



Čech–Delaunay gradient flow and homology inference for self-maps

U. Bauer¹ · H. Edelsbrunner² · G. Jabłoński² · M. Mrozek³

Received: 20 December 2019 / Accepted: 14 July 2020 / Published online: 30 August 2020
© The Author(s) 2020

Abstract

We call a continuous self-map that reveals itself through a discrete set of point-value pairs a *sampled dynamical system*. Capturing the available information with chain maps on Delaunay complexes, we use persistent homology to quantify the evidence of recurrent behavior. We establish a sampling theorem to recover the eigenspaces of the endomorphism on homology induced by the self-map. Using a combinatorial gradient flow arising from the discrete Morse theory for Čech and Delaunay complexes, we construct a chain map to transform the problem from the natural but expensive Čech complexes to the computationally efficient Delaunay triangulations. The fast chain map algorithm has applications beyond dynamical systems.

Keywords Computational topology · Persistent homology · Dynamical systems

Mathematics Subject Classification 37B99 · 55N31 · 65P99

This research has been supported by the DFG Collaborative Research Center SFB/TRR 109 “Discretization in Geometry and Dynamics”, by Polish MNiSzW Grant No. 2621/7.PR/12/2013/2, by the Polish National Science Center under Maestro Grant No. 2014/14/A/ST1/00453 and Grant No. DEC-2013/09/N/ST6/02995.

✉ U. Bauer
mail@ulrich-bauer.org
http://ulrich-bauer.org

H. Edelsbrunner
herbert.edelsbrunner@ist.ac.at
http://pub.ist.ac.at/~edels/

M. Mrozek
https://ww2.ii.uj.edu.pl/~mrozek/

¹ Department of Mathematics, Technical University of Munich, Boltzmannstraße 3, 85747 Garching b. München, Germany

² IST Austria (Institute of Science and Technology Austria), Am Campus 1, 3400 Klosterneuburg, Austria

³ Faculty of Mathematics and Computer Science, Jagiellonian University, ul. St. Łojasiewicza 6, PL-30-348 Kraków, Poland

1 Introduction

Suppose \mathbb{M} is a compact subset of \mathbb{R}^n and $f : \mathbb{M} \rightarrow \mathbb{M}$ is a continuous self-map with finite Lipschitz constant. We study the thus defined dynamical system in the setting in which f reveals itself through a *sample*, by which we mean a finite set $X \subseteq \mathbb{M}$, a self-map $g : X \rightarrow X$, and a real number ρ such that $\|g(x) - f(x)\| \leq \rho$ for every $x \in X$. We call ρ the *approximation constant* of the sample. Calling this setting a *sampled dynamical system*, we formalize a concept that appears already in Edelsbrunner et al. (2015). It is less demanding than the classical *discrete dynamical system*, in which time is discrete but space is not (Kaczynski et al. 2004). We believe that this relaxation is essential to make inroads into experimental studies, in which pairs $(x, f(x))$ can be observed individually, while the self-map remains in the dark. The approximation constant models the experimental uncertainty, but it is also needed to accommodate a finite sample. Consider for example the map $f : [0, 1] \rightarrow [0, 1]$ defined by $f(x) = \frac{x}{2}$. Letting u be the smallest positive value in a finite set $X \subseteq [0, 1]$, its image does not belong to X : $f(u) \notin X$. We call

$$\lambda = \max_{x, y \in X, x \neq y} \frac{\|g(x) - g(y)\|}{\|x - y\|} \quad (1)$$

the *Lipschitz constant* of g . It is not necessarily close to the Lipschitz constant of f , even in the case in which the ρ -neighborhoods of the points in X cover \mathbb{M} . However, Kirszbraun proved that for every $g : X \rightarrow X$ there is a continuous extension $f_0 : \mathbb{M} \rightarrow \mathbb{M}$ that has the same Lipschitz constant. Specifically, this is a consequence of the more general Kirszbraun Extension Property (Kirszbraun 1934; Wells and Williams 1975). Let \mathbb{F} be a fixed field and let $H(\mathbb{M}; \mathbb{F})$ denote the homology of \mathbb{M} with coefficients in \mathbb{F} . Hence, $H(\mathbb{M}; \mathbb{F})$ is a vector space. Throughout the paper we only use homology with coefficients in the field \mathbb{F} , so we abbreviate the notation to $H(\mathbb{M})$. The map f_0 induces a linear map $H(f_0) : H(\mathbb{M}) \rightarrow H(\mathbb{M})$. A natural characterization of this linear map are the *t-eigenvectors*. They capture homology classes invariant under the self-map up to a multiplicative factor t , called an *eigenvalue*. The *t-eigenvectors* span the *t-eigenspace* of the map. Starting with a finite filtration of the domain of the map, we get *t-eigenspaces* at every step, connected by linear maps, and therefore a finite path in the category of vector spaces, called an *eigenspace module*. The Stability Theorem in Edelsbrunner et al. (2015) implies a connection between the dynamics of g and f_0 , namely that for every eigenvalue t the interleaving distance between the eigenspace modules induced by g and by f_0 is at most the Hausdorff distance between the graph of g and that of f_0 . Furthermore, the Inference Theorem in the same paper implies that for small enough ρ and any eigenvalue, the eigenspace module for g gives the correct dimension of the corresponding eigenspace of the endomorphism between the homology groups of \mathbb{M} induced by f_0 .

1.1 Prior work and results

We employ the discrete Morse theory for Čech and Delaunay complexes developed in Bauer and Edelsbrunner (2017) to address the computational problem of estimating the homology of a self-map from a finite sample. Our results continue the program started in Edelsbrunner et al. (2015), with the declared goal to embed the concept of persistent homology in the computational approach to dynamical systems. Specifically, we contribute by improving the computation of persistent recurrent dynamics. This improvement is based on several interacting innovations, which lead to better theoretical guarantees as well as better computational efficiency than in Edelsbrunner et al. (2015):

1. We use the parallel filtrations of Čech and Delaunay complexes and the existence of a collapse from the former to the latter established in Bauer and Edelsbrunner (2017) to define chain maps between Delaunay complexes.
2. We construct the chain maps by implementing the collapse implicitly, avoiding the prohibitive construction of the Čech complex.
3. We establish inference results with a less stringent sampling condition than given in Edelsbrunner et al. (2015), depending only on the self-map and the domain.

The improved computational efficiency derives primarily from the use of Delaunay rather than Čech or Vietoris–Rips complexes. Indeed, in the targeted 2-dimensional case, the size of the Delaunay triangulation is at most six times the number of data points, while the Čech and Vietoris–Rips complexes reach exponential size for large radii. The improved theoretical guarantees rely on the use of chain maps that avoid the information loss caused by the interaction of local expansion and partial maps observed in Edelsbrunner et al. (2015). The improvements are obtained using refined mathematical and computational methods as mentioned above.

We first explain how we use Čech complexes, namely as an intermediate step to construct the chain maps from one Delaunay complex to another. Recall the Kirszbraun intersection property for balls established by Gromov (1987): letting Q be a finite set of points in \mathbb{R}^n , and $g : Q \rightarrow \mathbb{R}^n$ a map that satisfies $\|g(x) - g(y)\| \leq \|x - y\|$ for all $x, y \in Q$, then

$$\bigcap_{x \in Q} B_r(x) \neq \emptyset \implies \bigcap_{x \in Q} B_r(g(x)) \neq \emptyset, \tag{2}$$

in which $B_r(x)$ is the closed ball with radius r and center x . Similarly, if we weaken the condition to $\|g(x) - g(y)\| \leq \lambda \|x - y\|$, for some $\lambda > 1$, then the common intersection of the balls $B_{\lambda r}(g(x))$ is non-empty. This implies that the image of the Delaunay complex for radius r includes in the Čech complex for radius λr . To return to the Delaunay triangulation, we exploit the collapsibility of the Čech complex for radius λr to the Delaunay complex of radius λr recently established in Bauer and Edelsbrunner (2017). We second explain how we collapse without explicit construction of the Čech complex. Starting with a simplex, we use a modification of Welzl’s miniball algorithm (Welzl 1991) to follow the flow induced by the collapse step by step until we arrive at the Delaunay complex, where the image of the simplex is now a chain. The




Fig. 1 In each column, we get the eigenspace by comparing the inclusion between Delaunay–Čech complexes with the chain map obtained with the Čech complex as intermediary. The map $D\check{C}ech_r(X) \rightarrow \check{C}ech_{\lambda,r}(X)$ is induced by g , while the map $\check{C}ech_{\lambda,r}(X) \rightarrow D\check{C}ech_{\lambda,r}(X)$ is a simplicial collapse, and similarly for s instead of r

expected running time for a single step is linear in the number of points, so we have a fast algorithm provided the number of steps in the collapse is not large. While we do not have a bound on this number, our computational experiments provide evidence that it is typically small.

We give a global picture of our algorithm in Fig. 1. In the top row, we see a filtration of Delaunay–Čech complexes, which are convenient substitutes for the better known Delaunay complexes (also called alpha complexes) with the same homotopy type. The left map down from the top row is inclusion, and the right map down is the chain map induced by g . As indicated, the right map is composed of the inclusion into the Čech complex and the discrete flow induced by the collapse. In the bottom row, we see the eigenspace module computed by comparing the left and right vertical maps.

1.2 Outline

Section 2 describes the background in discrete Morse theory, its application to Čech and Delaunay complexes, and its extension to persistent homology. Section 3 addresses the algorithmic aspects of our method, which include the proof of collapsibility and the generalization of the miniball algorithm. Section 4 explains the circumstances under which the eigenspace of the self-map can be obtained from the eigenspace module of the discrete sample. Section 5 presents the results of our computational experiments, comparing them with the algorithm in Edelsbrunner et al. (2015). Section 6 concludes this paper.

2 Background

In this section, we introduce concepts from discrete Morse Theory (Forman 1998) and apply them to Čech as well as to Delaunay complexes of finite point sets (Bauer and Edelsbrunner 2017). We begin with the definition of the complexes and finish by complementing the picture with the theory of persistent homology.




Fig. 2 An example illustrating the difference between Delaunay and Delaunay–Čech complexes. The Delaunay complex for the given radius has three vertices and two edges. In contrast, the Delaunay–Čech complex is the full simplex on the three vertices (not shown), as all simplices are Delaunay and enclosed by a sphere of radius r

2.1 Geometric complexes

Our approach to dynamical systems is based on Čech complexes and Delaunay complexes—two common ingredients in topological data analysis—and the Delaunay–Čech complexes, which offer a convenient computational short-cut.

Čech complexes

Let $X \subseteq \mathbb{R}^n$ be finite, $r \geq 0$, and $B_r(x)$ be the closed ball of points at distance r or less from $x \in X$. The Čech complex of X for radius r consists of all subsets of X for which the balls of radius r have a non-empty common intersection:

$$\check{C}ech_r(X) = \{Q \subseteq X \mid \bigcap_{x \in Q} B_r(x) \neq \emptyset\}; \tag{3}$$

it is isomorphic to the nerve of the balls of radius r centered at the points in X . Equivalently, $\check{C}ech_r(X)$ consist of all subsets $Q \subseteq X$ having an enclosing sphere of radius at most r . For r smaller than half the distance between the two closest points, $\check{C}ech_r(X) = X$, and for r larger than $\sqrt{2}/2$ times the distance between the two farthest points, $\check{C}ech_r(X)$ is the full simplex on the vertices X , denoted by $\Delta(X)$. The size of $\Delta(X)$ is exponential in the size of X , which motivates the following construction.

Delaunay triangulations

The Voronoi domain of a point $x \in X$ consists of all points $u \in \mathbb{R}^n$ for which x minimizes the distance from u : $dom(x, X) = \{u \in \mathbb{R}^n \mid \|x - u\| \leq \|y - u\|, \text{ for all } y \in X\}$. The Voronoi tessellation of X is the set of Voronoi domains $dom(x, X)$ with $x \in X$. Assuming general position of the points in X , any $p + 1$ Voronoi domains are either disjoint or they intersect in a common $(n - p)$ -dimensional face. The Delaunay triangulation of X consists of all subsets of X for which the Voronoi domains have a

non-empty common intersection:

$$\text{Del}(X) = \{Q \subseteq X \mid \bigcap_{x \in Q} \text{dom}(x, X) \neq \emptyset\}; \quad (4)$$

it is isomorphic to the nerve of the Voronoi tessellation. Equivalently, $\text{Del}_r(X)$ consist of all subsets $Q \subseteq X$ having an empty circumsphere (containing no points of X in its interior). Again assuming general position, the Delaunay triangulation is an n -dimensional simplicial complex with natural geometric realization in \mathbb{R}^n . The Upper Bound Theorem for convex polytopes implies that the number of simplices in $\text{Del}(X)$ is at most some constant times $\text{card } X$ to the power $\lceil n/2 \rceil$. In $n = 2$ dimensions, this is linear in $\text{card } X$, which compares favorably to the exponentially many simplices in the Čech complexes.

Delaunay–Čech complexes

To combine the small size of the Delaunay triangulation with the scale-dependence of the Čech complex, we define the *Delaunay–Čech complex* of X for radius r as the intersection of the two:

$$\text{D}\check{\text{C}}\text{ech}_r(X) = \check{\text{C}}\text{ech}_r(X) \cap \text{Del}(X). \quad (5)$$

Observe that the Delaunay triangulation effectively curbs the explosive growth of simplex numbers, but does so only if the points are in general position. We will therefore assume that the points in X are in general position, justifying the assumption with computational simulation that enforce this assumption (Edelsbrunner and Mücke 1990).

Delaunay complexes

There is a more direct way to select subcomplexes of the Delaunay triangulation using r as a parameter. Specifically, the *Delaunay complex* of X for radius r consists of all subsets of X for which the restriction of the Voronoi domains to the balls of radius r have a non-empty common intersection:

$$\text{Del}_r(X) = \{Q \subseteq X \mid \bigcap_{x \in Q} [\text{dom}(x, X) \cap B_r(x)] \neq \emptyset\}; \quad (6)$$

it is isomorphic to the nerve of the restricted Voronoi domains. Equivalently, $\text{Del}_r(X)$ consist of all subsets $Q \subseteq X$ having an empty circumsphere of radius at most r . The Delaunay complexes, also known as *alpha complexes*, are the better known relatives of the Delaunay–Čech complexes. We use the They satisfy $\text{Del}_r(X) \subseteq \text{D}\check{\text{C}}\text{ech}_r(X)$, and it is easy to exhibit sets X and radii r for which the two complexes are different. See Fig. 2 for an illustrating example. As proved in Bauer and Edelsbrunner (2017), the Delaunay complex has the same homotopy type as the Delaunay–Čech complex

for the same radius. This is indeed the reason we can freely use the latter as a substitute of the former.

2.2 Radius functions

Structural properties of the geometric complexes are conveniently expressed in terms of their radius functions. In each case, the function maps a simplex to the smallest radius, r , for which the simplex belongs to the complex:

$$\mathcal{R}_C(Q) = \min\{r \mid Q \in \check{C}ech_r(X)\}, \tag{7}$$

$$\mathcal{R}_{DC}(Q) = \min\{r \mid Q \in D\check{C}ech_r(X)\}, \tag{8}$$

$$\mathcal{R}_D(Q) = \min\{r \mid Q \in Del_r(X)\}. \tag{9}$$

All three functions are *monotonic*, by which we mean that the radius assigned to any simplex is greater than or equal to the radii assigned to its faces. This property is sufficient to define their persistence diagrams, as we will see shortly. However, we will need more, namely compatible discrete gradients of the radius functions. After introducing the discrete Morse theory of Forman (1998) as a framework within which discrete gradients can be defined, we will return to the question of compatibility.

Discrete Morse theory

In a nutshell, a monotonic function on a simplicial complex, $F: K \rightarrow \mathbb{R}$, is a discrete Morse function if any two contiguous sublevel sets differ by a single elementary collapse or a critical simplex. We are now more precise. A *pair* consists of two simplices, $P \subseteq Q$, with dimensions $\dim Q = 1 + \dim P$. A *discrete vector field* is a partition, V , of K into pairs and singletons. It is *acyclic* if there is a monotonic function, $F: K \rightarrow \mathbb{R}$, with $F(P) = F(Q)$ iff P and Q belong to a pair in V . Such a function F is called a *discrete Morse function*, and V is its *discrete gradient*. A simplex is *critical* if it is in a singleton of V , and it is *non-critical* if it belongs to a pair of V .

The reason for our interest in this formalism is its connection to the homotopy type of complexes. To explain suppose $Q \in K$ maximizes F . If Q belongs to a pair $(P, R) \in V$, then we can remove both and obtain a smaller simplicial complex, $K \setminus \{P, R\}$. We refer to this operation as an *elementary collapse*, we say K *collapses* to the smaller complex, denoted $K \searrow K \setminus \{P, R\}$, and we note that both complexes have the same homotopy type. If on the other hand Q is a critical simplex, its removal changes the homotopy type of the complex.

Collapsing the geometric complexes

The radius functions are not necessarily discrete Morse functions, but they are amenable to discrete gradients. To explain what we mean, consider a monotonic function, $F: K \rightarrow \mathbb{R}$, and call $Q \in K$ *critical* if $F(Q)$ is different from the values of all proper faces and cofaces of Q . We say that an acyclic partition of K into pairs and singletons is *compatible* with F if every sublevel set of F is a union of pairs and

singletons in this partition, and Q is in a singleton of the partition iff Q is a critical simplex of F . The proof of collapsibility in Bauer and Edelsbrunner (2017) hinges on the fact that there is an acyclic partition, V , of $\Delta(X)$ that is simultaneously compatible with \mathcal{R}_C , \mathcal{R}_D , and \mathcal{R}_{DC} . Indeed, the existence of this acyclic partition is at the core of the proof of Theorem 5.10 in Bauer and Edelsbrunner (2017), which asserts that

$$\check{C}ech_r(X) \searrow D\check{C}ech_r(X) \searrow Del_r(X) \tag{10}$$

for every finite set $X \subseteq \mathbb{R}^n$ in general position, and every $r \geq 0$. Observe that this implies that the three radius functions have the same set of critical simplices. Indeed, these are the sets $Q \subseteq X$ for which the smallest enclosing sphere passes through all points of Q and no point of X lies inside this sphere.

2.3 Persistent homology

In its original conception, persistent homology starts with a filtration of a topological space, it applies the homology functor for coefficients in a field \mathbb{F} , and it decomposes the resulting sequence of vector spaces into indecomposable summands (Edelsbrunner et al. 2002; Zomorodian and Carlsson 2005). This decomposition is unique and has an intuitive interpretation in terms of births and deaths of homology classes. We flesh out the idea using the filtration of Delaunay–Čech complexes as an example.

Let $X \subseteq \mathbb{R}^n$ be finite and in general position, and recall that $\mathcal{R}_{DC}: Del(X) \rightarrow \mathbb{R}$ is the radius function whose sublevel sets are the Delaunay–Čech complexes. \mathcal{R}_{DC} is monotonic but not necessarily discrete Morse. The Delaunay triangulation is finite, which implies that \mathcal{R}_{DC} has only finitely many sublevel sets. To index them consecutively, we write $r_1 < r_2 < \dots < r_N$ for the values and $K_i = \mathcal{R}_{DC}^{-1}[0, r_i]$ for the i -th Delaunay–Čech complex of X . Applying the homology functor, we get

$$0 = H(K_1) \rightarrow H(K_2) \rightarrow \dots \rightarrow H(K_N), \tag{11}$$

in which we write $H(K_i)$ for the direct sum of the homology groups of all dimensions. Together with the maps $h_{i,j}: H(K_i) \rightarrow H(K_j)$ induced by the inclusions $K_i \subseteq K_j$, which are linear, we call this diagram the *persistent homology* of the filtration. More generally, a diagram of vector spaces with this shape is called a *persistence module*. Such a module is *indecomposable* if all vector spaces are trivial, except for an interval of 1-dimensional vector spaces, $\mathbb{F} \rightarrow \mathbb{F} \rightarrow \dots \rightarrow \mathbb{F}$, that are connected by isomorphisms. Indeed, (11), and more generally, any persistence module of finite-dimensional vector spaces, can be written as the direct sum of indecomposable modules, and this decomposition is essentially unique. See Edelsbrunner et al. (2015, Basis Lemma) for a constructive proof. If an interval starts at position i and ends at position $j - 1$, then we say there is a homology class *born* at K_i that *dies entering* K_j . To allow for the case $j - 1 = N$, we introduce $r_{N+1} = \infty$ and represent the interval by the *birth-death pair* (r_i, r_j) . Its *dimension* is the homological degree in which the class arises, and its *persistence* is $r_j - r_i$.

By construction, the rank of $H(K_i)$ is the number of indecomposable modules whose intervals cover r_i . It is readily computed from the multiset of birth-death pairs, which we call the *persistence diagram* of the radius function, denoted $\text{Dgm}(\mathcal{R}_{\text{DC}})$. More generally, we can use this diagram to compute the rank of the image of $h_{i,j}$ for $i \leq j$; see e.g. Edelsbrunner and Harer (2010, p. 152).

3 Computing the Čech–Delaunay gradient flow

The main algorithmic challenge we face in this paper is the local computation of the gradient that induces the collapse of the Čech to the Delaunay–Čech complex. Specifically, we trace chains through the collapse, using their images to construct the chain map that is central to our analysis. We explain the algorithm in three stages: first sketching the relevant steps of the existence proof, second describing how we compute minimum separating spheres, and third explaining the discrete flow that constructs the chain map. Once we arrive at the eigenspaces, we compute their persistent homology with the software implementing the algorithms in Edelsbrunner et al. (2015).

3.1 Computing separating spheres

At the core of the discrete gradient flow is the construction of smallest separating spheres, which are defined as follows. Let $X \subseteq \mathbb{R}^n$ be a finite set of points in general position, and let $A \subseteq X$ be a subset. An $(n - 1)$ -dimensional sphere *separates* another subset $Q \subseteq X$ from A if

- all points of Q lie inside or on the sphere, and
- all points of A lie outside or on the sphere.

If a point belongs to both A and Q , then it must lie on the separating sphere. Given Q and A , a separating sphere may or may not exist, and if it exists, then there is a unique *smallest separating sphere*, which we denote $S(Q, A)$.

The smallest separating sphere can be characterized in geometric terms as follows. For a sphere S , write $\text{Incl } S, \text{Excl } S \subseteq X$ for the subsets of enclosed and excluded points, with $\text{On } S = \text{Incl } S \cap \text{Excl } S$. Now assume that S is the smallest circumsphere of the points $\text{On } S$, i.e., the center z of S lies in their affine hull:

$$z = \sum_{x \in \text{On } S} \rho_x x \quad \text{with} \quad 1 = \sum_{x \in \text{On } S} \rho_x.$$

By general position, the affine combination is unique, and $\rho_x \neq 0$ for all $x \in \text{On } S$. We call

$$\begin{aligned} \text{Front } S &= \{x \in \text{On } S \mid \rho_x > 0\}, \\ \text{Back } S &= \{x \in \text{On } S \mid \rho_x < 0\} \end{aligned}$$

the *front face* and the *back face* of $\text{On } S$, respectively. The following lemma states necessary and sufficient conditions for a sphere to be a smallest separating sphere. It is

a special case of the general Karush–Kuhn–Tucker conditions, expressed in geometric and combinatorial terms.

Lemma 1 (Combinatorial KKT Conditions Bauer and Edelsbrunner 2017) *Let X be a finite set of points in general spherical position, and let $Q, A \subseteq X$. A sphere S satisfies $S = S(Q, A)$ iff*

- (i) S is the smallest circumsphere of the points $\text{On } S$,
- (ii) $\text{Front } S \subseteq Q \subseteq \text{Incl } S$, and
- (iii) $\text{Back } S \subseteq A \subseteq \text{Excl } S$.

Based on these optimality conditions, we can state a recursive formula for the smallest separating sphere.

Lemma 2 *Assume that $S(Q, A)$ exists. If $x \in Q$, then*

$$S(Q, A) = \begin{cases} S(Q \setminus \{x\}, A) & \text{if that sphere encloses } x, \\ S(Q, A \cup \{x\}) & \text{otherwise.} \end{cases}$$

Similarly, if $x \in A$, then

$$S(Q, A) = \begin{cases} S(Q, A \setminus \{x\}) & \text{if that sphere excludes } x, \\ S(Q \cup \{x\}, A) & \text{otherwise.} \end{cases}$$

Proof We only show the first part, with $x \in Q$, the other part being analogous.

First, assume that $S := S(Q \setminus \{x\}, A)$ encloses x . Then we have $Q \subseteq \text{Incl } S$, and thus $S(Q, A) = S$ by Lemma 1.

On the other hand, if $S(Q \setminus \{x\}, A)$ does not enclose x , then we must have $S := S(Q, A) \neq S(Q \setminus \{x\}, A)$, and thus Lemma 1 gives $\text{Front } S \not\subseteq Q \setminus \{x\}$. But Lemma 1 also gives $\text{Front } S \subseteq Q$, and so we must have $x \in \text{Front } S$. Since $\text{Front } S \subseteq \text{On } S \subseteq \text{Excl } S$, it follows that $A \cup \{x\} \subseteq \text{Excl } S$, and thus $S(Q, A \cup \{x\}) = S$ by Lemma 1. \square

We now turn these results into an algorithm for computing the smallest separating sphere of sets $Q, A \subseteq X$, or deciding that no separating sphere exists. We pattern the algorithm after the randomized algorithm for the smallest enclosing sphere described in Welzl (1991), which we recall first.

Welzl’s randomized miniball algorithm

The smallest enclosing sphere of a set $Q \subseteq \mathbb{R}^n$ is determined by at most $n + 1$ of the points. In other words, there is a subset $R \subseteq Q$ of at most $n + 1$ points such that the smallest enclosing sphere of R is also the smallest enclosing sphere of Q . The algorithm below makes essential use of this observation. It partitions Q into two disjoint subsets: R containing the points we know lie on the smallest enclosing sphere, and $P = Q \setminus R$. Initially, $R = \emptyset$ and $P = Q$. In a general step, the algorithm removes a random point from P and tests whether it lies on or inside the recursively computed smallest enclosing sphere of the remaining points. If yes, the point is discarded, and if no, the point is added to R .

```

1  sphere ENCLOSE( $P$ ,  $R$ ):
2    if  $P = \emptyset$  then let  $S$  be the smallest circumsphere of  $R$ 
3      else choose a random point  $x \in P$ ;
4           $S = \text{ENCLOSE}(P \setminus \{x\}, R)$ ;
5          if  $x$  outside  $S$  then  $S = \text{ENCLOSE}(P \setminus \{x\}, R \cup \{x\})$ ;
6  return  $S$ .

```

Since the algorithm makes random choices, its running time is a random variable. Remarkably, the expected running time is linear in the number of points in Q , and the reason is the high probability that the randomly chosen point, x , lies inside the recursively computed smallest enclosing sphere and can therefore be discarded.

Generalization to smallest separating spheres

Rather than enclosing spheres, we need separating spheres to compute the collapse. Here we get an additional case, when the sphere does not exist, which we indicate by returning NULL. As before, we work with two sets of points: R containing the points we know lie on the smallest separating sphere, and P containing the rest. Initially, $R = Q \cap A$ and $P = (Q \cup A) \setminus R$. Each point has enough memory to remember whether it belongs to Q and thus needs to lie on or inside the sphere, or to A and thus needs to lie on our outside the sphere. We say the point *contradicts* S if it lies on the wrong side.

```

1  sphere SEPARATE( $P$ ,  $R$ ):
2    if  $\text{card } R > n + 1$  then return NULL;
3    if  $P = \emptyset$  then let  $S$  be the smallest circumsphere of  $R$ 
4      else choose a random point  $x \in P$ ;
5           $S = \text{SEPARATE}(P \setminus \{x\}, R)$ ;
6          if  $x$  contradicts  $S$  then  $S = \text{SEPARATE}(P \setminus \{x\}, R \cup \{x\})$ ;
7  return  $S$ .

```

Since the smallest separating sphere is again determined by at most $n + 1$ of the points, the expected running time of the algorithm is linear in the number of points, as before. The correctness of the algorithm is warranted by Lemma 2.

Iterative version with move-to-front heuristic

Because finding separating spheres is at the core of our algorithm, we are motivated to improve its running time, even if it is only by a constant factor. Following the advise in Gärtner (1999), we turn the tail-recursion into an iteration and combine this with a move-to-front heuristic. Indeed, if a point contradicts the current sphere, it is likely that it does the same to a later computed sphere. The earlier the point is tested, the faster this new sphere can be rejected. Storing the points in a linear list, early testing of this point can be enforced by moving it to the front of the list. Write \mathcal{L} for the list, which contains all points of $Q \cup A$, and write $\mathcal{L}(i)$ for the point stored at the i -th location. As before, each point remembers whether it belongs to Q , to A , or to both. In addition,

we mark the points we know lie on the smallest separating sphere as members of R , initializing this set to $R = Q \cap A$. Furthermore, we initialize $m = \text{card}(Q \cup A)$.

```

1  sphere MOVETOFRONT( $\mathcal{L}, m, R$ ):
2    if  $\text{card } R > n + 1$  then return NULL;
3    let  $S$  be smallest circumsphere of  $R$ ;
4    for  $i = 1$  to  $m$  do
5      if  $x = \mathcal{L}(i)$  contradicts  $S$  then  $S = \text{MOVETOFRONT}(\mathcal{L}, i - 1, R \cup \{x\})$ ;
6      if  $S = \text{NULL}$  then return NULL;
7      move  $x$  to front of  $\mathcal{L}$ ;
8    return  $S$ .
```

Section 5 will present experimental evidence that the move-to-front heuristic accelerates the computations.

3.2 Collapsing non-Delaunay simplices

Recall that the collapsing sequence in (10) is facilitated by a discrete gradient, W , that is compatible with all three radius functions. To collapse a Čech complex to the Delaunay–Čech complex, we only need the pairs in W that partition the difference: $\check{\text{Cech}}_r(X) \setminus \text{D}\check{\text{Cech}}_r(X) \subseteq \Delta(X) \setminus \text{Del}(X)$. This difference is indeed partitioned solely by pairs because all singletons contain critical simplices, which belong to $\text{Del}(X)$. The discrete gradient on the full simplex $\Delta(X)$ determined by those non-Delaunay pairs will be denoted by V .

Following Bauer and Edelsbrunner (2017, Lemma 5.8), we note that every pair of the discrete gradient V is of the form (P, R) with $P \subseteq R \subseteq X$ and $R \setminus P = \{x\}$ for a unique vertex $x \in R$. In other words, $(P, R) \in V$ uniquely determines the vertex in which the two simplices differ, and given $Q \in \{P, R\}$ together with this vertex, we can recover the pair as $(P, R) = (Q \setminus \{x\}, Q \cup \{x\})$. We therefore introduce the map $\psi: \Delta(X) \setminus \text{Del}(X) \rightarrow X$ defined by mapping the non-Delaunay simplex Q to the corresponding vertex, $\psi(Q) = x$, and we use this map to represent the discrete gradient V .

We now describe the construction of the map ψ from Bauer and Edelsbrunner (2017) that defines the discrete gradient V , whose pairs partition the non-Delaunay simplices. To this end, we choose an arbitrary but fixed total ordering x_1, x_2, \dots, x_N of the points in X . For each $0 \leq j \leq N$, we write $X_j = \{x_i \mid i \leq j\}$ for the prefix. Given a non-Delaunay simplex $Q \in \Delta(X) \setminus \text{Del}(X)$, let $E_Q \subseteq X$ be the subset of points that lie on or outside of the smallest enclosing sphere of Q , and for each $0 \leq j \leq N$, define $A_j = E_Q \cup X_j$. The sequence A_0, A_1, \dots, A_N starts with just the exterior points, $A_0 = E_Q$, and ends with all points, $A_N = X$. Since $Q \notin \text{Del}(X)$, there is a minimal index $j \leq N$ such that Q and A_j do not permit a separating sphere. We use the corresponding vertex x_j to define $\psi(Q) = x_j$. To compute $\psi(Q)$, it thus suffices to iterate through the sequence A_0, A_1, \dots, A_N and find the first index j such that there is no sphere separating Q from A_j . This can be determined using the algorithm described in Sect. 3.1.

3.3 Constructing the chain map

We now have the necessary prerequisites for constructing the chain map. Specifically, given a cycle in $D\check{C}ech_r(X)$, we are interested in computing its image, which is a cycle in $D\check{C}ech_s(X)$, with $r \leq s \leq \rho + \lambda r$. The construction of the chain map is an application of the discrete Morse theoretic formalism of a discrete gradient flow and the corresponding stabilization map, which we now review.

We follow the notation in Forman (1998), in which the discrete gradient flow is formulated as a map on chains. Let K be a simplicial complex and V a discrete gradient on K . In our situation, $K = \check{C}ech_r(X)$, and V contains the pairs defined by the map ψ introduced in Sect. 3.2, which partition $\check{C}ech_r(X) \setminus D\check{C}ech_r(X)$. It is convenient to consider the discrete gradient as a chain map. Fixing an orientation on each simplex, this chain map is defined by linear extension of the map on the oriented simplices given by

$$V(P) = \begin{cases} \pm R & \text{if } (P, R) \in V, \\ 0 & \text{otherwise,} \end{cases} \tag{12}$$

where the sign is chosen so that P appears with coefficient -1 in the boundary of R . In terms of the map ψ defining the gradient V as discussed in Sect. 3.2, this definition can be rewritten as

$$V(P) = \begin{cases} \pm(P \cup \{\psi(P)\}) & \text{if } \psi(P) \notin P, \\ 0 & \text{otherwise,} \end{cases} \tag{13}$$

This map sends every oriented p -simplex to 0 or to an oriented $(p + 1)$ -simplex. The linear extension yields a homomorphism $V : C(K) \rightarrow C(K)$, which maps every p -chain to a possibly trivial $(p + 1)$ -chain. Recall that the boundary map, $\partial : C(K) \rightarrow C(K)$, sends every p -chain to a possibly trivial $(p - 1)$ -chain. We use both to introduce $\Phi : C(K) \rightarrow C(K)$ defined by

$$\Phi(c) = c + \partial(V(c)) + V(\partial(c)), \tag{14}$$

in which c is a p -chain and its image, $\Phi(c)$, is a possibly trivial p -chain. We call Φ the *discrete gradient flow* induced by V . Importantly, it commutes with the boundary map: $\partial\Phi = \Phi\partial$, which makes it a chain map; see Forman (Forman 1998, Theorem 6.4). Moreover, the iteration of Φ stabilizes in the sense that $\Phi^M = \Phi^N$ for large enough M and N (Forman 1998, Theorem 7.2). We call this chain map the *stabilization map* of Φ and denote it by Φ^∞ .

In this paper, we apply the discrete flow exclusively to cycles. In other words, $c \in C(K)$ satisfies $\partial c = 0$, which simplifies the above formula (14) to

$$\Phi(c) = c + \partial(V(c)). \tag{15}$$

In order to evaluate the stabilization map Φ^∞ , we simply iterate Φ until it stabilizes. The most demanding step in each iteration is the computation of smallest separating spheres, as discussed in Sect. 3.1.

4 Eigenspace Inference

We use the chain maps connecting the Delaunay–Čech complexes to construct a persistence module of eigenspaces from the sample $g : X \rightarrow X$, and specify properties of the sampled dynamical system under which the eigenspaces of the underlying self-map can be inferred from this module. Because of this specific goal, we typically work with coefficients in a finite field of larger order, in contrast to the typical setup in applied topology, where homology is often taken with coefficients in the field \mathbb{Z}_2 .

4.1 Eigenspaces

Given a finite set $X \subseteq \mathbb{M} \subseteq \mathbb{R}^n$, we recall that $\mathcal{R}_{DC} : \text{Del}(X) \rightarrow \mathbb{R}$ is the radius function whose sublevel sets are the Delaunay–Čech complexes of X . Let $r_1 < r_2 < \dots < r_N$ be the values of \mathcal{R}_{DC} , and write $D\check{\text{Cech}}_r(X) = \mathcal{R}_{DC}^{-1}[0, r]$ for the Delaunay–Čech complex at radius r . We construct the persistence diagram of this filtration, denoted $\text{Dgm}(\mathcal{R}_{DC})$, which is a multi-set of intervals of the form $[r_i, r_j)$. For each such interval, there is a unique homology class born at $D\check{\text{Cech}}_{r_i}(X)$ that maps to 0 when it dies entering $D\check{\text{Cech}}_{r_j}(X)$, and the collection of such classes gives a basis for the homology group of every complex in the filtration.

To define the eigenspace, for each r we consider two maps between homology groups, $\iota_r, \kappa_r : H(D\check{\text{Cech}}_r(X)) \rightarrow H(D\check{\text{Cech}}_{r+q}(X))$, in which ι_r is induced by the inclusion $D\check{\text{Cech}}_r(X) \subseteq D\check{\text{Cech}}_{r+q}(X)$, κ_r is induced by the chain map composed of g followed by the stabilization map Φ^∞ , and $q \geq 0$ is chosen such that all generators of $H(D\check{\text{Cech}}_r(X))$ have images under the chain map κ in $H(D\check{\text{Cech}}_{r+q}(X))$. It is convenient to represent ι_r and κ_r by matrices that write the images of the generators of $H(D\check{\text{Cech}}_r(X))$ in terms of the generators of $H(D\check{\text{Cech}}_{r+q}(X))$. Following Edelsbrunner et al. (2015), we consider the *generalized eigenspace* of the two maps for an eigenvalue t :

$$E^t(\kappa_r, \iota_r) = \ker(\kappa_r - t \cdot \iota_r) / (\ker \kappa_r \cap \ker \iota_r). \tag{16}$$

In words, $E^t(\kappa_r, \iota_r)$ is generated by the cycles in $D\check{\text{Cech}}_r(X)$ whose images under κ_r are homologous to t times their images under ι_r . Note that this is a slight modification of the classic eigenvalue problem in which the image and the range are identical. This is not the case for κ_r , so we compare it to ι_r to get the eigenspace. The maps between the eigenspaces,

$$e^t_{r,s} = E^t(\kappa_r, \iota_r) \rightarrow E^t(\kappa_s, \iota_s) \tag{17}$$

are obtained as restrictions of the maps $h_{r,s} : H(D\check{\text{Cech}}_r(X)) \rightarrow H(D\check{\text{Cech}}_s(X))$ induced by inclusion. For fixed $t \in \mathbb{F}$, we have a sequence of eigenspaces,

$$0 \rightarrow E^t(\kappa_{r_1}, \iota_{r_1}) \rightarrow E^t(\kappa_{r_2}, \iota_{r_2}) \rightarrow \dots \rightarrow E^t(\kappa_{r_N}, \iota_{r_N}), \tag{18}$$

which together with the maps e_{r_i, r_j}^t form a persistence module. Recall from Sect. 2.3 that this persistence module has an essentially unique interval decomposition. We can therefore compute the persistence diagram, which we refer to as the *eigenspace diagram* of g for eigenvalue t , denoted $\text{Egm}(g, t)$.

4.2 Maps between nerves

We will relate the eigenspace of f for t with the eigenspace module in two steps. The second step will use results about nerves of covers, which we now review.

Let \mathbb{X} be a topological space and $\mathcal{U} = (U_i)_{i \in I}$ a cover of \mathbb{X} . \mathcal{U} is *closed* or *open* if every U_i is closed or open, respectively, and \mathcal{U} is *good* if the common intersection of any subset of cover elements is empty or contractible. Recall that the *nerve* of \mathcal{U} is the collection of subsets with non-empty common intersection:

$$N(\mathcal{U}) = \{B \subseteq \mathcal{U} \mid \bigcap B \neq \emptyset\}. \tag{19}$$

Calling B a *simplex*, the nerve is an abstract simplicial complex. A *partition of unity subordinate* to \mathcal{U} is a collection of continuous nonnegative functions $\phi_i : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ such that $\sum_{i \in I} \phi_i(x) = 1$ for every $x \in \mathbb{X}$, and the support of ϕ_i is contained in U_i for every $i \in I$. Assuming a geometric realization of the nerve in which v_i denotes the vertex that represents the subset $U_i \in \mathcal{U}$, we introduce the map

$$r : \mathbb{X} \rightarrow |N(\mathcal{U})| \text{ defined by } r(x) = \sum_{i \in I} \phi_i(x) \cdot v_i. \tag{20}$$

The Nerve Theorem as stated in Hatcher (2002) asserts that r is a homotopy equivalence provided \mathcal{U} is a good cover that has a subordinate partition of unity. Such a partition exists for example if \mathcal{U} is open and \mathbb{X} is paracompact, which includes $\mathbb{X} \subseteq \mathbb{R}^n$. We expand on the Nerve Theorem, using the map r from (20) to relate a continuous map with a corresponding simplicial map between nerves.

Lemma 3 *Let $\mathcal{U} = (U_i)_{i \in I}$ and $\mathcal{V} = (V_j)_{j \in J}$ be open covers of spaces \mathbb{X} and \mathbb{Y} with corresponding subordinate partitions of unity. Let $f : \mathbb{X} \rightarrow \mathbb{Y}$ be continuous, let $g : I \rightarrow J$ be such that $f(U_i) \subseteq V_{g(i)}$ for every $i \in I$, and write $h : |N(\mathcal{U})| \rightarrow |N(\mathcal{V})|$ for the linear simplicial map induced by g . Then the diagram*

$$\begin{array}{ccc} \mathbb{X} & \xrightarrow{f} & \mathbb{Y} \\ r \downarrow & & \downarrow s \\ |N(\mathcal{U})| & \xrightarrow{h} & |N(\mathcal{V})| \end{array} \tag{21}$$

commutes up to homotopy, in which r and s are constructed as in (20).

Proof Let $x \in \mathbb{X}$, and let $\tau(x) = \text{conv}\{w_j \in J \mid f(x) \in V_j\}$, where w_j denotes the vertex corresponding to the subset $V_j \in \mathcal{V}$. Note that we have $s(f(x)) \in \tau(x)$ by construction of s . Similarly let $\sigma(x) = \text{conv}\{v_i \in I \mid x \in U_i\}$ and note that

$r(x) \in \sigma(x)$ by construction of r . By assumption on the map g , $x \in U_i$ implies $f(x) \in V_{g(i)}$. Equivalently, if v_i is a vertex of $\sigma(x)$, then $h(v_i) = w_{g(i)}$ is a vertex of $\tau(x)$. This implies that $h(r(x)) \in \tau(x)$. Hence, $s \circ f \simeq h \circ r$ by a straight-line homotopy between $s(f(x))$ and $h(r(x))$ within $\tau(x)$. \square

We note that the commutativity up to homotopy of the diagram (21) does not require the covers of \mathbb{X} and \mathbb{Y} to be good. See also Chazal and Oudot (2008, Lemma 3.4) and Ferry et al. (2014, Proposition 4.2) for related statements about the functoriality of the nerve of a cover.

4.3 Inference

We now relate the eigenspace $E^t(f)$ of the self-map f with a generalized eigenspace obtained from the sample g . The value of this comparison derives from the assumption that f remains unknown, beyond g , so its eigenspace can be approached only indirectly, through the properties of g . We begin by recalling the assumptions:

- $f: \mathbb{M} \rightarrow \mathbb{M}$ is a continuous self-map with Lipschitz constant λ ;
- $g: X \rightarrow X$ is a finite sample of f with approximation constant ρ ;
- the Hausdorff distance between X and \mathbb{M} is $\delta = d_M(X, \mathbb{M})$.

Note that this implies $\|g(x) - f(y)\| \leq \rho + \lambda\|x - y\|$ since the left-hand side is at most $\|g(x) - f(x)\| + \|f(x) - f(y)\|$. Setting $\eta = \rho + \lambda\delta$, we note that

$$f(B_\delta(x)) \subseteq B_\eta(g(x)) \tag{22}$$

for all $x \in X$. Hence g defines a simplicial map from $\check{C}ech_\delta(X)$ to $\check{C}ech_\eta(X)$, and we get two maps in homology,

$$\gamma, J: H(\check{C}ech_\delta(X)) \rightarrow H(\check{C}ech_\eta(X)), \tag{23}$$

in which γ is induced by g and J is induced by inclusion.

We now consider the generalized eigenspace of the two maps for an eigenvalue t :

$$E^t(\gamma, J) = \ker(\gamma - t \cdot J) / (\ker \gamma \cap \ker J), \tag{24}$$

noting that this is a special case of the setting considered in Sect. 4.1. We show that under appropriate conditions this generalized eigenspace is isomorphic to $E^t(f)$. We need some definitions to prepare the first step. Recall that $B_\delta(x)$ is the closed ball with radius δ centered at $x \in \mathbb{R}^n$. For $\mathbb{M} \subseteq \mathbb{R}^n$, we call $\mathbb{M}_\delta = \bigcup_{x \in \mathbb{M}} B_\delta(x)$ the δ -neighborhood of \mathbb{M} . By the Kirszbraun Extension Property (Kirszbraun 1934; Wells and Williams 1975), $f: \mathbb{M} \rightarrow \mathbb{M}$ extends to a map $f_\delta: \mathbb{M}_\delta \rightarrow \mathbb{M}_\delta$ with the same Lipschitz constant. Similarly, f extends to a map $f_\theta: \mathbb{M}_\theta \rightarrow \mathbb{M}_\theta$, again with the same Lipschitz constant, in which $\theta = \max(2\delta, \eta)$, with $\eta = \rho + \lambda\delta$ as before. The following diagram organizes the homology groups of the spaces relevant to our argument. Apart

from f_* , f_{δ_*} , and f_{θ_*} , any map in the diagram is induced by inclusion.

$$\begin{array}{ccccc}
 & & \mathbb{H}(X_\delta) & \xrightarrow{\iota} & \mathbb{H}(X_\theta) & & \\
 & \nearrow & & \searrow a & \nearrow b & \searrow & \\
 \mathbb{H}(\mathbb{M}) & \longrightarrow & \mathbb{H}(\mathbb{M}_\delta) & \longrightarrow & \mathbb{H}(\mathbb{M}_\theta) & & \\
 \downarrow f_* & & \downarrow f_{\delta_*} & & \downarrow f_{\theta_*} & & (25) \\
 \mathbb{H}(\mathbb{M}) & \longrightarrow & \mathbb{H}(\mathbb{M}_\delta) & \longrightarrow & \mathbb{H}(\mathbb{M}_\theta) & & \\
 & \searrow & \nearrow a & \searrow b & \nearrow & & \\
 & & \mathbb{H}(X_\delta) & \xrightarrow{\iota} & \mathbb{H}(X_\theta) & &
 \end{array}$$

Consider $\iota: \mathbb{H}(X_\delta) \rightarrow \mathbb{H}(X_\theta)$, let $\iota = b \circ a$ with $a: \mathbb{H}(X_\delta) \rightarrow \mathbb{H}(\mathbb{M}_\delta)$ and $b: \mathbb{H}(\mathbb{M}_\delta) \rightarrow \mathbb{H}(X_\theta)$, and define $\phi = b \circ f_{\delta_*} \circ a: \mathbb{H}(X_\delta) \rightarrow \mathbb{H}(X_\theta)$. To compare ϕ with ι , we consider their eigenspace,

$$E^t(\phi, \iota) = \ker(\phi - t \cdot \iota) / (\ker \phi \cap \ker \iota). \tag{26}$$

We claim that this eigenspace is isomorphic to the one considered in (24).

Lemma 4 $E^t(\phi, \iota) \cong E^t(\gamma, J)$.

Proof By finiteness of X , there is $\varepsilon > 0$ such that the inclusion of X_δ in the interior of $X_{\delta+\varepsilon}$ is a homotopy equivalence and $\check{C}ech_\delta(X)$ is isomorphic to the nerve of the cover of $X_{\delta+\varepsilon}$ by open balls of radius $\delta + \varepsilon$. We can thus apply (20) and get two commutative diagrams via Lemma 3:

$$\begin{array}{ccc}
 \mathbb{H}(X_\delta) \xrightarrow{\phi} \mathbb{H}(X_\theta) & & \mathbb{H}(X_\delta) \xrightarrow{\iota} \mathbb{H}(X_\theta) \\
 \cong \downarrow & & \cong \downarrow \\
 \mathbb{H}(\check{C}ech_\delta(X)) \xrightarrow{\gamma} \mathbb{H}(\check{C}ech_\theta(X)) & & \mathbb{H}(\check{C}ech_\delta(X)) \xrightarrow{J} \mathbb{H}(\check{C}ech_\theta(X))
 \end{array} \tag{27}$$

The diagrams imply $\phi \cong \gamma$ and $\iota \cong J$, so the eigenspaces are also isomorphic, as claimed. \square

For the second step, we add two assumptions: that the map $\mathbb{H}(\mathbb{M}) \rightarrow \mathbb{H}(\mathbb{M}_\delta)$ induced by inclusion is an isomorphism, and that the induced map $\mathbb{H}(\mathbb{M}_\delta) \rightarrow \mathbb{H}(\mathbb{M}_\theta)$ is a monomorphism. This implies that a is surjective and that b is injective; see (25). We claim that under the combined assumptions, the eigenspace of $f: \mathbb{M} \rightarrow \mathbb{M}$ for $t \in \mathbb{F}$ is isomorphic to the eigenspace considered in Lemma 4.

Lemma 5 $E^t(f) \cong E^t(\phi, \iota)$.

Proof We have $\ker a \subseteq \ker \phi$ simply because $\phi = b \circ f_{\delta_*} \circ a$, and we have $\ker a = \ker \iota$ because $\iota = b \circ a$ with b injective. This implies $\ker \phi \cap \ker \iota = \ker a$. Hence,

$$E^t(\phi, \iota) = \ker(\phi - t \cdot \iota) / (\ker \phi \cap \ker \iota) \tag{28}$$

$$= \ker(b \circ f_{\delta_*} \circ a - t \cdot b \circ a) / \ker a \tag{29}$$

$$\cong \ker(b \circ f_{\delta_*} - t \cdot b). \tag{30}$$

Since b is injective, the kernel in (30) is isomorphic to $E^t(f_{\delta_*})$. This concludes the proof since $H(\mathbb{M}) \cong H(\mathbb{M}_\delta)$, by assumption, and therefore $E^t(f_{\delta_*}) \cong E^t(f)$. \square

Summarizing Lemmas 4 and 5, we obtain a sampling theorem for inferring the eigenspace of the given self-map from the sampled eigenspace module (18).

Theorem 1 *Let $f : \mathbb{M} \rightarrow \mathbb{M}$ be a self-map with Lipschitz constant λ , and let $g : X \rightarrow X$ be a finite sample of f with approximation error ρ and Hausdorff distance $\delta = d_H(X, \mathbb{M})$. Suppose that the inclusion $\mathbb{M} \hookrightarrow \mathbb{M}_\delta$ induces an isomorphism in homology, while the inclusion $\mathbb{M}_\delta \hookrightarrow \mathbb{M}_\theta$ for $\theta = \max(2\delta, \rho + \lambda\delta)$ induces a monomorphism. Then the dimension of the eigenspace $E^t(f)$ equals the dimension of the generalized eigenspace $E^t(\gamma, j)$.*

5 Computational experiments

In this section, we analyze the performance of our algorithm experimentally and compare the results with those reported in Edelsbrunner et al. (2015). For ease of reference, we call the algorithm in Edelsbrunner et al. (2015) the *Vietoris–Rips* or *VR-method* and the algorithm in this paper the *Delaunay–Čech* or *DC-method*. We begin with the introduction of the case-studies — self-maps on a circle and a torus — and end with statistics collected during our experiments.

5.1 Expanding circle map

The first case-study is an expanding map from the circle to itself. To add noise, we extend it to a self-map on the plane, $f : \mathbb{C} \rightarrow \mathbb{C}$ defined by $f(z) = z^2$. While traversing the circle once, the image under f travels around the circle twice. To generate the data, we randomly chose N points on the unit circle, and letting z_i be the i -th such point, we pick a point x_i from an isotropic Gaussian distribution with center z_i and width $\sigma = 0.1$. Note that while the noise from a Gaussian distribution is unbounded, for large enough N and sufficiently small σ (in dependence on N), a random sample noisy still has a high probability of satisfying the sampling conditions from Sect. 4. Write X for the set of points x_i , and let the image of x_i be the point $g(x_i) \in X$ that is closest to x_i^2 . As explained earlier, we construct the filtration of Delaunay–Čech complexes of X and compute eigenspace diagrams for all eigenvalues in a sufficiently large finite field to avoid aliasing effects. Our choice is $\mathbb{F} = \mathbb{Z}_{1009}$. Recall that the definition of the eigenspace module in Sect. 4.1 required a choice of $q \geq 0$. For our computations, we always chose the smallest admissible value.




Fig. 3 Left: The most persistent cycle in the Delaunay–Čech complex for points sampled near the unit circle. Right: The image of the cycle after following the discrete flow from the Čech complex back to a Delaunay–Čech complex. As expected, the map doubles the winding number

Drawing $N = 100$ points, we compare the DČ-method of this paper with the VR-method in Edelsbrunner et al. (2015). For eigenvalue $t = 2$, both methods give a non-empty eigenspace diagram consisting of a single point. Figure 3 illustrates the results by showing the generating cycle computed with the DČ-method on the left and its image on the right.

5.2 Torus maps

The second case-study consists of three self-maps on the torus, which we construct as a quotient of the Cartesian plane; see Fig. 4. For $i = 1, 2, 3$, the map $f_i : [0, 1)^2 \rightarrow [0, 1)^2$ sends a point $x = (x_1, x_2)^T$ to $f_i(x) = A_i x$, in which

$$A_1 = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad A_3 = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

The 1-dimensional homology group of the torus has only two generating cycles. Letting one wrap around the torus in meridian direction and the other in longitudinal direction, we see that f_1 doubles both generators, f_2 exchanges the generators, and f_3 adds them but also preserves the first generator.

Correspondingly, f_1 has two eigenvectors for the eigenvalue $t = 2$, f_2 has two distinct eigenvalues $t = 1$ and $t = -1$, and f_3 has only one eigenvector for $t = 1$. The input data for our algorithm, X , consists of 100 points uniformly chosen in $[0, 1)^2$.

To define the image of a point $x \in X$, we compute the point $A_i x$ and let the image be the nearest point $g_i(x) \in X$. The eigenspace diagrams of f_1, f_2, f_3 for selected eigenvalues are shown in the last three panels of Fig. 5.




Fig. 4 The periodic Delaunay triangulation on the left and its embedding in \mathbb{R}^3 on the right. The blue cycle wraps around the torus once in meridian and once in longitudinal direction. It represents an eigenvector of f_1 for eigenvalue $t = 2$. Its image wraps around the torus twice in meridian and twice in longitudinal direction (not shown)

5.3 Accuracy

To study how accurate the two methods are, we look at *false positives* and *false negatives*, and the persistence of the recurrent features of the underlying smooth maps.

Circle map

Repeating the circle map experiment with $N = 100$ points ten times, we show the superimposed twenty eigenspace diagrams (ten each for the two methods) in the upper left panel of Fig. 5. Points of the VR-method are marked blue while points of the DČ-method are marked red. The eigenvector for $t = 2$ is detected each time. However, the DČ-method detects the recurrence consistently earlier than the VR-method, with smaller birth and death values but also with smaller average persistence. The shift of the birth values is easy to rationalize: a cycle arises for the same radius in both filtrations, but remains without image in the VR-method until the radius is large enough to capture the image of every edge in the cycle. The shift of the death value is more difficult to explain and perhaps related to the fact that the DČ-method maps a cycle in one complex, K_r , to a later complex, K_s with $r \leq s \leq \rho + \lambda r$ in the filtration of Delaunay–Čech complexes. Monitoring r and s in 100 runs for a range of number of points, we show the average Lipschitz constant and the average ratio $\frac{s}{r}$ in Table 1.

There are no false negatives in this experiment, but we see a small number of false positives reported by the VR-method (the points in the upper right corner of the first panel in Fig. 5, all for eigenvalues $t \neq 2$). This indicates that the VR-method is more susceptible to noise than the DČ-method. To support our claim, we compute the eigenspace diagrams using the DČ-method with increased noise, and indeed find no false positives; see Fig. 6.




Fig. 5 Top left panel: the superimposed eigenspace diagrams of the expanding circle map for ten randomly chosen sets of 100 points each. The intervals are plotted as points whose coordinates are the birth and death values of the corresponding homology classes. Points for the VR-method are blue and points of the DC-method are red. The only points with non-negligible persistence belong to eigenvalue $t = 2$, and we get exactly one such point for each eigenspace diagram. Top right panel: the eigenspace diagrams of f_1 for a few eigenvalues. The most persistent classes are represented by points on the upper edge of the panel, indicating that their intervals last all the way to the last complex in the filtration. Here we see two such points, which correspond to the intrinsic 1-dimensional homology of the torus. Bottom left panel: the eigenspace diagrams of f_2 for a few eigenvalues. There are two intervals that exists during most of the filtration, one for eigenvalue $t = 1$ and the other for eigenvalue $t = -1$. They have the same birth and death and are therefore visible as two identical points on the upper edge of the panel. Bottom right panel: the eigenspace diagrams of f_3 for a few eigenvalues. There is only one significant eigenvector for $t = 1$

Table 1 The average Lipschitz constant, λ , and the average shift, $\frac{s}{r}$, for points sampling the circle map

	$N = 100$	200	300	400	500
Average λ	1.99	2.00	2.05	2.03	2.04
Average s/r	1.13	1.14	1.12	1.14	1.16
Average λ	2.65	3.54	3.98	4.22	5.42
Average s/r	1.33	1.57	1.71	1.64	1.91

Top two rows: no noise. Bottom two rows: 2-dimensional Gaussian noise with standard deviation $\sigma = 0.1$ in both directions




Fig. 6 The superimposed eigenspace diagrams computed with the $\text{D}\check{\text{C}}$ -method of the expanding circle map for randomly chosen sets of 200 points each with isotropic Gaussian noise with increasing width σ . In each run, the only non-empty eigenspace diagram is for $t = 2$, and this diagram contains exactly one point

Torus maps

The situation is similar for the three torus maps, whose eigenspace diagrams are shown in the next three panels of Fig. 5. The eigenvectors of f_1, f_2, f_3 are represented by points on the upper edges of the panels, indicating that their corresponding homology classes last until the last complex in the filtration. This is different in the VR-method because the Vietoris–Rips complex for large radii is less predictable than the Delaunay–Čech complex. In contrast to the circle map, we observe false positives also in the $\text{D}\check{\text{C}}$ -method. They show up as points with small to moderate persistence in the three diagrams. We also have false positives in the VR-method, but the results are difficult to compare because for complexity reasons we could not run the algorithm beyond $N = 200$ points. As another indication of improved accuracy of the $\text{D}\check{\text{C}}$ -method, we note that the eigenspace diagrams we observe in our experiments do not suffer the problem of abundant eigenvalues discussed in Edelsbrunner et al. (2015, Section 6.4).

5.4 Runtime analysis

We analyze the running time of the $\text{D}\check{\text{C}}$ -method for sets of N points, with N varying from 100 to 10000. For the persistent homology computation, we use coefficients in the field \mathbb{Z}_{1009} . The time is measured on a notebook class computer with 2.6GHz Intel Core i7-6600U processor and 16GB RAM.

Overall running time

We begin with a brief comparison of the two methods, first of the overall running time for computing eigenspace diagrams; see Table 2. As mentioned earlier, the VR-

Table 2 Time needed to compute the eigenspace diagram of the expanding circle map for N points sampled near the unit circle

Time (s)	$N = 100$	150	200	500	1000	1500	2000	2500
VR-method	157.41	986.60	–	–	–	–	–	–
DČ-method	0.07	0.12	0.21	0.92	3.66	8.36	14.53	22.35

For $N \geq 200$, the VR-method needs more than half an hour, at which time we terminated the process




Fig. 7 The time needed to compute the eigenspace diagram of the expanding circle map with the DČ-method as a function of the number of sampled points. We also show the amount of time spent to compute separating spheres, which is more than half the overall running time. The time for computing the Delaunay–Čech complexes and the persistence diagrams is less than 0.5 s in all cases and therefore not shown. To estimate the asymptotic behavior, we use the least squares technique to fit lines to the log-log data points; see the right panel. Excluding the results for data with less than $N = 5000$ points we get slopes 2.66 and 2.71, which suggests that the experimental running time of our algorithm is between quadratic and cubic in the input size

method uses Vietoris–Rips complexes, which grow fast with the number of points and the radius. We could therefore run this method for $N = 100$ and 150 points only, terminating the run for $N = 200$ points after half an hour.

To get a better feeling for the running time of the DČ-method, we plot the results in Fig. 7, adding curves to indicate the asymptotic experimental performance. The outcome suggests that the computational complexity of the DČ-method is between quadratic and cubic in the number of points. We note that more than half of the time is used to compute smallest separating spheres.

Flowing an edge

To gain further insight into the time needed to flow a cycle from the Čech to the Delaunay–Čech complex, we present statistics for collapsing a random edges in a variety of settings. The edges are constructed from 100, 1000, 10000 points chosen along the unit circle with added Gaussian noise, and from 100, 1000, 10000 points chosen uniformly in $[0, 1)^2$.

For each data set, we pick two points at random and monitor the effort it takes to flow this edge from the Čech complex to the Delaunay–Čech complex. Specifically, we iterate Φ on each edge individually until the result stabilizes. The statistics in

Table 3 Statistics for flowing 1000 randomly chosen edges from the Čech to the Delaunay–Čech complex

	Circle			Square		
	$N = 100$	1000	10000	100	1000	10,000
#iterations: avg	5.27	9.09	14.70	5.47	11.98	14.60
max	9.00	13.00	19.00	9.00	16.00	17.00
#tests: avg	1.23	1.17	1.21	1.60	1.32	1.20
max	8.00	5.00	4.00	15.00	16.00	5.00

Top two rows: the average and maximum number of iterations of Φ to flow an edge from the Čech to the Delaunay–Čech complex. Bottom two rows: the average and maximum number of points tested to find a set for which the separating sphere does not exist

Table 3 shows how many times Φ is iterated and how many points are tested inside each call to compute the discrete gradient. The statistics for the circle and the square are similar, with consistently larger numbers when we pick the edges in the square.

Smallest separating spheres

Our analysis shows that the $\check{D}\check{C}$ -method spends most of the time computing smallest separating spheres. For this reason, we compare the straightforward implementation (function SEPARATE), with the heuristic improvement (function MOVETOFRONT). We generate the points in $[0, 1]^2$ as described above. For both functions, we randomly pick 10000 edges from the Čech complex and another 10000 edges from the Delaunay–Čech complex, and we test for each edge whether or not there exists a sphere that separates the edge from the rest of the points. Figure 8 shows that the running time of both functions depends linearly on the number of points, which is to be expected. The best-fit linear functions suggest that the move-to-front heuristic is faster than the more naive extension of the miniball algorithm to finding smallest separating spheres. The difference is more pronounced for edges of the Čech complex (left panel) for which we expect more points inside the circumscribed spheres and an early contradiction to the existence of a separating sphere. In contrast, the difference in performance is negligible for edges sampled from the Delaunay–Čech complex, for which separating spheres exist by construction.

6 Discussion

The main contributions of this paper are the construction of a filtration-preserving chain map from a Čech filtration to the corresponding Čech–Delaunay filtration, the construction of a geometrically meaningful chain self-map map on a Delaunay triangulation from a self map on a point set, and its application to computing eigenspaces of sampled dynamical systems. Following the proof of collapsibility in Bauer and Edelsbrunner (2017), we get an efficient algorithm for the chain map though implicit treatment of the Čech complex. The reported research raises a number of questions:




Fig. 8 Left: the time needed to compute 10,000 smallest separating spheres for randomly chosen edges from the Čech complex constructed on points sampled uniformly from $[0, 1]^2$. Right: the time needed to compute 10000 smallest separating spheres for edges of the Delaunay–Čech complex constructed on points sampled uniformly from $[0, 1]^2$

- Can we give theoretical upper bounds on the number of individual collapses needed to flow a cycle to its image under the stabilization map of the Čech–Delaunay gradient flow?
- Can the computation of smallest separating spheres be further improved by customizing the procedure to small sets inside the sphere, or by taking advantage of the coherence between successive calls?

We expect that the fast chain map algorithm has applications beyond this paper, including to the transport of structural information between meshes, and to the visualization of topological information shared by related high-dimensional dataset.

Funding Open Access funding provided by Projekt DEAL.

Compliance with ethical standards

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, 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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bauer, U., Edelsbrunner, H.: The Morse theory of Čech and Delaunay complexes. *Trans. Am. Math. Soc.* **369**(5), 3741–3762 (2017)
- Chazal, F., Oudot, S.Y.: Towards persistence-based reconstruction in Euclidean spaces. In: *Computational Geometry (SCG'08)*, pp. 232–241. ACM, New York (2008)
- Edelsbrunner, H., Harer, J.: *Computational Topology: An Introduction*. American Mathematical Society, Providence (2010)
- Edelsbrunner, H., Mücke, E.P.: Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.* **9**(1), 66–104 (1990)
- Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. *Discrete Comput. Geom.* **28**(4), 511–533 (2002)
- Edelsbrunner, H., Jabłoński, G., Mrozek, M.: The persistent homology of a self-map. *Found. Comput. Math.* **15**(5), 1213–1244 (2015)
- Ferry, S., Mischakow, K., Nanda, V.: Reconstructing functions from random samples. *J. Comput. Dyn.* **1**(2), 233–248 (2014)
- Forman, R.: Morse theory for cell complexes. *Adv. Math.* **134**(1), 90–145 (1998)
- Gärtner, B.: Fast and robust smallest enclosing balls. In: *Proceedings of the 7th Annual European Symposium on Algorithms, ESA '99*, pp. 325–338. Springer, Berlin (1999)
- Gromov, M.: Monotonicity of the volume of intersection of balls. In: *Geometrical Aspects of Functional Analysis (1985/1986)*, Volume 1267 of *Lecture Notes in Mathematics*, pp. 1–4. Springer, Berlin (1987)
- Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2002)
- Kaczynski, T., Mischakow, K., Mrozek, M.: *Computational Homology*. Applied Mathematical Sciences, vol. 157. Springer, New York (2004)
- Kirszbraun, M.: Über die zusammenziehende und Lipschitzsche Transformationen. *Fundam. Math.* **22**(1), 77–108 (1934)
- Wells, J.H., Williams, L.R.: *Embeddings and Extensions in Analysis*. Springer, New York (1975). *Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 84*
- Welzl, E.: Smallest enclosing disks (balls and ellipsoids). In: *New Results and New Trends in Computer Science (Graz, 1991)*, Volume 555 of *Lecture Notes in Computer Science*, pp. 359–370. Springer, Berlin (1991)
- Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discrete Comput. Geom.* **33**(2), 249–274 (2005)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.