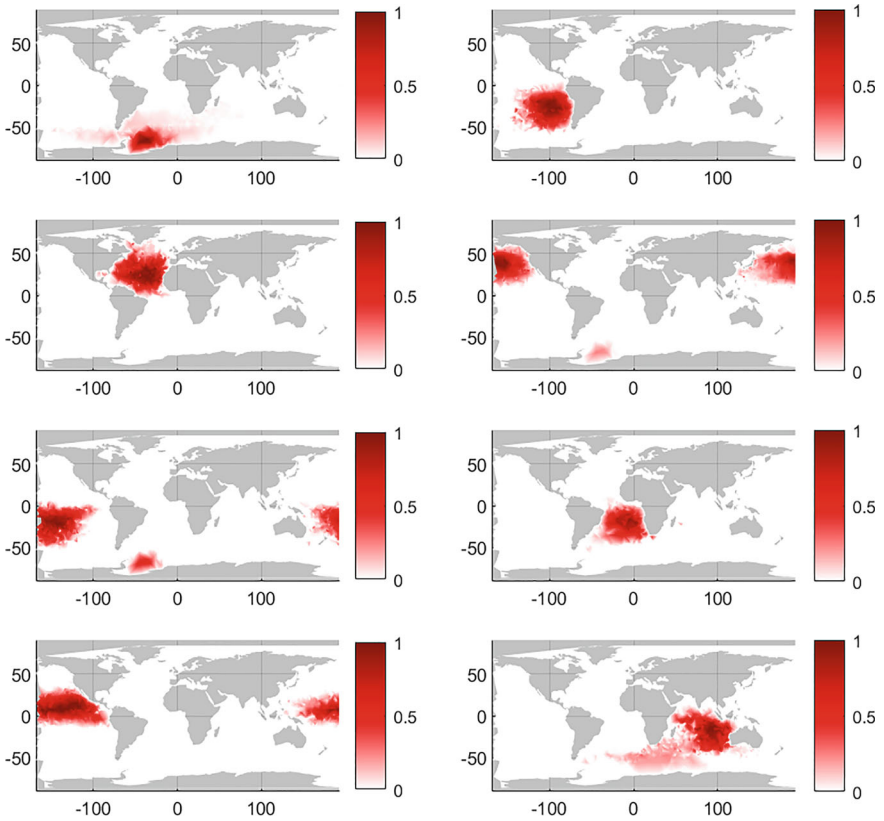

Fig. 9 Result of SEBA applied to the eight eigenfunctions shown in Fig. 7 to construct eight SEBA vectors $s_1^D, \ldots, s_8^D$. Each SEBA vector is supported on an individual coherent region of the ocean at January 2011, each of which remains coherent under ocean currents at 1000 m depth over a six-year duration from 2011–2017. Image from Abernathey et al. (2022)

6 Weights, Dynamic Metrics, and Zero-Noise Limits

Measure-preserving dynamics on weighted Riemannian manifolds

We have mainly discussed the special situation where we have a flat manifold M embedded in Euclidean space and Φ^t preserves Euclidean volume. We may also consider a dynamic Laplacian on a sequence of weighted Riemannian manifolds (M_t, h_t, m_t), where $M_t = \Phi^t(M)$, m_t is a Riemannian metric on M_t and $h_t : M_t \to \mathbb{R}^+$ is a positive weight function. The measure of $A \subset M$ is $\int_A h_0 \, dV_{m_0}$, where V_{m_0} is the volume measure induced by the metric m_0. The dynamics Φ_t can be made *measure preserving* by defining positive weight functions h_t so that the measure of a set $A \subset M$ equals the measure of $\Phi^t(A)$ on M_t: $\int_{\Phi^t A} h_t \, dV_{m_t} = \int_A h_0 \, dV_{m_0}$, for $0 < t \leq \tau$. One now defines

$$\Delta^D f := \frac{1}{\tau}\int_0^\tau \Delta_{h_t,g_t} f \, dt, \tag{21}$$

where $\Delta_{h,g} f = \frac{1}{h}\mathrm{div}(h\nabla_g f)$ is the h-weighted Laplace–Beltrami operator for g.

In the context of fluid flow, this weighted setting is useful for finding finite-time coherent sets in compressible fluid flows, where Euclidean volume is not preserved, but fluid mass is conserved. The weight functions h_t track the evolution of an initial fluid density h_0 (with respect to V_{m_0}) under Φ^t. This weighted setup also covers the situation where one is interested in coherence of a passive tracer carried by the flow and this passive tracer is initialised with a spatial concentration different to the fluid density; e.g. in the volume-preserving case, the fluid density would be uniform. One sets h_0 to be the initial density of the passive tracer, computed with respect to V_{m_0}, and future passive tracer densities are tracked by the weight functions h_t. We refer the reader to Froyland and Kwok (2020) for the relevant constructions and Froyland and Junge (2018) for numerical computations.

Is the dynamic Laplacian a Laplace–Beltrami operator?

One may ask whether there is a dynamic Riemannian metric g^D so that $\Delta^D = \Delta_{g^D}$. Let us return to the simplified setting of a flat manifold $M \subset \mathbb{R}^d$ endowed with the Euclidean metric e and Φ^t Lebesgue volume preserving; i.e. $V_{g_t} = V_e$ for all $t \in [0, \tau]$. The appropriate construction for g^D is to set $g^D = \left(\frac{1}{\tau}\int_0^\tau g_t^{-1}\, dt\right)^{-1}$, where g^{-1} represents the dual metric, which is a symmetric, positive-definite, bilinear form acting on pairs of covectors in cotangent space Karrasch and Keller (2020) and Karrasch and Schilling (2021). One then obtains that Δ^D is the weighted Laplace–Beltrami operator Δ_{θ,g^D}, where θ is the Radon-Nikodym derivative dV_e/dV_{g^D}. This operator acts on the weighted manifold (M, θ, g^D), where the role of the weight θ is to transform the weighted volume of a set $A \subset M$, namely $\int_A \theta \, dV_{g^D}$, into the Lebesgue volume of A, $\int_A \mathbf{1}\, dV_e$. Similar results hold for the measure-preserving case in the previous paragraph. See Karrasch and Keller (2020), Karrasch and Schilling (2021) and Froyland and Lee (2025) for details.

The dynamic Laplacian as a zero-noise limit

In Sect. 3.1 we discussed the connection between dynamic isoperimetric problems and finite-time coherent sets by considering advective dynamics with small-noise diffusive dynamics. In fact, there is a formal connection in the following sense. Finite-time coherent sets were initially introduced and motivated by minimally mixing behaviour of subsets in the presence of small noise Froyland et al. (2010), and Froyland (2013). The classical transfer operator approach Froyland (2013) uses leading singular vectors of a diffusive transfer operator $\mathcal{L}_\epsilon = \mathcal{D}'_\epsilon \circ \Phi^\tau_* \circ \mathcal{D}_\epsilon$, where $\mathcal{D}'_\epsilon$ and $\mathcal{D}_\epsilon$ are isotropic diffusion (of scale ϵ) operators on M and $\Phi^\tau(M)$, respectively. One may ask what happens in the $\epsilon \to 0$ limit. The naive approach of considering leading eigenvectors of $\lim_{\epsilon\to 0} \mathcal{L}^*_\epsilon \mathcal{L}_\epsilon$ is uninformative because this limit is the

identity operator. However, a finer analysis Froyland (2015) shows that the dynamic Laplacian emerges in an appropriate limit. For $f \in C^3(M)$

$$\lim_{\epsilon \to 0} \frac{(\mathcal{L}_\epsilon^* \mathcal{L}_\epsilon - I) f}{\epsilon^2} = C \cdot (\Delta_e + (\Phi^\tau)^* (\Delta_e (\Phi_*^t f)) = C \cdot \left(\Delta_{g_0} + \Delta_{g_\tau}\right)/2,$$

for a constant C depending on the common ϵ-scaled kernel of the diffusion operators $\mathcal{D}_\epsilon$. The RHS above is a dynamic Laplace operator where the average of Laplace–Beltrami operators is taken at the end points of the time interval $[0, \tau]$. This result was proved for general weighted Riemannian manifolds Froyland and Kwok (2020) and the limit shown to be uniform in the C^0-topology for f ranging over $C^3(M)$. Related calculations are in Banisch and Koltai (2017) and Karrasch and Schilling (2021).

7 When Finite-Time Coherent Sets Come and Go: the Inflated Dynamic Laplacian

The dynamic Laplacian tracks dynamics over a time interval $[0, \tau]$ and is an average of Laplace–Beltrami operators Δ_{g_t} with respect to metrics that are pullbacks under the dynamics. Because of this averaging property, the dynamic Laplacian succeeds in constructing dynamic isoperimetric properties of sets, and consequently finite-time coherent sets for the interval $[0, \tau]$. In several real-world situations, one may have complex dynamics in which many finite-time coherent sets emerge and disappear, each with their own distinct lifetimes. The dynamic Laplacian may be deployed on distinct time intervals to either separately find these sets and/or stitch distinct computations together across time Schneide et al. (2022). A recent alternative is to create an *inflated dynamic Laplace operator* on the spacetime manifold $[0, \tau] \times M$ for a perhaps larger τ than many of the lifetimes of the coherent sets. On each time fibre $\{t\} \times M$, one applies Δ_{g_t} (acting in space), while the distinct time fibres are tied together with a one-dimensional Laplace operator on $[0, \tau]$. The eigenfunctions of this inflated dynamic Laplacian encode the appearance and disappearance of multiple finite-time coherent sets, whose individual lifetimes can be extracted by applying the SEBA algorithm in spacetime. Much of the theory of the dynamic Laplacian can be naturally modified for the inflated dynamic Laplacian; see Froyland and Koltai (2023) and Atnip et al. (2024) for details.

7.1 Clusters in Time-Varying Networks

The dynamic Laplacian and the inflated dynamic Laplacian have counterparts on time-varying networks. For example, one may consider permutation dynamics on the nodes of the network and develop a corresponding dynamic spectral theory for

dynamic graph cuts Froyland and Kwok (2015). The concept of a supra-Laplacian Gómez et al. (2013) and Sole-Ribalta et al. (2013) is related to a spacetime network version of the inflated dynamic Laplacian and recent work Froyland et al. (2024) has shown how spacetime spectral clustering in time-varying networks can automatically detect the appearance and disappearance of clusters.

Acknowledgements The author thanks Yat-Long Lee for a careful reading of the manuscript. GF's research is partially supported by an Australian Research Council Discovery Project.

References

Abernathey, R., Bladwell, C., Froyland, G., Sakellariou, K.: Deep Lagrangian connectivity in the global ocean inferred from Argo floats. J. Phys. Oceanogr. **52**(5), 951–963 (2022)

Atnip, J., Froyland, G., Koltai, P.: An inflated dynamic Laplacian to track the emergence and disappearance of semi-material coherent sets (2024). arXiv:2403.10360

Banisch, R., Koltai, P.: Understanding the geometry of transport: Diffusion maps for Lagrangian trajectory data unravel coherent sets. Chaos Interdiscip. J. Nonlinear Sci. **27**(3), (2017)

Buser, P.: A note on the isoperimetric constant. Annales scientifiques de l'École Normale Supérieure **Ser. 4, 15**(2), 213–230 (1982)

Cheeger, J.: A lower bound for the smallest eigenvalue of the Laplacian. In: Problems in Analysis: a Symposium in Honor of Salomon Bochner, pp. 195–199. Princeton University Press (1970)

Courant, R., Hilbert, D.: Methods of Mathematical Physics, vol. 1. Interscience Publishers, New York (1953)

Froyland, G., Junge, O.: On fast computation of finite-time coherent sets using radial basis functions. Chaos **25**(8), 087409 (2015)

Froyland, G., Kalia, M., Koltai, P.: Multi-set spectral clustering of time-evolving networks using the supra-Laplacian (2024). arXiv:2409.11984

Froyland, G., Lee, Y.L.: In preparation

Froyland, G., Padberg-Gehle, K.: A rough-and-ready cluster-based approach for extracting finite-time coherent sets from sparse and incomplete trajectory data. Chaos Interdiscip. J. Nonlinear Sci. **25**(8), 087406 (2015)

Froyland, G., Padberg-Gehle, K.: Almost-invariant and finite-time coherent sets: directionality, duration, and diffusion. In: Ergodic Theory, Open Dynamics, and Coherent Structures, pp. 171–216 (2014)

Froyland, G., Rock, C.P.: Higher Cheeger ratios of features in Laplace-Beltrami eigenfunctions. Appl. Comput. Harmon. Anal. **74**, 101710 (2025)

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

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

Froyland, G.: Dynamic isoperimetry and the geometry of Lagrangian coherent structures. Nonlinearity **28**, 3587–3622 (2015)

Froyland, G., Junge, O.: Robust FEM-based extraction of finite-time coherent sets using scattered, sparse, and incomplete trajectories. SIAM J. Appl. Dyn. Syst. **17**(2), 1891–1924 (2018)

Froyland, G., Koltai, P.: Detecting the birth and death of finite-time coherent sets. Commun. Pure Appl. Math. **76**(12), 3642–3684 (2023)

Froyland, G., Kwok, E.: Partitions of networks that are robust to vertex permutation dynamics. Spec. Matrices **3**, 22–42 (2015)

Froyland, G., Kwok, E.: A dynamic Laplacian for identifying Lagrangian coherent structures on weighted Riemannian manifolds. J. Nonlinear Sci. **30**, 1889–1971 (2020)
Froyland, G., Rock, C.P., Sakellariou, K.: Sparse eigenbasis approximation: multiple feature extraction across spatiotemporal scales with application to coherent set identification. Commun. Nonlinear Sci. Numer. Simul. **77**, 81–107 (2019)
Gómez, S., Díaz-Guilera, A., Gómez-Gardeñes, J., Pérez-Vicente, C.J., Moreno, Y., Arenas, A.: Diffusion dynamics on multiplex networks. Phys. Rev. Lett. **110**(2), (2013)
Hadjighasem, A., Karrasch, D., Teramoto, H., Haller, G.: Spectral-clustering approach to Lagrangian vortex detection. Phys. Rev. E **93**(6), 063107 (2016)
Haller, G., Beron-Vera, F.J.: Coherent Lagrangian vortices: the black holes of turbulence. J. Fluid Mech. **731**, R4 (2013)
Karrasch, D., Schilling, N.: A Lagrangian perspective on nonautonomous advection-diffusion processes in the low-diffusivity limit (2021). arXiv:2102.04777
Karrasch, D., Keller, J.: A geometric heat-flow theory of Lagrangian coherent structures. J. Nonlinear Sci. **30**(4), 1849–1888 (2020)
Lablée, O.: Spectral Theory in Riemannian Geometry. EMS Textbooks in Mathematics. European Mathematical Society (2015)
Mathew, G., Mezić, I., Petzold, L.: A multiscale measure for mixing. Phys. D **211**(1–2), 23–46 (2005)
Miclo, L.: On hyperboundedness and spectrum of Markov operators. Invent. Math. **200**(1), 311–343 (2015)
Mosovsky, B.A., Meiss, J.D.: Transport in transitory dynamical systems. SIAM J. Appl. Dyn. Syst. **10**(1), 35–65 (2011)
Ottino, J.M.: The Kinematics of Mixing: Stretching, Chaos, and Transport, vol. 3. Cambridge University Press (1989)
Padberg-Gehle, K., Schneide, C.: Network-based study of Lagrangian transport and mixing. Nonlinear Process. Geophys. **24**(4), 661–671 (2017)
Pierrehumbert, R.T.: Tracer microstructure in the large-eddy dominated regime. Chaos Solitons Fractals **4**(6), 1091–1110 (1994)
Schneide, C., Vieweg, P.P., Schumacher, J., Padberg-Gehle, K.: Evolutionary clustering of Lagrangian trajectories in turbulent Rayleigh–Bénard convection flows. Chaos **32**(1), (2022)
Sole-Ribalta, A., De Domenico, M., Kouvaris, N.E., Diaz-Guilera, A., Gomez, S., Arenas, A.: Spectral properties of the Laplacian of multiplex networks. Phys. Rev. E **88**(3), 032807 (2013)
Thiffeault, J.L.: Using multiscale norms to quantify mixing and transport. Nonlinearity **25**(2), R1 (2012)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.